Transform/Subband Representations for Multidimensional Signals with Arbitrary Regions of Support

by

John G. Apostolopoulos

S.M., Massachusetts Institute of Technology (1991)
S.B., Massachusetts Institute of Technology (1989)

Submitted to the Department of Electrical Engineering and Computer Science in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 1997

Author ______________________________
Department of Electrical Engineering and Computer Science
May 27, 1997

Certified by ______________________________
Jae S. Lim
Professor of Electrical Engineering
Thesis Supervisor

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

Transform/Subband Representations for Multidimensional Signals with Arbitrary Regions of Support

by
John G. Apostolopoulos

Submitted to the Department of Electrical Engineering and Computer Science
on May 27, 1997, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering

Abstract

Transform/subband representations are a basic building block for many signal processing algorithms and applications. Until recently, most of the research in the signal processing community has focused on developing representations for infinite-length signals, with simple extensions to finite-length 1-D and rectangular-support 2-D signals. However, many applications today and in the near future will entail efficient processing of multidimensional signals defined over arbitrary (non-rectangular) regions of support. One increasingly important class of signals in which this is evident is the class of arbitrarily shaped objects or regions in an image or video scene. Conventional approaches developed for 1-D problems are inappropriate for 2-D or M-D, as they are either computationally impractical or do not achieve adequate performance.

We present a novel framework for creating critically sampled, perfect reconstruction transform/subband representations for discrete 1-D, 2-D, and general M-D signals with arbitrary regions of support. Our method determines a basis for a given signal by selecting an appropriate subset of vectors from a basis defined over a conveniently chosen superset space. In this manner, the resulting representation can inherit some of the important properties of the basis defined over the superset space. In particular, we develop and explore two novel wavelet-type representations for 1-D, 2-D, and M-D signals with arbitrary supports. These representations preserve many of the important properties associated with wavelet representations of infinite-length signals. Our work appears promising for a number of applications in 1-D, 2-D, and M-D signal processing problems. We briefly explore one such direction, namely, its application for 2-D object/region-based compression.

Thesis Supervisor: Jae S. Lim
Title: Professor of Electrical Engineering

Acknowledgements

During my graduate study at MIT I have had the privilege to interact with a select group of people who have contributed to my personal and professional growth. Professor Jae Lim took me in his research group when I was a beginning graduate student. Working in that group has had a profound effect on me. He guided me through my early years and through the HDTV standardization process. In recent years his faith allowed me significant freedom in identifying and attempting to solve important problems. This freedom has enabled me to grow intellectually. For these and many other reasons I am extremely grateful.

Professor Gilbert Strang has influenced me both directly through our discussions and his insightful comments, and indirectly through the many wonderful books that he has written. Professor George Verghese graciously joined my committee even though the topic was somewhat far from his principal research areas, and in the brief time that we interacted he provided a number of helpful suggestions. Professors William Schreiber and Peter Hagelstein have both been very generous in sharing their knowledge and advice over the years.

MIT is filled with many wonderful people who I have had the good fortune to work with as colleagues, and more importantly, to get to know as friends. The members of the Advanced Telecommunications and Signal Processing (ATSP) Group and the Digital Signal Processing Group (DSPG) have helped make my graduate school days a very enjoyable and rewarding experience. There are a number of people who I would like to give a special thanks. Peter Monta was very influential and extremely helpful when I first joined the group. His brilliance, enthusiasm, and modesty have been inspiring. Shiufun Cheung is a good friend who was always available to share his broad knowledge in academic and aesthetic milieus. The generosity and camaraderie of Haralabos (Babis) Papadopoulos have been high points of recent years. Mike Polley and Rajni Aggarwal have been great friends since my undergraduate days. I will always have many fond memories of our times together, from debating the feasibility of some new technology, to enjoying home brew and fine Texan/Indian cuisine. One can't ask for better friends.

A special thanks goes to Cindy LeBlanc for making my life much easier during my graduate years. She makes ATSP a pleasant place to work and keeps the group running smoothly.

I would like to thank the MIT Advanced Television Research Program, and the AT&T Foundation for an AT&T Bell Laboratories PhD Fellowship that supported my graduate education.

I would like to thank a very dear friend, Susie Wee. Her kindness, support, laughter, understanding, and patience have been an integral part of my life over these years.

Finally, my parents, Professor George and Maria Apostolopoulos, instilled in me an appreciation for the essentials of life, education, and creative thought. Their love, support, and encouragement have enabled me to become what I am today. It is to them that I dedicate this thesis. In addition, I would like to thank my younger brother, Dr. Nicholas Apostolopoulos, for his constant support, dynamic advice, story-telling abilities, and updates on da Bulls and da Bears.

Dedication

To my parents,

George and Maria Apostolopoulos.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Transform/subband representations form an essential element of many signal processing algorithms and applications – including restoration, enhancement, compression, and analysis [1, 2, 3, 4]. These representations are typically designed for signals with convenient regions of support (ROS), such as infinite-extent 1-D or 2-D signals or finite-length 1-D or rectangular support 2-D signals. However, many important applications involve signals that do not have convenient supports. For example, most high-level representations of images or video incorporate 2-D or 3-D models which decompose the scene into arbitrarily shaped objects or regions. Medical imaging often results in 2-D or 3-D imagery where the relevant information is localized over an arbitrarily shaped region. Furthermore, many areas of scientific research involve signals which are defined over complex domains. High quality processing of these signals with arbitrary ROS's is very important for the success of these applications.

In this thesis we consider the problem of creating efficient transform/subband representations for discrete 1-D, 2-D, and general M-D signals defined over arbitrary regions of support (A-ROS). An important example of the signals that we would like to represent are the objects or regions within an image or video scene. Throughout this work, we assume that the ROS of the signal is given, and the goal is to represent the signal's amplitude over the ROS. We wish to create a representation for signals with arbitrary supports that achieves a similar high level of performance and low level of complexity as is achieved by conventional representations for signals with convenient supports. High performance means that the representation should accurately represent typical signals using only a small fraction of the total number of coefficients. Low complexity is of prime importance since even a small M-D signal can make most algorithms impractical, e.g. a 3-D signal of only 30-sample "diameter" may require processing up to 27,000 samples.

This introduction begins by briefly describing the importance of signals with arbitrary ROS's. A number of comments are then made on the conventional approaches used for processing signals with simple and convenient supports. Finally, an overview of the remainder of the thesis is given.

1.1 Importance of Signals with Arbitrary Regions of Support

Signals with A-ROS's arise in many important theoretical areas and practical applications. Practical systems (implementations) process a finite portion of a signal at each step. Analytical formulations are often simpler to develop in the context of infinite-length signals, however the desired processing is performed on finite-length signals. For instance, in 1-D, the signal may be of finite (possibly arbitrary) length, or may be of infinite length but may be acquired or segmented into finite portions for appropriate processing.

Many of the issues and problems that arise when processing signals with A-ROS are much clearer in the context of multidimensional signals. We will briefly consider three general areas which involve processing signals with A-ROS.

Image and video processing An image is a 2-D signal with square or rectangular ROS. Video is a 3-D signal with rectangular spatial support and possibly infinite temporal support. Most conventional image/video processing algorithms, such as compression, involve block-based or overlapped-block-based schemes. These algorithms partition the image into blocks (possibly overlapping) and each block is processed independently of the others. All conventional Block-DCT, lapped transform, and wavelet image/video compression algorithms can be viewed as block-based or overlapped block-based schemes. Block-based processing achieves good performance while allowing architectural simplicity. However, these block-based schemes do not exploit (and in fact neglect) the actual content of the image or video. In effect, these approaches implicitly assume a source model of moving square blocks. However, a typical video scene is not composed of moving square blocks. Therefore, block-based schemes impose an artificial structure on the video signal and then try to encode this structure, as opposed to recognizing the structure inherent to a particular video scene and attempting to exploit it.

An improved source model may be developed by identifying and efficiently representing the structure that exists within a video scene. For example, since real scenes contain objects, a promising source model is two or three dimensional moving objects. This approach may provide a much closer match to the structure in a video scene than the aforementioned block-based schemes. This source model is illustrated in Figure 1.1.

An improved source model of this form leads to a more natural signal representation which facilitates higher-level processing. For example, the content in the scene may be identified [1], and

[1]The objects may be identified in a sophisticated manner, such as "this is a car" and "this is a person", or a relatively simple manner, such as "these objects are moving or active" while "these are static".

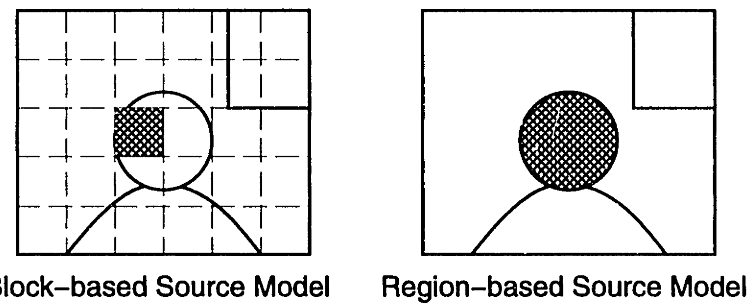

Figure 1.1: Block-based versus object- or region-based source models.

the scene may be processed based on its content, providing a form of *content-based processing*. This content-based image/video representation may facilitate novel algorithms and applications. The individual objects/regions in an image or video may be individually extracted and processed, with different processing applied to each. For example, this representation may lead to potentially improved compression, and also enables direct access and manipulation/processing of the objects from the compressed bitstream – providing a simple and natural object-based scalability (at the bitstream level).

Object- or region-based image/video processing appears very promising, however two fundamental issues must be addressed before it will be successful: (1) the video scene must be appropriately segmented into objects, and (2) each of the objects must be processed appropriately. The type of processing would depend on the particular application. For example, high quality compression [5, 6, 7] requires efficient coding of the segmentation information and of the object interiors, i.e. the shape (ROS) and amplitude of the signal. A high performance and computationally efficient transform/subband approach for representing signals with A-ROS's would greatly assist the coding of the object interiors. Furthermore, a representation of this form may facilitate many other types of processing for the individual objects, e.g. restoration, extrapolation, texture analysis, etc.

An example of a 2-D region extracted from an image is shown in Figure 1.2. This example corresponds to the simple case where the ROS corresponds to a single, contiguous arbitrarily shaped region. More complex supports may arise in a number of situations. For example, partial occlusion may result in the object's support being composed of a number of separate regions.

A video signal can be decomposed into 3-D signals in a number of different ways. Conceptually the simplest approach is to identify the various 3-D objects in the scene. Alternatively, the projection of a 3-D object onto the 2-D image plane can be identified, and its evolution can be tracked with time. The resulting arbitrarily shaped 3-D signal can be thought of as a 2-D signal that changes with time, specifically both the intensity within the region and the region shape can change with time.

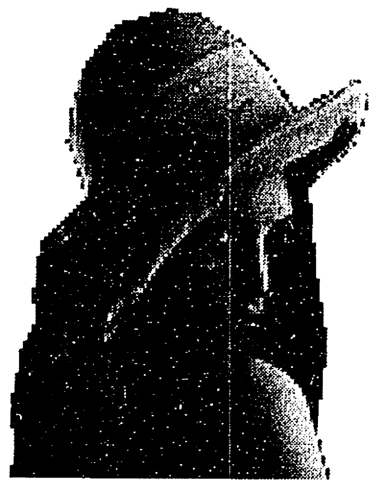

Figure 1.2: An arbitrarily shaped 2-D signal taken from an image.

Medical imaging Many areas of medical imaging involve the acquisition and processing of 2-D or 3-D arbitrarily shaped objects. For example, in computer tomography a number of line projections (typically from an X-ray source) are used to construct an image of the interior of the human body. Magnetic resonance imaging (MRI) is another approach that enables 3-D imaging of the human body. In both cases, the 2-D and 3-D reconstructed images correspond to signals localized over an arbitrarily shaped region in space. For example, the human head is localized in 3-D space, and one 2-D slice showing the interior of the head is shown in Figure 1.3. An appropriate transform/subband representation for signals with arbitrarily shaped ROS's may be beneficial for further processing of the reconstructed 3-D imagery. Also note that the evolution of a 3-D signal over time (e.g. the beating of a heart) corresponds to a 4-D signal with arbitrarily shaped support, and an appropriate representation may also be useful for processing these signals.

Problems defined over arbitrarily shaped domains Many problems in engineering and applied science involve solving problems over an arbitrarily shaped domain. Many of these problems require the solution of a partial differential equation (PDE) subject to boundary conditions over a given domain. For example, a basic problem in electrostatics is to solve Laplace's or Poisson's equation over a given domain with some boundary conditions. Similar problems exist throughout mechanics, fluid dynamics, and many other areas. Their analysis provides information on if

Figure 1.3: A 2-D signal with an arbitrarily shaped ROS that corresponds to one MRI slice of a human brain. A number of slices stacked on top of one another provide a 3-D signal with arbitrarily shaped ROS that describes the interior of a human brain.

an engine will work, if an airplane will fly, and a prediction of the weather forecast. These phenomena are typically modeled by sophisticated partial differential equations, with complicated connectivities and media characteristics, and over sometimes complex domains. Closed form solutions for these problems can only be found in extremely simple cases. In general, one is limited to numerically approximating the solution.

The various numerical approaches for attacking these problems correspond to different methods for discretizing the continuous problems [8]. They convert the infinite-dimensional continuous problem into a finite-dimensional problem which can be processed by a computer. Two classic approaches are finite difference and finite element methods. In finite differences the differential operators are discretized. In the finite element method the solution space is discretized, i.e. a solution is determined over a (finite-dimensional) subspace of the actual (infinite-dimensional) solution space. The finite element method is especially useful as it is easily adapted to complex domains. A finite element is a piecewise polynomial defined over a local region of the domain. The finite element method begins by defining finite elements over each local region of the domain. The true solution is then approximated as a linear combination of the finite elements, e.g. the approximate solution is equal to the projection of the true solution onto the subspace formed by the span of the finite elements. The finite element method is therefore similar to transform/subband representations in that both attempt to represent a signal as a linear combination of a set of vectors.

Wavelets are currently being studied as another finite-element-type method for attacking these problems (e.g., see [9, 10]). Typically, the piecewise polynomial finite elements are replaced by scaling functions (wavelet lowpass functions) and a multiresolution approach is employed to solve the problem (analogous to the classical multigrid methods [11]). Currently it is an open question as to whether wavelets will provide improved performance over the traditional finite element methods. One difficulty is that the wavelet approach can not adapt to the complicated problem geometries as easily as the finite element method. The ability to create wavelet representations for signals with A-ROS may increase the applicability of wavelet-based methods for certain classes of problems.

The three general areas described above show that signals defined over arbitrary regions of support arise in a number of very important problems. Considerable research has been focused on some of these areas over the years. For instance, the numerical solution of partial differential equations over (somewhat) arbitrary domains correspond to a class of problems that have been crucial to various industries for many decades. Many high performance, very sophisticated approaches have been developed for this class of problems.

Our primary motivation for this work is to represent the objects and regions in image and video. This problem has received little attention until relatively recently, and poses a number of distinct differences from the PDE-based problems. Foremost are the performance criteria: for PDE-based problems the solution typically must very accurately model the real world, otherwise there may be unacceptable consequences; meanwhile for image/video, the criteria depends on the application, but in compression the goal is to approximate the signal (with a given bit budget) so that it is visually pleasing – any inaccuracies, while a nuisance, are not catastrophic. To achieve the high performance required for the PDE problems mandates the use of very sophisticated and highly complex processing. Image and video applications are typically constrained to low or medium complexity approaches. High performance approaches for image and video exist in principle, but their complexity severely limits their usefulness. Therefore, while our work may potentially be useful in other areas, our primary motivation is for high-quality, low-complexity processing of image and video signals.

1.2 Remarks on Conventional Processing of Signals with Convenient Supports

The primary goal in many early signal and image processing applications was to appropriately process the interior of the signal. The theory was developed and algorithms were designed for infinite-length signals, and simple methods for tackling the boundary problems of finite-length signals were developed and appeared to be adequate. Furthermore, for block-based processing such as the Block-DCT [1], the problem to a large extent could be avoided. Only for filtering (convolution) or lapped transforms did the boundary issues arise, but even in these cases the effects were typically quite small. High-quality boundary processing for 1-D signals has recently received considerable attention, and this will be reviewed in Section 2.2. As previously discussed, the 2-D/M-D problems being examined are becoming more complex requiring the creation of novel approaches for their solution. In the following paragraphs we will briefly consider how the problems being addressed have evolved, and how this has impacted the solution methods.

In many applications, the signal ROS can be explicitly chosen to reduce or eliminate the boundary problems. For instance, 1-D signals can be chosen to have a convenient length, e.g. 256, 512, or 1024 samples. This is especially evident in image processing, where the majority of the experiments published in the literature are performed on images of size 256×256 or 512×512 pixels. Furthermore, many image or video compression algorithms have restrictions that the source material have rectangular support with horizontal and vertical dimensions which are integral multiples of 8 or 16 pixels. These signal restrictions alleviate the boundary problems for block-based transform/subband schemes, such as the Block-DCT.

In some applications it may be desirable to use basis functions that are longer than the subsampling factor (i.e. that overlap neighboring regions), such as lapped transforms or general subband filtering schemes. With these transform/subband approaches, the boundary problem still exists, even with the proposed size restrictions. However, the ramifications in typical applications may be minor. When processing a 1-D signal the problem arises only at the boundaries, and the longer the signal the smaller the boundary effect. An important example is audio compression, where the boundary problem only affects a few milliseconds at the beginning and end of the audio segment. For a typical image processing scenario, the affected boundary area lies along the sides of the image. Typically only a small percentage of the total number of pixels are affected by the boundary processing, and also any difficulties appear at the sides of the pictures – far from the crucial center viewing area of the image. In both of these cases, any inefficiency or inadequacy of the boundary processing only affects a small percentage of the total number of samples, thereby limiting the ill-effects on the algorithm. In addition, harmful effects (potential distortions) are "hidden" to a certain extent, since they are localized in perceptually less-crucial areas of the signal.

Arbitrarily shaped objects/regions in an image or video do not possess these appealing properties. Their regions of support typically are not convenient in size or shape, and in most cases the ROS would be completely out of our control. In addition, the region may have a shape in which all of the samples in the region are affected by the processing at the boundary; therefore any improper boundary processing can have a drastic effect. Furthermore, the object (and its boundary) may be placed in the middle of the image, prominently displaying any potential distortion.

1.3 Outline of the Thesis

The goal of this thesis is to create efficient transform/subband representations for discrete 1-D, 2-D, and general M-D signals defined over arbitrary regions of support. Specifically, we wish to create a representation that achieves a comparable high level of performance and low level of complexity for signals with arbitrary supports, as is achieved by conventional representations of signals with convenient supports. Most of the signals that we are interested in have arbitrarily shaped supports, or some extension thereof. However, as we will show, a representation that is applicable to any arbitrary ROS greatly simplifies a number of issues, and may potentially lead to applications in other areas.

Chapter 2 begins by presenting a brief overview of transform/subband representations. In particular we briefly consider the case of a 1-D two-channel filterbank, which is an important building block within this thesis. Chapter 2 continues by reviewing the previous research in representing arbitrary-length 1-D signals and 2-D signals with arbitrarily shaped supports. Since there has been significant research on representing these types of signals, we provide a brief overview of the various proposed approaches and attempt to unify and classify them.

In Chapter 3, we introduce a general approach for creating critically sampled, perfect reconstruction transform/subband representations for signals with arbitrary regions of support. This general approach determines a basis for a given signal by selecting an appropriate subset of vectors from a basis defined over a convenient superset space. This general approach may be based on any linear transform/subband representation (DFT, DCT, wavelet, etc.) and provides a significant amount of freedom in its creation.

In Chapter 4, we focus on creating a wavelet (multiresolution) representation based on the general approach presented in Chapter 3. A wavelet-type representation appears to be a natural approach for representing signals with arbitrarily shaped supports, and is especially promising for the arbitrarily shaped objects or regions in an image or video. In particular, we propose two different wavelet-based representations. The first representation is designed to fully preserve the wavelet's important polynomial accuracy property. The resulting representation provides some interesting theoretical features relative to previous approaches, however fully preserving polynomial accuracy leads to two disadvantages that may limit its practical usefulness.

The second wavelet representation is designed to represent signals with completely arbitrary supports. This representation provides a number of compelling properties. Foremost, for appropriate filters, it provides a basis over any possible arbitrary support, i.e. the ROS can have any shape, can contain holes, can be disconnected, and can include isolated samples. While it may not be beneficial to apply this representation for all possible arbitrary supports, the fact that there are no constraints can greatly simplify a number of applications. In addition, for a simple alignment of the filters, the selection of the vectors to provide the basis is trivially related to the signal's support – thereby providing an extremely simple approach for determining a basis. Furthermore, for typical signals, the resulting basis preserves many of the general characteristics provided by wavelet representations of signals with convenient supports. Overall, both theoretical and empirical analysis suggests that this second representation can provide a high-quality basis for representing signals with arbitrary supports.

Chapter 5 presents a preliminary investigation of the application of the proposed representation to image and video compression. We begin with an overview of conventional image/video compression, and then discuss the motivation for object/region-based compression. The problem of coding the amplitude ("interior" or "texture") of the arbitrarily shaped objects/regions in an image or video scene is then discussed, and a preliminary comparison between our second wavelet representation and other possible approaches is given.

Finally, Chapter 6 summarizes the principal contributions of this thesis, and suggests some possible directions for future study.

Chapter 2

Background and Previous Research

This chapter provides a brief review of conventional transform/subband representations for 1-D arbitrary-length signals and 2-D signals with arbitrarily shaped supports. Since there has been significant research on representing these types of signals, we provide a brief overview of the various proposed approaches and attempt to unify them in a common framework. In particular, we examine many of the approaches from a vector space framework which illustrates many of their fundamental similarities and differences. Currently there is a large and growing literature on processing signals with arbitrarily shaped supports; we primarily focus on those approaches that are most relevant to the problem at hand, and mention the popular approaches in current applications.

This chapter begins with a very brief overview of transform/subband representations in order to introduce our perspective and our nomenclature. We then consider the familiar cases of arbitrary-length 1-D signals and 2-D or M-D signals with square or rectangular supports. The basic problem of 2-D or M-D signals with arbitrarily shaped supports is then considered. These problems may involve bounded (finite) supports, or supports that are bounded in only one direction, e.g. a causal or anticausal 1-D signal, or 3-D signals such as video or the solution of a time-dependent partial differential equation, where the signal is bounded spatially, and evolves in time. Throughout this chapter, the signal's support is assumed to be composed of a single contiguous region. This chapter concludes with a brief summary highlighting some of the advantages and disadvantages of the various approaches.

2.1 Overview of Transform/Subband Approaches

This section provides a very brief overview of transform/subband representations. We limit our discussion since there exists a number of excellent texts on the topic [2, 3, 4]. In particular, we briefly examine the 1-D two-channel filterbank which is an important building block within this thesis. We consider the two-channel filterbank from both transform and subband filtering perspectives, and in so doing introduce some of the important concepts and terminology that is used throughout the remainder of the thesis.

Transform and subband approaches express a signal as a linear combination of a set of vectors. The transform and subband signal representations are highly related. For instance, under certain conditions the two representations contain the same information, simply arranged differently. For example the discrete cosine transform (DCT) can be computed like a transform as the inner product between the DCT basis vectors and the signal, or like a subband filterbank where time-reversed DCT basis vectors are convolved with the signal and the filtered outputs are downsampled. In this case, the transform and subband approaches yield the same information, simply arranged differently.

The transform and subband approaches also have some differences. The transform approach stresses the linear expansion or linear algebra (vector space) viewpoint, while the subband approach stresses the filtering or signal processing viewpoint. Both of these viewpoints provide valuable perspectives on transform/subband signal processing. Strictly speaking, the optimal transform scheme (the Karhunen-Loève transform) differs from the optimal subband filterbank scheme (ideal bandpass filterbank), since they are defined with respect to different criteria. The different algorithms typically used to compute each representation and the different optimal solutions, has caused the transform and subband approaches to be traditionally classified as separate entities. However, since both of these approaches express the signal in an equivalent manner (as a linear expansion of a set of vectors), we choose to collectively refer to these approaches as transform/subband representations. Along these lines, we will readily alternate between the transform and subband perspectives when appropriate and use the terms "filter" and "basis vector" accordingly.

2.1.1 Two-channel Filterbank

A conventional 1-D two-channel filterbank is shown in Figure 2.1. The original signal is filtered by a lowpass analysis filter $h_0(n)$ and a highpass analysis filter $h_1(n)$. The output of each filter is downsampled by a factor of two. The resulting outputs, $c_0(k)$ and $c_1(k)$, are referred to as the lowpass and highpass coefficients, or the lowpass and highpass subbands, respectively. The original signal may be reconstructed by upsampling the lowpass and highpass coefficients and sending them through appropriate lowpass and highpass synthesis filters, $g_0(n)$ and $g_1(n)$ respectively.

If the four filters are designed appropriately, the output of the two-channel filterbank is equal to the input (up to a simple delay and scaling), i.e. the filterbank provides *perfect reconstruction*. That is, we can decompose the signal into its lowpass and highpass subbands and then perfectly recover

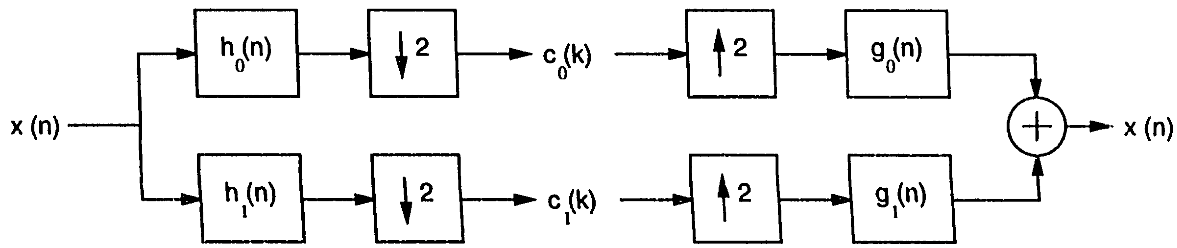

Figure 2.1: A 1-D two-channel filterbank.

the signal from these subbands. This subband representation is *critically sampled,* i.e. the total number of coefficients is equal to the total number of samples in the original signal. Specifically, the signal is decomposed into two coefficient streams, each running at half the sample rate of the original signal. This conserves the total sample rate between the signal and the coefficients.

When we consider the two-channel filterbank from a subband filtering point-of-view, the signal is analyzed by filtering and downsampling. The signal is then synthesized (reconstructed) by upsampling and filtering. The elements that comprise the two-channel filterbank are therefore the two analysis filters, $h_0(n)$ and $h_1(n)$, the two synthesis filters, $g_0(n)$ and $g_1(n)$, and the filter/downsample and upsample/filter operations.

The two-channel filterbank can also be examined from a transform point-of-view. In a conventional transform, the signal is analyzed (forward transform) by computing inner products between the signal and a set of analysis basis vectors. The signal is synthesized or reconstructed (inverse transform) by forming a linear combination of a set of synthesis basis vectors. In the case of an orthogonal transform, the analysis and synthesis basis vectors are identical, and for a biorthogonal transform they are duals of each other.

Using these ideas, the two-channel filterbank can also be examined from a transform point-of-view. Specifically, the signal is analyzed by computing inner products with a set of analysis basis vectors, where the analysis basis vectors are given by all even translates of the time-reversed lowpass and highpass analysis filters.

$$c_0(k) = < x(n), h_0(2k-n) > \quad (2.1)$$

$$c_1(k) = < x(n), h_1(2k-n) > \quad (2.2)$$

The signal is synthesized by a linear combination of a set of synthesis basis vectors, where the synthesis basis vectors are given by all even translates of the lowpass and highpass synthesis

Subband Filtering Point-of-view	Transform Point-of-view
Analysis filters $\{h_0(n), h_1(n)\}$	Analysis basis vectors $\{h_0(2k-n), h_1(2k-n)\}_{k\in J}$
Filter/downsample	Inner products
Synthesis filters $\{g_0(n), g_1(n)\}$	Synthesis basis vectors $\{g_0(n-2k), g_1(n-2k)\}_{k\in J}$
Upsample/filter	Linear combination

Table 2.1: The relationship between the transform and subband filtering points-of-view for analyzing the two-channel filterbank.

filters.

$$x(n) = \sum_k c_0(k) g_0(n-2k) + \sum_k c_1(k) g_1(n-2k) \tag{2.3}$$

The lowpass and highpass coefficients are given by the inner products between the signal and the lowpass and highpass analysis vectors, and these coefficients correspond to the weights for the linear combination of the lowpass and highpass synthesis vectors for reconstructing the signal.

The two-channel filterbank (as any transform/subband representation) can be viewed from both the transform and subband points-of-view. Each perspective provides many valuable insights. We can view the analysis as either lowpass/highpass filtering followed by downsampling, or as inner products with a set of analysis basis vectors. Similarly, the synthesis can be viewed as either upsampling followed by lowpass/highpass filtering, or as a linear combination of a set of synthesis basis vectors. These relationships are summarized in Table 2.1. We will readily switch back and forth between these perspectives throughout this thesis, referring to either filters or basis vectors as is most appropriate at each given instance.

Wavelet Representations The lowpass and highpass subbands that result from a two-channel filterbank can also be further decomposed by applying another two-channel filterbank to each of the subbands. Both the lowpass and highpass subbands can be further decomposed, or only one of them. An important case in signal processing is when a two-channel filterbank is recursively applied to only the lowpass subband. This is illustrated for a three-level decomposition in Figure 2.2. This type of filterbank is often referred to as a logarithmic filter bank (since the subbands have equal bandwidths on a logarithmic scale) or an octave-band filterbank (since the input bandwidth has been divided into octaves). This decomposition is also often referred to as a *discrete-time wavelet transform*, since there is an intimate connection between it and continuous-time wavelets. Recursively applying the two-channel filterbank to any of the subbands results in the creation of arbitrary binary-tree-structured filterbanks, referred to as *discrete-time wavelet packets.*

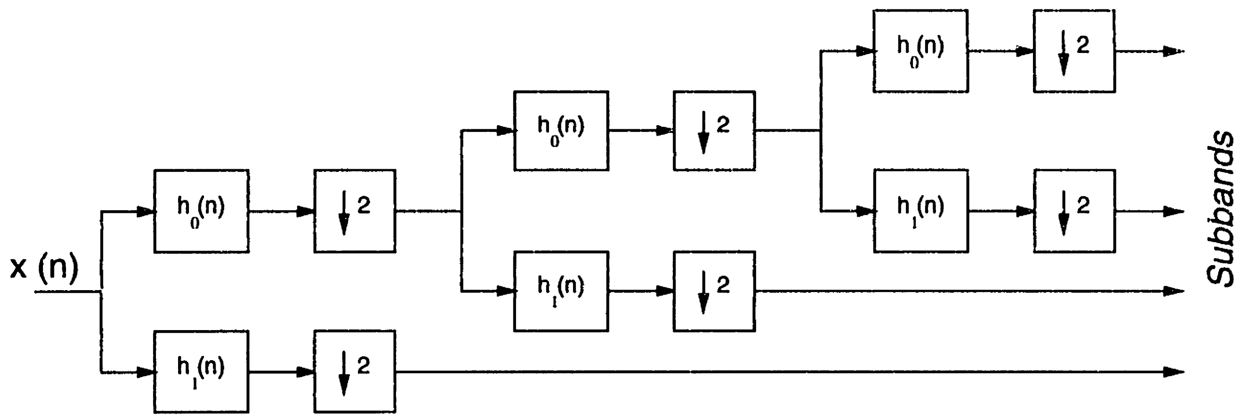

Figure 2.2: A three-level discrete-time wavelet decomposition created by iterating the two-channel filterbank on the lowpass subband.

Discrete-time filters are intimately related to continuous-time scaling functions and wavelets. For example, the discrete filters relate the continuous-time scaling functions and wavelets at different scales (via the dilation or refinement equation). An infinite iteration of an appropriate discrete lowpass filter converges in the limit to the continuous-time scaling function. Given the continuous-time scaling function, the continuous-time wavelet is produced by one application of the dilation equation with the discrete highpass filter. Further information concerning continuous-time wavelets and their connection with discrete-time filters can be found in [12, 2, 3, 13, 14, 15, 16, 4].

Extension to 2-D/M-D The simplest extension of the 1-D two-channel filterbank to 2-D signals defined over the entire 2-D plane is given by the separable 2×2-channel filterbank. This filterbank may be computed by separably applying the 1-D two-channel filterbank along the horizontal and vertical dimensions. The four 2-D analysis filters are given by the separable (tensor) products of the two 1-D analysis filters. Similarly, the four 2-D synthesis filters are given by the separable products of the two 1-D synthesis filters. One application of the 2×2-channel filterbank to a 2-D signal decomposes the signal into four subbands: low-horizontal/low-vertical (LL) frequency, low-horizontal/high-vertical (LH) frequency, high-horizontal/low-vertical (HL) frequency, and high-horizontal/high-vertical (HH) frequency. By iterating the 2×2-channel filterbank on the low-horizontal/low-vertical frequency subband one creates a 2-D discrete wavelet transform. These ideas extend straightforwardly to create separable decompositions for M-D signals. Nonseparable decompositions of 2-D/M-D signals also exist [3, 2], however we limit our discussion in this thesis to separable decompositions.

2.2 Arbitrary-Length 1-D Signals

In this section we briefly examine the problem of processing arbitrary-length 1-D signals. We specifically consider the case of arbitrary finite-length 1-D signals, however the basic problems and proposed solutions are also applicable to infinite-length one-sided signals, e.g. causal or anti-causal signals.

The simplest scenario is one in which a block transform (filter length equal to the subsampling factor equal to the blocksize) is applied to a signal whose length is constrained to be an integral multiple of the blocksize. In this case the signal boundaries can be lined up nicely with the transform blocks (and thus with the basis vectors) and the boundary problem is then avoided. As previously discussed, this is the classic case of Block-DCT processing of images and video, which forms the basis for all current image and video compression standards. While block transforms have been very successful, the constraint that the filter length equals the blocksize limits the performance. Improved performance can often be achieved by allowing longer filters.

In typical transform/subband approaches, such as general subband filtering, lapped transforms, and all wavelet schemes except for Haar, the filter length is greater than the subsampling factor. The basis functions in these representations do not conveniently end at the signal boundaries. In particular, some of the basis vectors lie on the boundary, with a portion of each of these basis vectors within the signal's ROS and a portion extending beyond. These basis functions will be referred to as *boundary basis vectors*. The basis vectors that are localized to the signal's interior, hereafter referred to as *interior basis vectors*, do not pose any problems and their coefficients may be computed in the conventional manner. The basis vectors that are completely outside the signal's support, referred to as *exterior basis vectors* also do not pose a problem since they have no effect on representing the signal. However, in order to determine the coefficients for the boundary basis functions, conventional approaches dictate a convolution (inner product) using samples both inside and outside the signal's ROS. The basic problem is that we are trying to use samples outside the signal's ROS, and these samples are undefined. This problem is illustrated for the case of a causal signal in Figure 2.3.

The previous proposals for overcoming this problem can be classified into two general classes [2]:

1. Extrapolation at the signal boundaries
2. Construction of boundary filters

The idea behind the first class of approaches is to *extrapolate or define the samples outside the signal's*

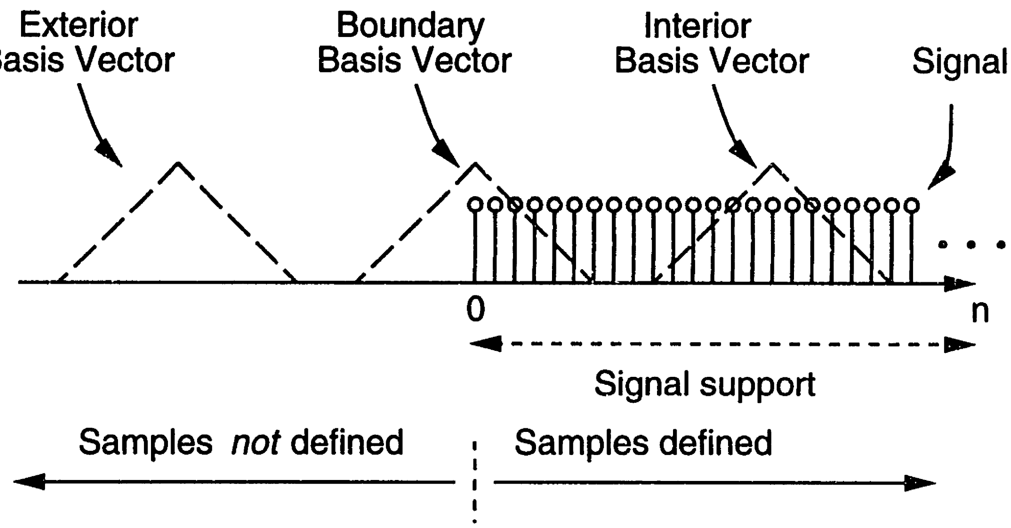

Figure 2.3: Difficulty with arbitrary-length 1-D signals: what should be the sample values outside the signal's ROS in order to compute the convolutions/inner products. A discrete-time signal whose amplitude is 1 for $n \geq 0$ and undefined for $n < 0$ is shown. Possible solutions: (1) define the sample values outside the ROS, (2) use different filters at the boundary that do not extend beyond the ROS.

ROS. Following the extrapolation, conventional convolution or inner products may be computed straightforwardly. The second class of approaches is based on the idea of *replacing the troublesome filters at the boundary by new filters that do not extend beyond the signal's ROS*. The difficulty at the boundary is that the boundary filters extend partially within and partially beyond the signal's ROS. This problem can be avoided by replacing the filters at the boundary by a new set of filters that are only defined within the signal's ROS. Thus, the signal's interior would be filtered by the original filters, while the boundary area would be filtered by new boundary filters that do not extend beyond the ROS.

Extrapolation There are many possible methods to extrapolate a signal; we will begin by discussing some of the simpler methods, and then generalize and point out some of the important properties. The simplest extrapolation approach is *zero-padding*, where we assume that all the samples outside the signal's ROS are equal to zero. While this is a very simple approach, ensuring perfect reconstruction typically requires an overcomplete representation, in which there are more coefficients than original signal samples. This problem may be overcome by *periodic extension*, where the signal would be replicated end over end (equivalently wrapped around) to form an infinite-length periodic signal. The coefficients would then be computed by circular convolutions instead of linear convolutions. A problem with both of these approaches is that they typically produce extraneous discontinuities at the boundaries: in zero-padding the signal goes abruptly from nonzero to zero, and in periodic extension the left end of the signal is connected to the right end which is typically very different. These discontinuities produce artificial high frequencies that can degrade performance in compression as well as in other possible applications.

Using a smoother extension can reduce the artificial high frequency components produced by the extrapolation. Smith and Eddins proposed the idea of *symmetric extension*, in which the original signal is concatenated with a flipped (reflected) version of itself and then periodically replicated [17]. The symmetric boundary condition produces a continuity of the signal amplitude across the boundaries, and thereby reduces many of the artificial high frequencies. The underlying concept is analogous to how the DCT provides improved energy compaction over the DFT [1]. With linear phase filters (symmetric or anti-symmetric) the symmetric extension method leads to a critically sampled representation that enables perfect reconstruction.

These ideas can be extended in a number of directions. For example, higher degrees of signal continuity at the boundary can be achieved by applying different extensions [18]. In addition, the general case of nonlinear phase filters can also be solved by applying an appropriate linear extrapolation of the signal based on the specific filters. Smooth signal extensions that produce critically sampled representations with perfect reconstruction capability have been examined by a number of authors and extended in many directions [19, 20, 21, 22, 23] and an overview may be found in [24].

Another smooth extrapolation method was proposed by Williams and Amaratunga [25, 26]. They considered the case of a wavelet transform, and they designed the extrapolation with the idea of preserving the polynomial accuracy of the wavelet transform. They propose to model the signal near each boundary using the least-squares fit polynomial of appropriate order, and then extend the polynomial to extrapolate the signal. Unlike the previously discussed approaches which was primarily motivated by the desire for improved compression, the work by Williams and Amaratunga was motivated by the desire for an appropriate treatment of boundary conditions in the numerical solution of ordinary and partial differential equations.

The various signal extrapolation approaches for processing a finite-length signal can be unified in the following manner. In each case we would like to apply a given transform/subband filtering scheme to a finite-length signal. The first step is an appropriate extrapolation of the signal so that the transform/subband scheme can be conveniently applied, as illustrated in Figure 2.4. The transform/subband is computed for the extrapolated signal, and the resulting coefficients are then sampled to provide a critically sampled representation. For reconstruction, the critically sampled set of coefficients are extrapolated to recover the coefficients before the sampling. The coefficients are then inverse transformed/subband filtered, and the original signal can be recovered by a simple windowing.

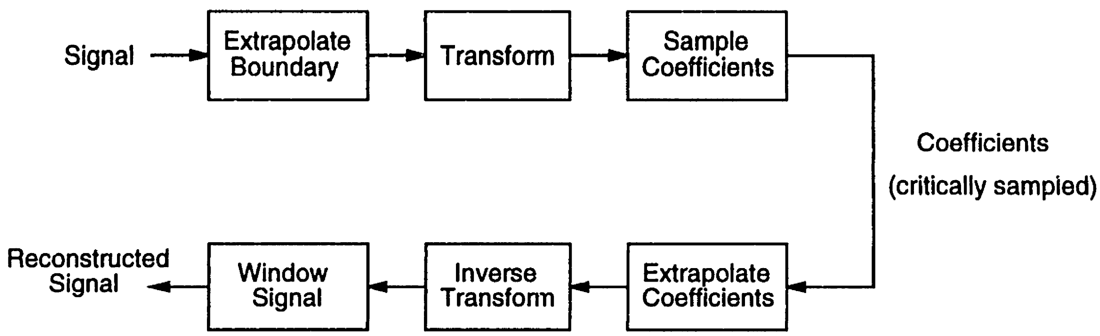

Figure 2.4: The general signal extrapolation method for creating a critically sampled perfect reconstruction transform/subband representation.

A number of general comments can be made about this class of approaches. If there is no sampling of the coefficients, then it is relatively straightforward to create a representation that enables perfect reconstruction. However, the resulting representation would be overcomplete: there would be more coefficients than original signal samples. If either the extrapolation or coefficient sampling is performed inappropriately, then perfect reconstruction may not be possible, i.e. the reconstructed signal can be distorted. In addition, the underlying assumption throughout this discussion has been that the conventional inverse transform/subband filtering will be applied for reconstruction. If we have the freedom to compute a "new inverse", then perfect reconstruction could be possible under a much wider range of extrapolations and coefficient samplings. However, computing and applying this new inverse can be a very complex process.

These various approaches trade off sampling rate, potential distortion, complexity, performance, and flexibility. As we will see, these tradeoffs exist throughout all the approaches to be discussed in this proposal. The best choice among these properties will typically depend on the particular application. In the case of finite-length 1-D signals, the signal extrapolation methods such as symmetric extension, and the boundary filter methods to be discussed next, provide simple and highly successful approaches for creating critically sampled perfect reconstruction representations.

Boundary Filter Construction Boundary filter approaches are based on the ideas that (1) the original filters can be used to represent the signal interior in the conventional manner, and (2) new boundary filters that do not extend beyond the boundary can be designed for and used to represent the signal at the boundaries. The new boundary filters replace the original filters at the boundaries. The construction of the new boundary filters must guarantee that the representation provides a basis, and that the basis has the desirable properties for high performance applications. We will examine the boundary filter methods in slightly more detail then the extrapolation methods, since they are important for our problem and also not as well known in the signal processing community.

Much of the original research on constructing boundary filters has been performed in the wavelet and mathematics communities, specifically in the context of constructing continuous-time bases over the unit interval $[0, 1]$. Lucid reviews of this work can be found in [2, 27].

One of the earliest constructions was provided by Meyer [28]. Meyer began with a wavelet basis over the entire real line and used it to construct a wavelet basis over the unit interval $[0, 1]$. Meyer's construction was simple and elegant, and based on Daubechies compactly supported and orthogonal scaling functions and wavelets. We discuss one scale of the decomposition, as the multiresolution decomposition follows naturally. Assume a scale sufficiently fine (basis functions sufficiently short) such that the boundary basis functions at the left end of the interval do not overlap with the boundary functions at the right end of the interval. Retain all the basis functions that are non-zero over some part of $[0, 1]$, and restrict them to $[0, 1]$. This leads to three sets of basis functions: left-boundary basis functions, interior basis functions, and right-boundary basis functions. These three sets of basis functions are mutually orthogonal, and their union spans the space. Each of the interior basis functions are orthogonal to each other. However, this does not hold for the left or right boundary basis functions since they have been clipped (their tails are not in the interval $[0, 1]$). In addition, the left and right boundary basis functions are redundant (extra functions exist in both sets). An orthogonal basis can be constructed by orthogonalizing the boundary functions at each end, for instance by applying the Gram-Schmidt procedure. In this manner, an orthogonal basis for the space can be constructed[1].

The fact that the orthogonalization process corresponds to a linear transformation is very important. This means that each of the new boundary basis functions corresponds to a linear combination of the original boundary basis functions. Since the original boundary basis functions satisfy a dilation equation from one scale to the next, the new boundary basis functions satisfy a new dilation equation. Specifically, the new boundary functions at a given scale can be expressed as a linear combination of the new boundary and interior functions at the next finer scale. In fact, there are a total of three dilation equations that relate the new left boundary functions, interior functions, and new right boundary functions, from scale to scale. However, the key point is that the new basis that Meyer has constructed still satisfies the dilation property, and therefore produces a multiresolution hierarchy of nested spaces.

[1] The actual process is slightly more complex than described: first the boundary scaling functions are orthogonalized, and the wavelets are then determined to span the orthogonal complement of the scaling function subspace.

Meyer's construction produces a multiresolution, orthonormal decomposition for $L^2([0,1])$, which also is an unconditional basis for a wide class of function spaces. Meyer's construction also has some disadvantages [29]. One potential difficulty is that the decomposition of a given scale into a lowpass and a highpass subspaces is in a sense uneven, as the dimension of the lowpass subspace is greater than the highpass subspace (there are more scaling functions than wavelets at each scale). It has been stated [29] that this may prevent wavelet packet constructions over $[0,1]$. However, this may be more of an inconvenience than a fundamental difficulty, as for example, wavelet packet constructions can still be made, though the tree may be imbalanced.

A more severe difficulty relates to the actual construction of the boundary basis functions. Since some of the boundary functions may be very small over $[0,1]$ (as only their tails overlap the interval) the construction may be very ill-conditioned, especially for longer filters. This problem with the "tails" also affects other possible applications. For example, a fundamental signal processing problem is to determine the appropriate extrapolation of a signal. In this case, we may want to extrapolate the signal from $[0,1]$ to the entire real line. Since the original boundary basis functions were defined outside $[0,1]$, one possibility is to try to use them in the extrapolation. A natural method to accomplish this is to reconnect the previously truncated portions outside $[0,1]$. However, some of the boundary basis functions with small "tails" within the signal's support (and large portions outside the support) may have very large amplitude coefficients. The resulting extrapolation will be skewered by these few high-amplitude terms.

A third difficulty relates to the frequency response of the constructed boundary filters. Wavelets typically correspond to bandpass filters with bandwidth of approximately an octave. However, the orthonomalization procedure produces boundary wavelets that are much more oscillatory, with much broader bandwidths.

The problem of finite-length discrete-time signals was examined by Herley and Vetterli [30, 31]. While Meyer applied Gram-Schmidt to the boundary basis functions to construct one orthogonal basis, Herley and Vetterli identified that there exists a *whole space* of orthogonal boundary solutions (this solution space contains any orthogonal basis for the space spanned by the boundary basis functions). Therefore there is considerable freedom in optimizing the boundary filters. For instance, one may construct the boundary filters that provide the best frequency selectivity. In addition to applying these ideas for processing finite-length signals, Herley and Vetterli also applied them to the problem of time-varying filterbanks. In this case, the goal is to switch between different filterbanks while preserving critical sampling, perfect reconstruction capability, and high performance. This approach enables nice transition regions between different filterbank topologies or simply different filter sets. They also showed how their time-varying discrete-time filters can be iterated to produce time-varying continuous-time bases. This illustrates the elegant connection between discrete-time and continuous-time bases: Meyer constructed a basis for the unit interval working solely in the continuous domain, while Herley and Vetterli constructed a basis for the unit interval by first constructing a finite-length discrete-time basis, and then iterating.

A very important property of wavelets is their ability to accurately represent polynomials up to a given order. Specifically, given a scaling function of accuracy p, all polynomials of degree $\leq p-1$ can be exactly expressed as a linear combination of the scaling functions. Daubechies filters of length N lead to scaling functions of accuracy $p = \frac{N}{2}$ and wavelets with p vanishing moments [2]. The polynomial accuracy property is critical for many of the wavelet approximation properties – which have directly led to many new theoretical results and practical applications. Therefore, it is very important to preserve this property. In the previously discussed proposals, Meyer's construction does preserve the polynomial accuracy property, however it has a number of potential difficulties. On the other hand, Herley and Vetterli's discrete-time approach does not guarantee the continuous-time polynomial accuracy property.

Cohen, Daubechies, and Vial propose a different approach for constructing a wavelet basis over the unit interval which preserves the polynomial accuracy property and also avoids some of the difficulties of Meyer's construction [29]. To preserve the polynomial accuracy property, they must construct boundary scaling functions such that their combination with the interior scaling functions would perfectly represent all polynomials on $[0, 1]$ up to a certain degree. Their construction was simple and direct – they explicitly found the functions needed to preserve the polynomial accuracy, and included them in the basis. Consider a single scale and Daubechies orthogonal scaling functions and wavelets (which can exactly represent all polynomials of degree $\leq p-1$). The interior scaling functions would perfectly represent the polynomials within the interior of the signal, the problem would only occur at the boundaries. Consider each polynomial to be approximated, beginning with the DC. Using only the internal scaling functions, the DC polynomial would be represented perfectly except for an error at each end of the interval. Their idea is that if these "error pieces" were part of the basis, then the DC polynomial would be represented perfectly – therefore add these pieces as boundary scaling functions. In this way the DC polynomial will be perfectly represented. Continuing in this manner, with the augmented set of scaling functions and a linear polynomial, there would once again be an error at each boundary. By adding these new "error pieces" as boundary scaling functions, one could then perfectly represent both DC and linear polynomials. This process would be continued up to and including polynomials of order $p-1$. An interesting note is that a total of p boundary scaling functions are required at each end to represent all polynomials of order $p-1$, but there is only

[2] Nonorthogonal filters lead to nonorthogonal scaling functions which can tradeoff orthogonality for higher accuracy with the same length.

space for $p - 1$ scaling functions. To alleviate this problem they discard the outermost interior scaling function at each end. The boundary scaling functions can then be orthogonalized, and the boundary wavelet functions can also be constructed and orthogonalized.

This construction leads to an orthogonal basis for any scale with approximation accuracy p. In addition, the decomposition is balanced in that there are the same number of scaling functions and wavelets at each scale (albeit at the cost of discarding the outside interior scaling functions). The construction is also much more numerically stable then Meyer's approach. Furthermore, the new boundary scaling functions and wavelets are less oscillatory than Meyer's. However, the boundary scaling functions partially have the form of polynomials (which they were designed to represent) as opposed to being similar to the internal scaling functions.

An important issue with this construction is that the polynomial accuracy property is only preserved for a single scale. Applying one-level of CDV's original approach (similar properties also hold for Meyer's approach) to a polynomial signal of order $\leq p - 1$ leads to all the signal energy concentrated in the lowpass band and the highpass band is zero. However, the lowpass coefficients are *not* a polynomial at the boundaries, and can even be spiky. Therefore, a second 2-band decomposition applied to the lowpass signal would have energy spread in both the lowpass and highpass subbands. The polynomial accuracy property is only preserved for a *single* level of decomposition.

CDV recognized this problem and showed that it is impossible to have an orthogonal representation that preserves the polynomial accuracy property at all levels. They then proposed to pre-conditioning the signal before applying their finite-length transform. The pre-conditioner corresponds to a linear transformation at the boundaries which converts polynomials at the boundary into signals that will produce zero highpass coefficients at all levels of the decomposition. This is illustrated in Figure 2.5 (bottom plots) where all the highpass coefficients are zero. Their pre-conditioning step (corresponding to a non-orthogonal matrix multiply) trades off orthogonality for the polynomial accuracy at all levels. In this way, they construct a biorthogonal multiresolution decomposition of the unit interval with polynomial accuracy p at every scale.

An undesirable aspect of this approach is that in order to have zero highpass coefficients at each level, the smooth polynomial is transformed into a non-smooth signal at the boundaries. A (smooth) polynomial signal will therefore have a (non-smooth) non-polynomial set of lowpass coefficients, as shown in Figure 2.5. A smooth signal will therefore have a non-smooth representation. This has a number of disadvantages. For example, in compression one would like a constant signal to have constant-amplitude lowpass subband coefficients, and not a spiky set. More importantly, it is somewhat disconcerting for a smooth signal to have a non-smooth representation. Otherwise,

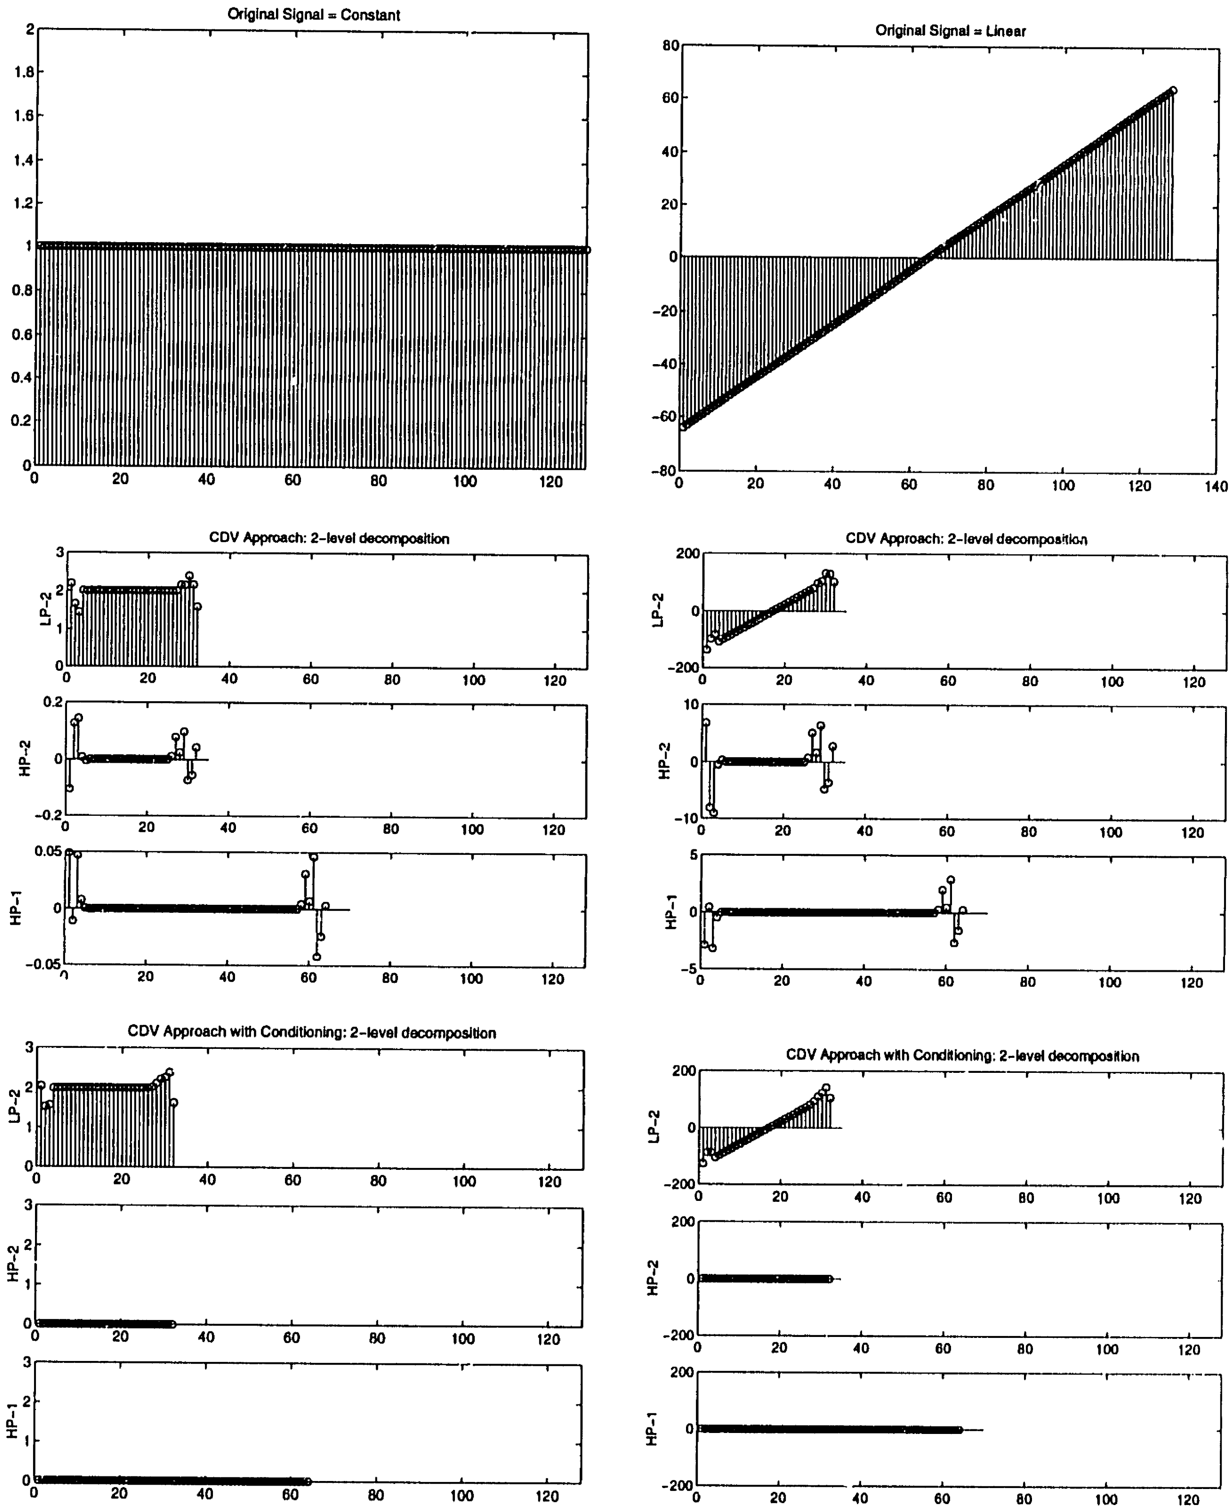

Figure 2.5: The constant signal (left column) and linear signal (right column) are represented using original CDV (middle row), and CDV with conditioning (bottom row). In each case a two-level wavelet transform is taken using Daubechies 14-tap filters. Notice that CDV has artificial energy in the highpass coefficients while CDV with conditioning has zero highpass coefficients. Both CDV and CDV with conditioning have non-smooth lowpass coefficient amplitudes.

some of the simplest and most important properties are disguised. For instance, a smooth signal can have a non-smooth representation as shown in Figure 2.6 (top two plots). Furthermore, when the representation appears smooth, for example when the lowpass coefficients are constant and all other coefficients are zero, the original signal may be far from smooth, as shown in Figure 2.6 (bottom two plots). It is highly advantageous to be able to simply characterize the properties of a signal by simply examining the properties of its transform. The lack of the smoothness property in CDV's approaches (both the original and with conditioning) makes this more difficult.

2.3 2-D Signals and Images with Rectangular Supports

The simplest class of 2-D signals to analyze are those with rectangular ROS. For example, images typically have square or rectangular support. Transform/subband analysis for these signals is straightforward as one may directly apply all the knowledge from 1-D arbitrary-length signals: a 2-D basis may be constructed as the *separable* (tensor) product of a 1-D basis. This construction enables the theoretical analysis from 1-D to carry over to 2-D, and furthermore provides an extremely nice computational structure. Specifically, a separable 2-D transform over a rectangular ROS may be computed by separably computing 1-D transforms of first the rows and then the columns. Decomposing the 2-D transform into a number of 1-D transforms leads to very low computational and memory requirements (this idea of decomposing a large transform into smaller transforms is also the basis for the FFT algorithm and many other "fast" algorithms).

Boundary processing for these signals is based on the simple boundary processing for 1-D finite-length signals. Most image processing transform/subband applications are based on separable approaches with circular convolution or symmetric extension at the boundaries. Separable decompositions of this form may be applied to any M-D signal with rectangular ROS.

We should briefly comment on the application of non-separable transform/subband schemes to 2-D or M-D signals. A non-separable filter has the advantage over a separable filter that for a given size it provides much more freedom. However, the design of a non-separable filter with a given set of properties is much more difficult, since the highly developed 1-D filter theory which is easily applied to separable filters has limited applicability to non-separable filters. This design difficulty is much more profound in the case of a critically sampled perfect reconstruction non-separable transform/subband scheme. In addition, the complexity of a non-separable scheme can be significantly higher, since it may not be decomposable into a number of smaller transforms. Furthermore, a nonseparable scheme may not lead to simple boundary processing for signals with simple, rectangular support. Some applications dictate the use of nonseparable processing

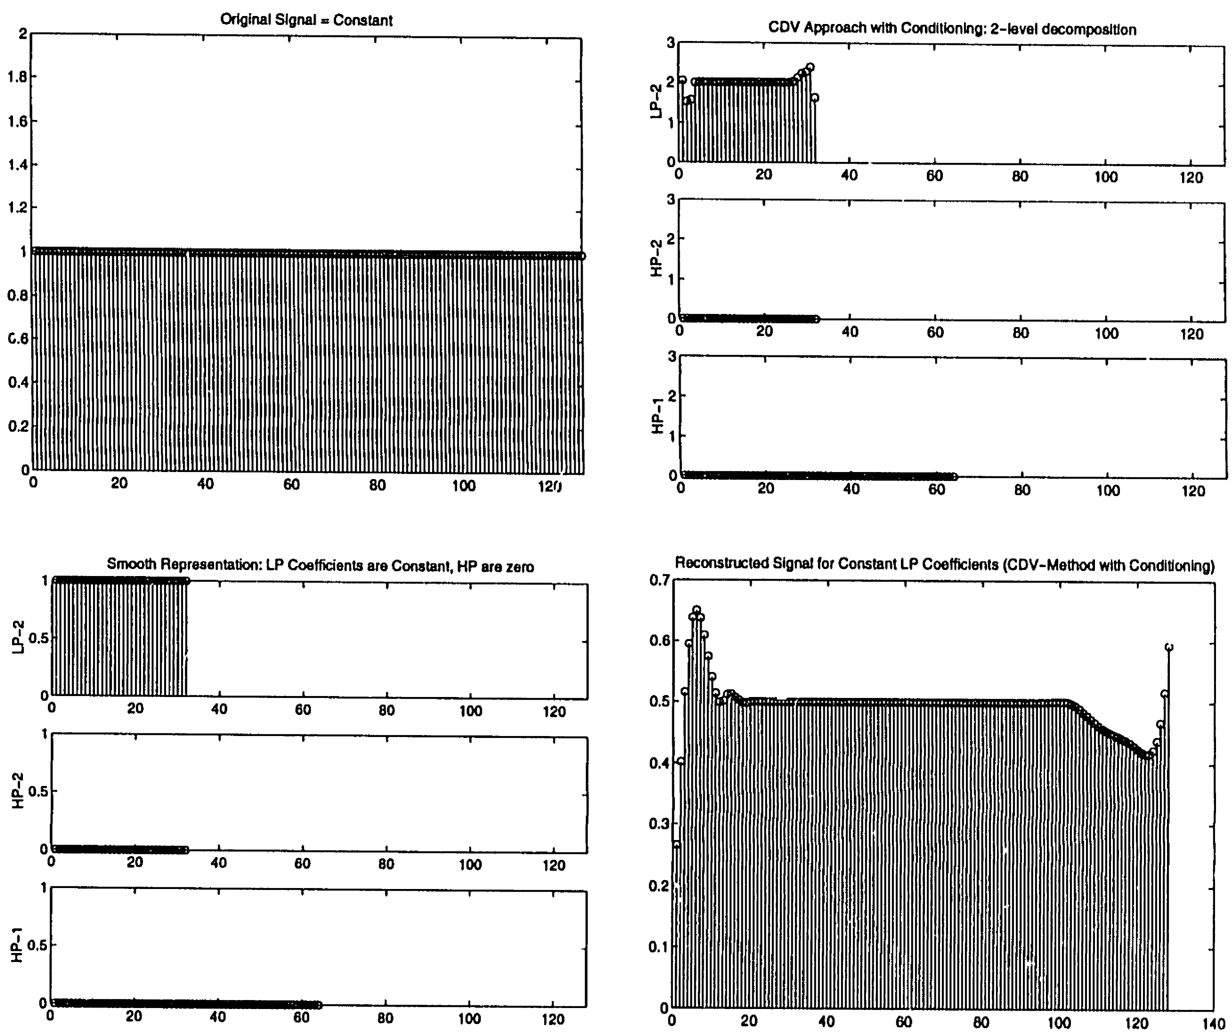

Figure 2.6: These figures illustrate the lack of smoothness for the CDV scheme (with preconditioning). The smooth constant signal (top left) has zero bandpass and highpass coefficients, but non-smooth lowpass coefficients (top right). The smooth coefficients (bottom left), with constant amplitude lowpass coefficients and zero bandpass and highpass coefficients, corresponds to a non-smooth signal (bottom right).

(e.g. when a non-rectangular frequency decomposition is desired), but generally non-separable processing is not required, and its slight improvement over simple separable processing does not usually justify its higher complexity. Therefore, throughout this work we only consider separable transform/subband schemes (where the 2-D basis vectors are tensor products of 1-D schemes).

2.4 2-D Signals with Arbitrarily Shaped Supports

The various 1-D proposals offer elegant, practical, and successful methods for processing arbitrary finite-length 1-D signals. They can also be straightforwardly applied to processing 2-D or M-D signals with rectangular regions of support. However, it is unclear how to extend these approaches to successfully represent 2-D or M-D signals with arbitrarily shaped regions of support (AS-ROS). One of the difficulties is due to the geometry. For example, it is relatively simple to determine a smooth extension for a 1-D signal, e.g. symmetric extension. However, it is unclear how to determine a smooth extension for a 2-D signal, as the extension should be smooth both horizontally and vertically. It is even more unclear how to determine a smooth extrapolation that would produce a critically sampled perfect reconstruction representation.

A practical difficulty relates to the immense size of typical 2-D or M-D signals. Even a simple 2-D object from an image can contain many tens of thousands of samples, and a 3-D signal could contain significantly more – and a large fraction of these samples may be explicitly involved in the boundary processing. The complexity issue may be clarified by comparing to the 1-D problem. For example, consider a 1-D signal of 1000 samples. As long as the filters are short then only a small fraction of the total number of signal samples would we involved in the boundary processing. Now consider a 2-D or M-D signal containing 1000 samples and a corresponding filter that is the separable product of the 1-D filter. Depending on the signal's ROS, a large fraction of the samples may be affected by the boundary processing. Therefore, it is not the total number of samples that affect the complexity, but the number of samples that are affected by the boundary processing [3]. The complexity can be intuitively related to the "size" of the boundary, e.g. the perimeter for a 2-D signal, or the surface area for a 3-D signal. For a fixed total number of signal samples, the larger the boundary, the greater number of samples affected by it, and the greater the complexity. Therefore, processing AS-ROS 2-D or M-D signals can be extremely complex. This complexity may prohibit the application of the 1-D approaches or limit them to special classes of problems.

[3]Of course, if we use vectors with global support then all of the signal samples are involved in the boundary processing.

In this section we will briefly overview the various proposed approaches for representing AS-ROS 2-D and M-D signals. In addition to the extension of the 1-D approaches to 2-D/M-D, there have also been a wide variety of other proposed approaches that attempt to achieve good performance at reasonable complexity. While our goal is a critically sampled perfect reconstruction representation, we will also examine the other approaches to determine if another approach is more appropriate for the 2-D/M-D problems. The various proposed approaches for representing signals with arbitrarily shaped supports (and also for the general case of any arbitrary support) can be classified into three groups:

- *Undercomplete representations:* The representation does not span the signal space and therefore does not guarantee perfect reconstruction. Typically the case when there are fewer vectors than that required for a basis.

- *Complete representations:* The representation corresponds to a basis for the signal space. Equivalent to a critically sampled perfect reconstruction representation.

- *Overcomplete representations:* The representation spans the signal space, but there are more vectors than required for a basis.

The familiar ideas of constructing a basis and performing an extrapolation will appear in the complete and overcomplete classes, respectively.

Throughout this section, the goals and difficulties may be visualized by reflecting back to some simple examples, e.g. consider the 2-D arbitrarily-shaped object from an image shown in Figure 2.7.

2.4.1 Undercomplete Representations

Historically, the earliest representation for signals defined over arbitrarily shaped ROS were based on undercomplete representations. For instance, a common approach was to approximate the signal as a linear combination of low-order polynomial basis vectors, i.e. to approximate the signal amplitude as a low-order polynomial [5]. Since the number of coefficients in the polynomial were typically far less than the number of samples in the signal, this representation did not have the perfect reconstruction capability. The use of higher-order polynomials was impractical as the complexity would be immense, and also the uniform sampling of the polynomials would lead to a highly ill-conditioned representation.

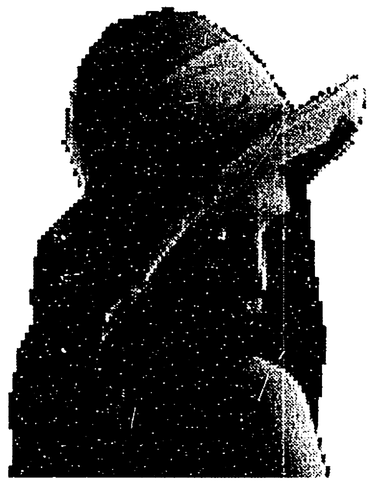

Figure 2.7: A 2-D signal with an arbitrarily shaped region of support.

2.4.2 Complete Representations

A complete representation is equivalent to a critically sampled perfect reconstruction representation. These representations provide a basis for the signal space. There are two groups of approaches that will be discussed within the class of complete representation. The first group results from the simple application of 1-D transforms to the rows and columns of an AS-ROS signal. The resulting representations are not genuine 2-D representations, but they require very low complexity. The second group of approaches construct genuine 2-D representations (bases) for the given 2-D AS-ROS signals. However, these representations can be quite complex because of the required basis construction.

Simple Application of 1-D Schemes These approaches are motivated by the fact that while it is quite difficult to construct a basis for a 2-D signal with arbitrarily shaped ROS, it is rather straightforward to construct a basis for an arbitrary finite-length 1-D signal. Therefore this becones the idea of converting the difficult 2-D problem into one or a number of 1-D problems. Possibly the simplest approach is to scan the 2-D signal into a 1-D signal (e.g. one row at a time) and then compute a 1-D transform. Alternating left-to-right and right-to-left scanning of the rows could reduce the artificial discontinuities in the resulting 1-D signal. Approaches of this form lead to very simple critically sampled perfect reconstruction representations which achieve low complexity. However, their primary drawback is their inability to exploit the 2-D structure that exists in the signal. This drawback severely limits their performance.

A popular set of approaches that attempt to exploit the 2-D structure are based on performing 1-D transforms of each row and each column. The "Shape-Adaptive DCT" (SA-DCT) by Sikora and Makai is probably the most well-known of these approaches [32]. These approaches begin by computing an invertible transform of each row. For instance, this may be an appropriate length DCT of each row, or in general any PR transform/subband approach for the given length row [32, 33]. The resulting coefficients can then be aligned (organized) in some fashion, and then an invertible transform is computed for each column. As long as each step of the processing is invertible (row processing/alignment/column processing) the entire approach is invertible. Also, in the event that the 2-D signal has a rectangular ROS, this method may reduce to the conventional block transform (depending on the details of alignment, etc.). This approach leads to a critically sampled perfect reconstruction representation that only requires short 1-D transforms, thereby enabling low computational and memory requirements. Another interesting approach along these lines and based on a wavelet transform is given in [34].

A disconcerting issue related to these approaches, is that while they utilize separable computations (1-D transforms are computed of the rows and columns) they are non-separable. For example, the resulting coefficients depend on the order of the operations (column first or row first) – the horizontal and vertical operations do not commute. This may be easily illustrated by examining a 2-column signal where the first column contains 3 pixels and the second contains 2 pixels. Computing first the DCT of each row will lead to a different final set of coefficients than computing first the DCT of each column. In addition, the representation does not express the signal as a linear combination of separable basis vectors, i.e. the basis vectors have non-rectangular ROS's (that depend on the order of operations) [4]. Furthermore, since full-length transforms are taken of each row and column, it appears that the entire transform uses basis vectors with global support, but this is actually not the case. In a conventional full-support 2-D Block-DCT, for example, every sample within the region is used to compute each coefficient value, and in the reconstruction each coefficient affects every sample in the reconstructed region. However, this is not necessarily true with these new proposals. Overall, these approaches do not fully exploit the 2-D signal structure, they exhibit strange behavior, and their theoretical analysis is much more cumbersome and complex. These difficulties make analysis of the representation of a given signal, as well as any subsequent processing of the representation, much more difficult. A genuine 2-D approach would avoid these problems.

[4] Note that the non-separable basis vectors result from the lack of commutability of the horizontal and vertical transforms, and do not result from them being defined over an arbitrarily shaped support.

Genuine 2-D Schemes There are many approaches for constructing a genuine 2-D basis for the signal space over the arbitrarily shaped ROS. One possibility is to model the signal as an appropriate stochastic process, and compute the KLT over the ROS (see, e.g. [35, 36]). Another well-known approach, proposed by Gilge, Engelhardt, and Mehlan, begins with a set of vectors that spans the signal space, and applies an orthogonalization procedure to construct an orthogonal basis over the signal space [37]. For example, if the ROS is circumscribed by a square, the vectors corresponding to the DCT basis over the square will span the signal space of not only the square, but also of the arbitrarily shaped region. By beginning with the set of DCT vectors and applying an orthogonalization algorithm such as Gram-Schmidt, an orthogonal basis for the region may be constructed.

Both of these approaches yield critically sampled representations, and the orthogonality may simplify subsequent processing such as quantization. However, they are extremely complex as they require the extremely complex operations of eigendecomposition and orthogonalization, respectively. Furthermore, even after the basis is constructed, since the basis is unstructured (can not be decomposed into smaller transforms), the coefficient computation will be extremely complex – requiring very large computational and memory requirements. The complexity is therefore both in terms of constructing the basis and then computing the coefficients of the signal with respect to the new basis. In a compression system, the decoder would also be extremely complex, since it must construct the basis in parallel with the encoder, and then reconstruct the signal given the coefficients. Additional difficulties with the orthogonalization include that the final basis will depend on the specific order that the basis vectors are orthogonalized, and also that the final basis is not guaranteed to have any other good properties besides orthonormality, i.e. it may perform badly for compression.

It is instructive to examine the elegant 1-D approaches of Meyer, Herley and Vetterli, and Cohen, Daubechies, and Vial, and determine their applicability for 2-D or M-D problems. These approaches consist of two basic elements: (1) retaining the interior basis vectors and (2) constructing a new set of boundary basis vectors to replace the troublesome original boundary basis vectors. The new boundary vectors are designed to guarantee a basis and provide some useful properties such as polynomials accuracy or good frequency selectivity. Unfortunately, there are a number of difficulties that arise if one attempts to extend these methods to 2-D or M-D problems.

A straightforward extension of the discussed approaches to 2-D or M-D would result in the loss of local support for the boundary basis vectors. To understand this, lets reexamine the 1-D case: if the original basis vectors had local support, the new boundary basis vectors would also be guaranteed to have local support. This results since the new boundary vectors are restricted to lie in the boundary regions at the ends of the signal. However, in 2-D the boundary exists around the entire 2-D support. The new boundary vectors can lie anywhere in this region, and may form a "ring" around the interior – completely losing the important local support property. One possible approach to preserve the local support of the boundary basis vectors is to partition the boundary region into local subspaces (local with respect to spatial support), and then perform an appropriate construction within each subspace.

Recently, the mathematicians Cohen, Dahmen, and DeVore have made major inroads on this problem [38]. They have shown that given a continuous domain $\Omega \in R^d$ satisfying some constraints, it is possible to construct a multiresolution decomposition over the domain which preserves the polynomial accuracy property, local support of the basis vectors, and stability of the basis. Their analysis is quite sophisticated, and it is too early to decide the practical applicability of their approach.

A practical difficulty with these approaches is the increased complexity for the higher dimensional problems. For example, in the 1-D scenario, there are only a small number of boundary possibilities (topologies), and these can be identified and all possible boundary filters can be precomputed and stored for application to any 1-D signal. However, in 2-D or M-D the number of boundary possibilities is immense, and it is impractical to store all the possible boundary filters [5]. Therefore the boundary filters would probably have to be individually constructed for each and every signal to be processed [6].

Overall, the elegant, practical, and highly successful 1-D approaches of constructing frequency-optimized boundary filters (Herley and Vetterli) or polynomial accuracy boundary filters (Cohen, Daubechies, and Vial) appear impractical to extend to 2-D or M-D problems, except for special cases that explicitly require the properties that they provide, and can afford the complexity.

Another class of approaches that appear promising are based on the *lifting scheme* proposed by Sweldens [39, 40, 41, 42]. Sweldens is a numerical analyst who considered the problem of developing wavelet-type representations for more general settings than the real line or the unit interval, e.g. to create wavelet bases over curves, surfaces, or general domains, and for irregular sampled data. Sweldens realized that it is difficult to apply in these general settings a conventional wavelet transform where the basis vectors are related to each other by translation and dilation. He proposed a representation where the basis vectors are not necessarily related to each other by translation and dilation, i.e. the basis vectors (and filters) vary spatially to adapt to the region

[5] By restricting the filters to very small supports the number of boundary filters can be greatly reduced, thereby reducing the required storage. However, the application of these filters would still be quite cumbersome.

[6] The local support of the basis vectors will make this slightly less complex than the full-support basis vectors of Gilge's original proposal, but the construction would still be very complex and cumbersome.

over which they are defined. These *second generation wavelets* preserve some of the conventional wavelet properties such as local support in space and frequency, accuracy in locally reproducing polynomials, and existence of fast algorithms. Overall, the lifting scheme appears promising for a number of applications, including creating wavelet bases over the sphere [43, 44]. The applicability of lifting for creating wavelet bases over arbitrary domains in R^d has been briefly mentioned [39, 40, 41], however no indepth studies have yet been performed. An important issue relates to the complexity of adapting to the support of the signal. In an application such as compression, the adaption must be performed at both encoder and decoder, and therefore its complexity is of prime concern.

2.4.3 Overcomplete Representations

Creating a transform/subband representation for an arbitrarily shaped ROS is a very complex process, while creating a representation for a square or rectangular ROS is quite straightforward. This naturally leads to the idea of embedding the given region inside a square or rectangle, and performing the processing over the square. For example, if the given region is circumscribed by a square, the vectors corresponding to the DCT basis over the square will span the signal space of not only the square, but also the region. This leads to the general idea of using a basis defined over a superset signal space to represent the given signal space. An overview and unification of different approaches of this type may be found in [45].

It is instructive to briefly compare this class of approaches to the complete (critically-sampled) approaches. In the critically sampled approaches, the difficulty lies in constructing the basis (producing a minimal set of vectors that spans the signal space), and to a lesser extant in computing the coefficients with respect to the basis. In an overcomplete representation as discussed above, the spanning is automatic and no basis construction is required. However, the overcomplete representation has more coefficients than samples in the original signal; the vectors are linearly dependent over the ROS and this linear dependence makes subsequent processing more complex [7]. Since overcomplete representations tradeoff complexity of constructing a basis for complexity of subsequent processing, it is important to consider the difficulty of subsequent processing.

In our work, we desire a representation that can accurately approximate the appropriate signals using only a small fraction of the total number of coefficients. This capability is important for compression as well as many other applications. To be more precise, we would like to determine

[7] Not all the subsequent processing is more complex. For instance in compression, the decoder for an overcomplete representation can be much simpler than for a critically sampled approach, since no basis construction is required.

the best approximation to a signal using only a small number of coefficients, where best is measured in the least-squares sense. An important note is that for a general overcomplete representation (without additional built-in structure), determining the k best coefficients for approximating a signal is NP-hard. Therefore, the general problem corresponds to a combinatorial optimization problem, and the optimal solution can not be guaranteed without examining all the possibilities [8]. A proof of the complexity of this problem in the context of a general overcomplete representation is given by Davis [46]. Therefore, the various approaches to be discussed are practical schemes to achieve high-quality, though suboptimal, solutions [9].

The topic of overcomplete representations has been around for a while, but it has recently received renewed interest in the signal processing and mathematics communities (also referred to as theory of frames or underdetermined linear systems). Considerable research has been performed by a number of authors in different disciplines, but we will briefly touch upon only the work that is relevant to the problem at hand.

The problem of representing a signal with a basis defined on a superset signal space leads to two problem formulations. The first formulation considers what is the best *extrapolation* to fill the area outside the region and within the square, such that a transform over the square would provide the desired coefficients. The second formulation directly considers the *overcomplete representation within the signal space*, and considers which of the many possible decompositions is the best to choose from. These two formulations appear different, however they are very similar, and their different perspectives lead to different insights and algorithms (see [47, 45]).

Extrapolation to a Circumscribing Square Similar to the 1-D case, the goal is a smooth extrapolation of the signal to a circumscribing square, such that the 2-D transform over the square will result in only a small number of high-amplitude coefficients. However, in this case it is unclear how to perform an extrapolation that would produce a critically-sampled representation, and we settle for an overcomplete representation.

A number of problems arise, including what is the optimal extrapolation; an extrapolation optimal for one region may be bad for another. Therefore, a fixed extrapolation scheme is not

[8] This is with respect to minimizing the l^2 norm (energy) of the error. For other norms, such as l^1, the optimal solution may be determined much more efficiently.

[9] It is interesting to note that from a practical point of view, the optimal solution may *not* be desirable. For example, depending on the properties of the representation (e.g. how overcomplete it is), the optimal set of k coefficients for a signal may be unstable: a very small change in the signal can lead to a completely different set of k optimal coefficients. This is undesirable from a general signal processing perspective, since we would like similar signals to have similar sets of coefficients. Also, the optimal set of coefficients may not be appropriate for compression, since their amplitudes may be exceedingly large. The suboptimal approaches to be described have the benefit of stabilizing the representation.

appropriate for all regions. Also, given an extrapolation, there is the question of how to determine the appropriate set of coefficients and their amplitudes.

Let us assume that an appropriate extrapolation to fill the square has been performed. Also, assume that the k coefficients to be retained are known, for instance they may be the k most energetic coefficients from the transform computed over the square. The coefficient amplitudes that result from computing a transform over the square will yield the coefficient amplitudes to minimize the MSE over the square. However, the goal is to minimize the MSE over the region, and not the square. To solve the problem one may explicitly set up and solve the normal equations for the least squares problem, but this is an extremely cumbersome and complex process. Chen, Civanlar, and Haskell recognized that the optimal coefficient amplitudes may be computed by an iterative process that utilizes the fast and separable processing of the superset basis and avoids setting up the linear system [48]. Specifically, they based the algorithm on the theory of projections onto convex sets (POCS) [49, 50]. Geometrically, the algorithm performs alternate projections onto the k-dim subspace defined by the selected vectors, and the affine subspace of possible extrapolations. The iteration will converge in the limit to the coefficient amplitudes that produce the least squares solution.

Although the theory of POCS was cited to guarantee convergence, given that the constraint sets are a subspace and an affine subspace, this is more specifically an example of the Alternating Projection Theorem. This observation is useful since one can then bring to bear a multitude of numerical linear algebra techniques for solving the problem or increasing the rate of convergence of the iterative algorithms, e.g. various relaxation algorithms, conjugate gradient methods, etc. [8].

The key idea, however, is that the fast and separable properties of the superset basis can be exploited to solve problems over arbitrarily shaped regions. To summarize, the POCS-based approach begins by (1) selecting the k important coefficients (possibly by performing an initial extrapolation), and then (2) applies the iterative scheme to compute the amplitudes of the selected coefficients that would minimize the MSE over the region. This approach achieves good performance, with low computation and memory requirements.

A limitation of the POCS-based approach is that the set of coefficients must be selected at the beginning and held fixed throughout the iterations. However, the best set of coefficients is unknown at the start. Apostolopoulos and Lim proposed to refine the coefficient set as the iteration progresses [47]. In this way, the iterative nature of the algorithm can be exploited to fine-tune not only the coefficient amplitudes, but also the set of retained coefficients. For example, at each iteration one can adaptively select and retain the k most energetic coefficients. Geometrically, the POCS-based approach of Chen, Civanlar, and Haskell converges to a solution by alternating projections onto the k-dim subspace chosen during the first iteration and the affine subspace of possible extrapolations. With adaptive iteration, we refine our choice of the k-dim subspace as the iteration proceeds.

Since the set of selected coefficients can change from iteration to iteration, it corresponds to a non-convex constraint set. Therefore the theory of POCS can no longer be used to guarantee convergence. However, in an application such as compression, the goal is to minimize the distortion, and not to converge to a set of coefficients. The distortion in this approach is guaranteed to monotonically decrease with each iteration [47] [10]. In addition, overrelaxation techniques (e.g. geometrically projecting beyond the affine subspace of possible extrapolations) can be applied with two important benefits: (1) they can significantly accelerate the convergence, and (2) they can also expand the coefficient search space leading to an improved selection of coefficients. Furthermore, quantization may be incorporated within the iteration, and the distortion is still guaranteed to monotonically decrease with each iteration. Overall, this approach is a generalization of the POCS-based approach, which provides improved selection, accelerated convergence, and enables the incorporation of quantization within the iteration.

Overcomplete Representation within the ROS In these approaches, the goal is to find the best linear expansion of the signal over an overcomplete (redundant) set of vectors. Once again the vectors are defined in terms of a basis over a circumscribing square, however only the portion of each vector over the signal's ROS is considered. Since the vectors are linearly dependent, choosing the best k vectors is quite complex. We will examine two sets of approaches for attacking this problem.

Unlike the case with an orthonormal set of vectors, the inner products between the signal and the vectors are not true indicators of which k vectors are the most important. However, they do indicate which is the single most important vector (equivalent to the optimal detection problem). This leads to a greedy, multi-step, successive-approximation algorithm which begins by selecting the most important vector and subtracting its contribution from the signal. The error or residual is then processed in a similar fashion for selecting the next most important vector, and this continues for k steps, until k vectors have been selected. The expansion coefficients are then determined by projecting the signal onto the subspace spanned by the k selected vectors.

Mallat, Zhang, and Davis developed and analyzed this algorithm in the general context of overcomplete representations and referred to it as Matching Pursuits (MP) [52, 53, 46] [11]. Since

[10] This approach can be posed in the framework of an alternate minimization algorithm, e.g. the theory of generalized projections (GP) which concerns iterative projections onto a convex and a non-convex constraint set [51].

[11] Algorithms similar in spirit to MP have also been developed in fields as diverse as astronomy (Clean algorithm [54,

overcomplete representations are a natural vehicle for representing signals with arbitrarily shaped ROS's, MP-type algorithms have also been developed independently for this problem by a number of authors, with the earliest possibly being Kaup and Aach [60] and Desai [61]. In addition, Chang and Messerschmitt extended the work of Kaup and Aach by introducing a number of orthogonalized versions [62]. At a later date, Apostolopoulos and Lim also developed an MP-type approach and orthogonalized versions [47]. They also propose how the search steps (the matching) can be approximated while utilizing the separable/fast properties of the underlying superset basis set. Once the matching is completed, the coefficient amplitudes may be quickly computed by an overrelaxed POCS-type scheme, or by a more sophisticated iteration. This incarnation of MP requires minimal memory and utilizes the separable/fast processing of the underlying superset basis [45].

In an application such as compression the coefficients must be quantized. Quantization in MP-type approaches is typically performed within the search process, i.e. at each step the largest coefficient is selected, quantized, and used to compute the residual, thereby preventing quantization noise buildup [61, 63]. This approach, of incorporating the quantization within the search process, has the additional benefit that it avoids the coefficient amplitude computation that normally follows the search. However, in some scenarios, improved performance may be achieved by performing the search, computing the coefficient amplitudes, and then performing an appropriate quantization.

These approaches approximate the signal within its ROS, and lead to an extrapolation of the signal outside its ROS. A smooth signal extrapolation may be useful for a number of applications. Toward this goal, Mallat suggested regularizing the wavelet coefficients in order to produce a smooth extrapolation [64]. Desai presents some results on this topic in [65]. Other work along similar lines include [].

The problem of selecting the best k vectors for approximating a signal from an overcomplete representation was previously mentioned to be an NP-hard problem. Therefore, matching pursuits-type approaches would typically provide a high-quality, though suboptimal, solution. A different approach and insightful perspective was proposed by Chen and Donoho [66, 67]. They propose to determine the basis (out of all possible bases in the overcomplete representation) whose expression of the signal has the minimum l^1-norm. They refer to this method as *Basis Pursuit*.

The l^1-norm has long been known to offer a sense of sparsity in the solution. Chen and Donoho cited its ability to recognize when a signal is composed of a small number of dictionary

55]), spectroscopy (see [56] and references therein), statistics (projection pursuit [57, 58]) and neural networks [59]. Various orthogonalized versions of MP have also been developed in these fields.

elements (vectors in the overcomplete representation). Furthermore, they describe how recent advances in solving large-scale linear-programming (LP) problems make the computation of the global optimal l^1 solution feasible. A few comments are in order to compare Matching Pursuits and Basis Pursuits. Matching Pursuits attempts to determine the set of k coefficients that minimize the l^2 error, and the global optimal is not guaranteed. Basis Pursuit determines the global minimum l^1-norm basis, out of all possible bases. Basis pursuit does not guarantee the optimal set of k coefficients, nor the optimal sparsity.

An interesting point is that the proposed computational methods for obtaining the minimum l^1 solution are based on the recently developed interior-point methods. Interior-point methods are iterative algorithms which involve projections at each step. Therefore, we may once again use the separable/fast properties of the superset basis to compute or approximate the projections at each stage. On another note, this leads to the idea that a three convex set POCS iteration may also lead to the minimum l^1 solution.

Overview of Overcomplete Representations These approaches, like any transform/subband scheme, attempt to represent a signal as a linear combination of vectors. The important steps are (1) choosing an initial set of vectors or dictionary, (2) searching for the best subset of vectors (best subspace) for approximating (or exactly representing) the given signal, and (3) determining the best linear combination of the selected vectors. All the approaches use a basis defined on a superset signal space; they differ in the methods used to search for the best subset of vectors and to compute their optimal coefficients.

A Structured Class of Overcomplete Representations The choice of a structured basis on a superset signal space provides a class of overcomplete representations with a highly structured dictionary. This leads to a number of desirable properties. For instance the min l^2-norm solution can be computed easily, and the basis for the affine solution space (space of possible extrapolations) is known (no SVD required). Utilizing the *separable and fast properties of the underlying superset basis* allows fast computation of coefficients, *without* the need for storing any vectors, thereby significantly reducing both computation and memory requirements. These are important practical considerations in light of the very-high dimensionality of the problem.

2.4.4 Additional Approaches

Signals with arbitrarily shaped supports arise quite frequently in medical imaging. A basic problem in medical imaging involves reconstructing an AS-ROS signal from line projections at a number of different angles. This naturally leads to the idea of a representation based on line projections. However, this does not appear to be the most practical method. Difficulties include complex reconstruction algorithms, difficulty in guaranteeing PR, potentially high overhead (possibly making it impractical for compression), and the representation in terms of line integrals may not be as useful - one often wants localized information, and the line projections provide global information. However, this approach may be useful for extremely large signals where a highly undercomplete representation is adequate, i.e. where a very small number of line projections may provide adequate information. Also some problems may start with line integrals, and it may be better to compress the line integrals directly, as opposed to reconstructing the signal using the line integrals, and then compress it.

There are also many other possible overcomplete approaches for representing AS-ROS 2-D/M-D signals. The general methods of iterative, signal-adaptive extrapolations, beckon the hereto untested idea of computing an extrapolation for a 2-D signal by iteratively performing extrapolations along the horizontal and vertical directions. An approach of this form may converge to a smooth extrapolation for the 2-D signal.

2.5 Summary

At this point, we will attempt to summarize the different approaches and point out their important properties, advantages, and disadvantages. We will specifically describe their applicability to representing 2-D/M-D AS-ROS signals. A side note in regard to terminology is in order. Creating a critically sampled (CS) representation that enables perfect reconstruction (PR) is, of course, equivalent to determining a basis. However, considering the problem in terms of creating a CS representation with PR is very beneficial in highlighting many of the difficulties. In addition, this viewpoint may aid in the development of an efficient computational algorithm, since the actual computation will most likely be performed in a subband filtering framework. Once again, both the transform and subband perspectives to the problem can be very beneficial.

In the important problem of 1-D signals and linear phase filters, the popular symmetric extension method provides a very simple and successful approach for producing a critically sampled representation that enables perfect reconstruction and has good energy compaction. For nonlinear phase filters, the situation is slightly more complicated, but once again one can perform an extrapolation that leads to a critically-sampled representation with the PR property. However, determining an extrapolation for a 2-D or M-D signal with an arbitrarily shaped support that enables CS and PR appears to be a much more complex problem.

Boundary filter construction is an elegant approach for creating a CS representation with PR and with nice boundary vectors for processing 1-D signals. Furthermore, it enables the creation of 1-D time-varying, CS filterbanks with PR. An important practical feature of this approach is that the complex boundary filter construction can be performed once for a given set of filters and stored for simple application. This is possible since there are a very small number of boundary cases for any given filter set. However, 2-D or M-D signals have a huge number of possible boundary scenarios, and it appears impractical to precompute and store the boundary filters for all the possible scenarios [12]. In addition, it is important to preserve the local support property of the boundary vectors, and this further complicates the 2-D/M-D construction problem. Overall, the construction of boundary filters is extremely computationally complex as well as cumbersome (indexing, etc.), thereby severely limiting its application to 2-D/M-D problems.

A critically sampled representation with PR property can easily be constructed for 2-D or M-D AS-ROS signals by applying 1-D schemes along each of the dimensions, e.g. along all the rows and all the columns. As long as each of the steps is invertible, the entire operation is invertible (enabling PR). By utilizing 1-D transforms, these approaches achieve very low computational and memory requirements. However they do not fully exploit the 2-D structure in the signal, and they attempt to do so in a somewhat awkward manner.

A natural approach for representing 2-D/M-D AS-ROS signals is by using a transform/subband scheme defined over a circumscribing square (superset basis). This guarantees PR capability and facilitates the use of the separable/fast properties of the underlying superset basis. However, the overcomplete nature of the representation makes subsequent processing more intricate. There are many open questions in regard to overcomplete representations, for example how can we efficiently determine which of the many possible expansions of the signal is best for a given application.

A fundamental issue that emerges is that while it is relatively easy to construct a CS PR representation for 1-D, it is much more difficult for 2-D or M-D. Mallat expressed this quite eloquently (though in a different context) when he said that it is relatively easy to be undercomplete or overcomplete, but it can be very difficult to ride the border between the two and be complete.

[12] It may also be desirable to choose the size of the filter to apply based on the size of the signal's support; this requires that all the boundary filters for all the possible filters be stored.

The various approaches that have been examined tradeoff a number of different properties: perfect reconstruction or possible distortion, sampling rate (under/critical/overcomplete), performance (quality of the basis vectors, e.g. orthogonality, approximation property, frequency response), flexibility, and complexity (both in terms of determining/constructing the representation and subsequent processing). Of course, these different factors are intimately tied together. The appropriate tradeoff may vary considerably from one application to another. In the next chapter, we propose a novel representation that achieves a tradeoff that we feel is beneficial for a number of different applications.

Chapter 3

General Approach for Representing Signals with Arbitrary Regions of Support

In this chapter we propose a novel approach for creating critically sampled, perfect reconstruction representations for discrete signals with arbitrary regions of support. The goal is to create a representation for signals with arbitrary supports which achieves comparable performance and complexity as compared to conventional approaches used for representing signals with convenient supports.

This chapter begins by providing an overview of the general approach, and then discusses the flexibility available in designing the representation, as well as some of its important properties. In the next chapter, the proposed general approach is used to create a number of wavelet-type representations which appear promising for a variety of problems.

3.1 Overview of General Approach

The goal of this work is to create critically sampled perfect reconstruction representations for signals with arbitrary regions of support. This is equivalent to determining a basis for the signal space defined by the ROS of the signal. In the previous chapter we saw that the theory for finite-length 1-D signals is relatively mature. There are two general classes of approaches for solving the 1-D problem: (1) performing an appropriate extrapolation at the signal boundaries, and (2) constructing a new set of filters to use at the boundaries. These approaches have been highly successful for 1-D applications. However, they do not appear applicable to 2-D/M-D signals – determining an appropriate extrapolation for a 2-D/M-D A-ROS signal is unclear, and constructing boundary filters for all supports appears very complex and impractical.

We would like to create a representation for A-ROS signals that provides critical sampling and perfect reconstruction with low complexity. We would like the complexity of the representation to be comparable to that of representations for signals with convenient supports. The representations for signals with convenient supports achieve low complexity because their associated bases are simply defined and easily computed via fast/separable processing.

We wish to create a representation that provides both the fast/separable processing from bases defined over convenient supports and the critical sampling and perfect reconstruction that follows from a basis defined over the A-ROS. This leads to a *third class* of approaches for determining a basis. These approaches are motivated by the idea that we can easily create representations defined over convenient supports, e.g. the entire line in 1-D or the entire plane in 2-D. Therefore, we use our knowledge and ability in representing signals with convenient support to produce a practical and successful method for representing signals with arbitrary supports.

To illustrate this approach, consider a 2-D signal with an arbitrary (finite or bounded) support and an initial transform/subband representation defined over a larger, more convenient support that contains the given signal. For example, the initial representation may be defined over the entire 2-D plane, or over a rectangular support that circumscribes the given signal. The vectors in this initial representation correspond to a basis over the space (support) that they are defined. In addition to spanning their conveniently defined space, they also span the signal space of the A-ROS signal. The vectors in the initial representation are actually overcomplete over the A-ROS, therefore an appropriate subset of them would provide a basis for the A-ROS. Thus, the goal is to *select an appropriate subset of these vectors which will provide a basis for the A-ROS signal.* The general approach is illustrated in Figure 3.1 and can be summarized as follows:

General Approach:

1. Begin with a basis defined over a convenient support that contains the signal's ROS.
2. Select an appropriate subset of the vectors to provide a basis over the A-ROS.
3. Compute the coefficients for the selected vectors.

3.2 Properties of General Approach

The proposed approach determines a basis for the A-ROS signal space by selecting an appropriate subset of vectors from a conveniently defined basis over a superset signal space. Beginning with a basis defined over a superset space guarantees that there exists (at least one) basis for the desired signal space. However, there are typically many possible subsets of vectors that lead to a basis. This follows since the number of vectors that overlap the ROS exceeds the number of signal samples in

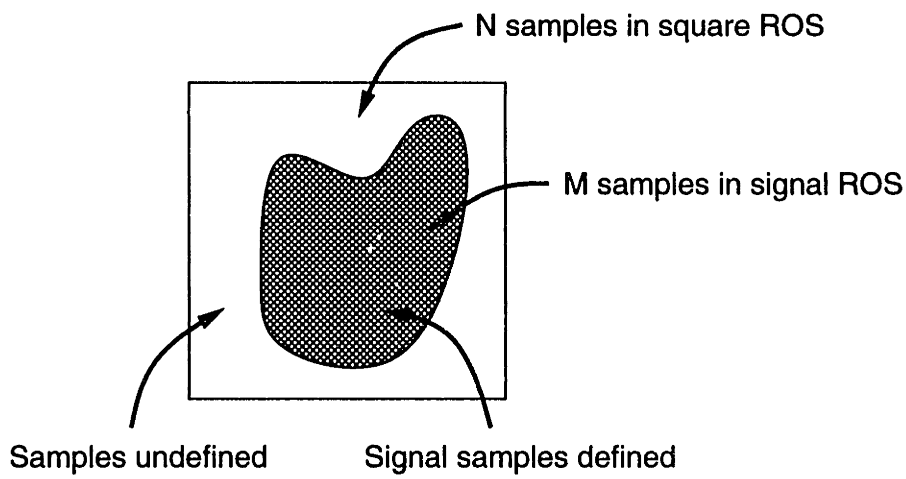

Figure 3.1: The proposed approach for creating a critically sampled perfect reconstruction representation for a signal with an A-ROS. The A-ROS signal contains a total of M samples and is shown in grey. Begin with a basis over a circumscribing square that contains the given A-ROS signal. The square contains a total of N samples, and therefore the basis over the square consists of a total of N vectors. We propose to select M out of the N vectors defined over the square such that the M selected vectors provide a basis for the A-ROS.

the ROS, i.e. the vectors are overcomplete in the signal space. Therefore, there are typically many possible subsets of vectors that lead to a basis, where each different subset leads to a different basis. Of course, given a particular performance criteria, some of these bases will perform better than others.

An important benefit of this approach is that the chosen basis for the A-ROS signal can inherit some of the approximation and computational properties from the superset basis. For example, the superset basis is typically designed to provide good approximation properties, e.g. for compression applications. Since a subset of these same vectors are used for representing the A-ROS signal, some portion of the approximation properties can be preserved. Furthermore, the basis vectors for the superset basis are simply defined and (for signals with convenient supports) their coefficients are easily computed via fast and separable computations. Since the same vectors are used for representing the AS signal, they are once again simple to define. In addition, it may be possible to exploit the fast/separable computations of the superset basis to compute the coefficient amplitudes for the A-ROS signal, as well as to reconstruct the A-ROS signal from the computed coefficient amplitudes.

An important difference between this approach and the previous proposals is that in this approach a basis can be determined without actually constructing it. Specifically, the basis is selected as opposed to constructed, i.e. each vector is either selected or discarded but it is not modified in any manner. In this manner we avoid the complexity required for constructing/modifying the vectors. Of course, in order for this to be useful we need efficient methods to select the appropriate subset of vectors to form the basis, to compute the coefficient amplitudes for a given signal, and also to reconstruct the signal from the representation.

The proposed approach provides a significant amount of flexibility, and there are many possible directions for investigation: initial transform/subband representation (e.g. DFT, DCT, LOT, wavelet); circumscribing ROS over which the initial representation is defined, and placement of the signal within the ROS; subset of vectors (which basis) to choose for the A-ROS signal. Furthermore, the selected vectors may be used for either "analyzing" or "synthesizing" the signal. In effect, we are creating a biorthogonal representation in which we explicitly choose *either* the analysis or the synthesis vectors. The flexibility in creating the representation can be summarized as follows:

1. Initial basis
   a. Initial transform/subband representation
   b. Superset ROS over which initial basis is defined
   c. Placement of signal within initial superset ROS
2. Selection of vectors (which basis)
3. Use of basis for analysis or synthesis

The goal is to exploit this flexibility in an appropriate manner to create a representation that is useful for the given application. In the next section we briefly consider the question of using the selected basis for analysis or synthesis.

3.3 Analysis versus Synthesis Basis

Selecting an appropriate subset of vectors corresponds to selecting a basis over the signal's ROS. If the basis is orthogonal, then it is used for both analyzing (transforming) and synthesizing (inverse transforming) the signal. However, in our representation the basis is biorthogonal and therefore there are two bases involved in representing a signal: the selected basis and its dual (inverse) basis. The selected basis may be used for either analyzing or synthesizing the signal, and its dual will implicitly be used for performing the other operation [1]. Let us briefly consider the merits of the analysis and synthesis approaches.

[1]The dual basis is said to be implicitly used, since it is never actually calculated. However, its effect is apparent when performing the other operations.

The analysis approach corresponds to selecting a subset of vectors for analyzing the given signal. Specifically, if the inner products between the selected vectors and the signal are linearly independent, then the signal is uniquely determined by the inner products. The analysis may also be viewed as sampling a subset of the coefficients. For example, assume that the signal is suitably extrapolated to fill the space over which the superset basis is defined. Then one can compute the transform over the superset space, and sample a subset of the coefficients. In this manner, the problem of selecting the vectors is equivalent to a sampling problem: what subset (sampling) of coefficients will uniquely represent the signal.

A number of issues arise with the analysis approach. Computing the coefficients requires that an extrapolation be performed to define the area outside the signal's ROS and within the superset space. Ignoring the extrapolation and only considering the vectors over the signal's ROS is equivalent to using a zero extrapolation. Depending on the application, however, a zero extrapolation may be inappropriate. For example, in image compression it is advantageous to have a smooth extrapolation that concentrates the energy into a small fraction of the coefficients. A zero extrapolation results in artificial discontinuities at the boundaries that will spread the signal energy throughout all the coefficients. The quality of the extrapolation can have a significant influence on the potential success of the representation. Different extrapolations produce different coefficient amplitudes, and determining the appropriate extrapolation for each signal appears to be a difficult problem.

In addition, the analysis approach explicitly determines the analysis vectors, and only implicitly determines the synthesis vectors (corresponding to the dual of the analysis vectors). Therefore reconstructing the signal from the coefficients requires either determining the synthesis vectors, or solving the linear system. This approach makes the reconstruction more complex.

The "synthesis" approach corresponds to selecting a subset of vectors for reconstructing or synthesizing the signal. The synthesis vectors are chosen, and then their amplitudes are computed. This is a natural approach because when performing a linear expansion of a signal, one is typically interested in the basic elements (synthesis vectors) that comprise the signal. The synthesis approach directly addresses this.

The synthesis approach does not require any extrapolation – the tricky issue of creating an appropriate extrapolation is avoided. In fact, the synthesis approach leads to (implicitly creates) an extrapolation of the reconstructed signal. The extrapolation is a function of the signal and the selected vectors, and may or may not be useful (e.g. smooth) depending on how the vectors were selected. In addition, selecting an appropriate subset of vectors for synthesizing the signal can lead to a smooth set of coefficient amplitudes. These amplitudes are typically more compression-friendly and possibly more useful for characterizing the signal properties.

Selecting the vectors for either analysis or synthesis implicitly determines the dual vectors for the other. However, the dual vectors may be very different and exhibit very different properties from the selected vectors. As is shown in Section 4.2, judiciously choosing the selected vectors so that they provide a given property can cause the dual vectors to provide another property. However, generally one has little control over the properties of the dual vectors. Therefore, based on a set of desirable properties for a given application, it is best to choose analysis or synthesis based on which is the most appropriate for providing those properties.

For many applications, the synthesis approach is more natural and intuitive – we are selecting the vectors (choosing the basis) with which to express the signal. The synthesis vectors can be chosen so that they provide some desirable properties. In addition, the synthesis approach concentrates the complexity at the analysis operation (e.g. at the encoder) and not at the synthesis operation (e.g. at the decoder), which is often a favorable distribution of complexity. Therefore, throughout the remainder of this work, we will focus on the selection of an appropriate subset of vectors for synthesizing the signal. However, much of the discussion also directly applies to the selection of a subset of vectors for analyzing the signal.

3.4 Summary

The proposed approach is conceptually quite straightforward, yet it leads to a rich set of possible representations. The desired representation (basis) for an A-ROS signal should be chosen so that it provides high performance with low complexity. The definition of high performance varies with the individual application, however it often corresponds to good approximation capability, i.e. the signals of interest can be accurately approximated with only a small fraction of the total number of coefficients. Low complexity corresponds to the selection of vectors being a simple function of the signal's ROS, the coefficient computation being relatively fast, and the signal reconstruction being relatively fast.

The proposed approach can be used with any initial transform/subband representation, e.g. DFT, DCT, LOT, wavelet, etc. In the next chapter, we focus on creating a wavelet-type representation. Wavelet-type representations illustrate both the elegance and simplicity of the general approach, and they also appear very promising for a number of practical applications. In particular, we develop two different wavelet representations that provide a number of desirable properties for signals with varying degrees of arbitrary supports.

Chapter 4

Wavelet-type Representations

Wavelet representations have been studied extensively in the literature. These multiresolution representations appear to provide high performance in various image processing applications, as well as in many other applications. Their success in representing images make them a promising choice for representing the arbitrarily shaped objects or regions within images. In addition, a wavelet representation where the vectors have local support appears to be a natural choice for representing signals with arbitrarily shaped regions of support. By chosing FIR filters, the local support of the wavelet vectors will typically result in only a small fraction of the vectors interacting with the boundary; most of the vectors will be completely inside or completely outside the signal's ROS, thereby simplifying the following steps.

We focus on the problem of determining a single-level wavelet decomposition for a signal with A-ROS. Once a single level can be processed, the approach can be recursively applied to any of the subbands to create a wavelet or wavelet packet type decomposition. The choice of a wavelet-type representation with FIR filters yields a relatively simple and clean incarnation of our proposed general approach. To recall, the proposed approach begins with a basis over a superset space that contains the given signal's ROS, and then selects a subset of vectors from the superset basis such that they provide a basis over the ROS. For example, in the case of a 2-D single-level wavelet transform (2×2-channel filterbank), all the vectors are periodic translations of four basic vectors: the LL, LH, HL, and HH synthesis filters. As a result, the ROS of the initial wavelet representation has minimal effect — it can be defined over the entire 2-D plane or only over a rectangular region, in both cases providing the same set of vectors. In addition, the placement of the signal within the initial ROS also is simpler with the wavelet, since the periodic placement of the wavelet vectors lead to a relatively small number of distinct situations. For example, for the 2-D one-level wavelet transform, there are only four distinct placements of the signal with respect to the vectors. These are significant advantages over global representations such as a global DFT or DCT, where every possible initial ROS corresponds to a completely different set of vectors and every placement of the signal within the initial ROS corresponds to a different situation. These properties greatly simplify the analysis and computational issues.

The success of this approach depends on overcoming two distinct problems. The first problem relates to determining whether there exists a subset of vectors from the basis over the superset space that will provide a basis with desirable characteristics for a signal with arbitrary support. The second problem relates to developing computationally efficient algorithms for selecting the subset of vectors and computing their coefficients. These questions can be summarized as:

- Does there exist a subset of vectors that provides a high-quality basis? What are the desirable and achievable properties of this basis? Is there a selection algorithm that guarantees this basis for every ROS, or for every ROS in a reasonable class of ROS's?
- How do we efficiently select the vectors based on the ROS? How do we efficiently compute the expansion coefficients for a given signal?

This chapter begins by providing an overview of creating a wavelet-type representation using our proposed general approach. Even within the framework of a wavelet-type representation, there are potentially a large number of possible representations that may be created which provide a variety of desirable properties. We present two different wavelet representations, each of which provides some interesting attributes, and discuss their relative merits. A number of examples are provided to illustrate the potential applications of these representations. Finally, the computational issues are examined.

4.1 Overview

We can both motivate the wavelet-based approach as well as illuminate some of its salient features by considering the problem of creating a basis for a 2-D signal with an arbitrarily shaped region of support. To simplify the analysis, we will focus on Daubechies orthogonal FIR (closest to linear phase) filters [12] and their use as *synthesis vectors*. However, the discussion extends straightforwardly to the use of the vectors for analyzing the signal, and for the case of biorthogonal filters. For brevity, we will refer to a signal defined over an arbitrarily shaped region of support as an AS-ROS signal.

Consider the problem of determining a single-level wavelet decomposition of an AS-ROS signal. Once a single level can be processed, the approach can be recursively applied to any of the subbands to create a wavelet or wavelet packet type decomposition. Assume that an AS-ROS signal in the 2-D plane and an initial wavelet transform defined over the entire 2-D plane are given. The wavelet vectors can be split into three groups: (1) the *interior vectors* which lie entirely inside the signal's ROS, (2) the *boundary vectors* which lie partially inside and partially outside the ROS, and (3) the *exterior vectors* which lie entirely outside the ROS. The three subspaces spanned by these groups of vectors will be referred to as V_{int}, V_{bnd}, and V_{ext} respectively. Let $V_{\text{AS-ROS}}$ be the signal space and $V_{\text{bnd-trunc}}$ be the subspace spanned by the boundary vectors when truncated to the AS ROS:

$$V_{\text{2D-plane}} = V_{\text{int}} \oplus V_{\text{bnd}} \oplus V_{\text{ext}} \tag{4.1}$$

$$V_{\text{AS-ROS}} = V_{\text{int}} \oplus V_{\text{bnd-trunc}} \,. \tag{4.2}$$

Determining a basis for $V_{\text{AS-ROS}}$ is equivalent to determining a basis for both V_{int} and $V_{\text{bnd-trunc}}$ — thus for orthogonal vectors the problem can be decomposed into two lower-dimensional, independent problems. All the interior vectors are selected, otherwise there will be a "hole" in the representation. An appropriate subset of boundary vectors are also selected to span $V_{\text{bnd-trunc}}$. Specifically, only the portion of each boundary vector within the ROS is relevant. Since the original filters are orthogonal, V_{int} is orthogonal to $V_{\text{bnd-trunc}}$, however the set of truncated boundary vectors are not mutually orthogonal. The problem reduces to selecting a set of boundary vectors that provide a basis for $V_{\text{bnd-trunc}}$. Note that while this is similar to [29, 30] where a basis for $V_{\text{bnd-trunc}}$ is explicitly constructed, the novelty of the proposed approach is that we create a basis by simply selecting or discarding each boundary vector, without modifying any of the vectors.

Given a 2-D AS-ROS signal and an initial wavelet representation defined over the entire 2-D plane, we select a subset of the vectors from the initial representation to provide a basis for the AS-ROS signal. Typically there are many possible subsets of vectors that lead to a basis, however some of these bases have better properties than others. We desire a selection algorithm that (1) guarantees a basis and (2) the basis is of high-quality for the desired application. In the following we consider some of the possible properties that an appropriately chosen basis may potentially provide.

4.1.1 Selecting the Vectors/Choosing a Basis: Desirable Properties

The wavelet decomposition of infinite-extent signals (e.g. infinite-length 1-D signals or 2-D signals defined over the entire 2-D plane) provides a number of important properties that lead to its usefulness and success. It is desirable to retain as many of these properties as possible within our representation for finite-size signals. The important properties provided by a wavelet transform of infinite-extent signals include:

- Perfect reconstruction
- Critical sampling
- Multiresolution
- Local support
- Polynomial accuracy
- Frequency selectivity
- Orthogonality
- Conditioning (stability)
- Fast/separable computations

The first four properties extend naturally to the proposed representation. Determining a basis is equivalent to providing critical sampling and perfect reconstruction. The representation is multiresolution since the vectors at one level can be expressed as a simple linear combination of vectors at a finer level. Local support is preserved since the vectors used in the representation are subsets of the initial vectors, which have local support. The last five properties are considerably more difficult to achieve. The possibility of preserving polynomial accuracy is unclear, as well as the frequency selectivity of the boundary vectors. Orthogonality is not possible, except for extremely special ROS's and filters, i.e. even if the initial wavelet transform over the 2-D plane is orthogonal, the resulting representation for a signal with arbitrary region of support will virtually always be biorthogonal.

The issue of conditioning does not typically arise for the popular discrete transform/subband representations. The conditioning or stability of a representation describes the change in the representation that results from a change in the signal. In many applications it is desirable for similar signals to have similar representations. For example, a constant signal and a constant signal with a small amount of noise should have similar representations. Likewise, it is desirable that a small change in the signal will lead to a small change in the signal's representation, and vice versa. For example, in compression, it is desirable that if the coefficients are quantized with a small error, the reconstructed signal will also exhibit a small error. When this is true, the representation is referred to as well-conditioned. A conventional orthonormal transform is perfectly conditioned, i.e. a change in the signal will result in a change of exactly the same energy in the representation. Since the vast majority of transform/subband representations considered in the literature are orthonormal, or close-to-orthonormal, the issue of conditioning typically does not arise and is not considered. However, conditioning should be examined for any biorthogonal representation, and is an important consideration for our proposed representation.

Furthermore, it is also important to consider the effects on the proposed representation of changes in the ROS over which the signal is defined. This issue of possible changes in the signal's support does not arise in conventional transform/subband representations of signals with convenient supports. It is desirable that a small change in the ROS will lead to a small change in the set of selected vectors, and also that a small change in the signal's amplitude over the ROS would lead to a small change in the amplitudes of the selected vectors. The representation should therefore be stable to changes in both the support of the signal and the signal's amplitude over the support.

The existence of computationally efficient algorithms has been a crucial element in the success of transform/subband approaches. The analysis of 1-D signals has been facilitated by fast 1-D algorithms, and 2-D signals with rectangular support can be efficiently processed by separably applying the 1-D algorithm along each of the dimensions. For many M-D problems, the ability to perform separable processing (or other very low complexity methods) has been a critical enabling factor for their solution. Thus, the proposed approach requires efficient algorithms for selecting the subset of vectors for a given ROS, and computing their coefficients. While the non-rectangular supports appear to prohibit the use of separable processing, some type of computationally efficient algorithm is necessary to maximize the usefulness of the proposed approach.

An important additional property for our work, that does not arise for infinite-extent signal, is flexibility in the range of possible supports that may be represented. One basic and important class of signal supports are those with contiguous supports. These include finite-length 1-D signals or 2-D/M-D signals with arbitrarily shaped regions of support, e.g the arbitrarily shaped 2-D objects within an image. For these signals, it is desirable to have flexibility in the possible lengths of the 1-D supports and in the shapes of the 2-D/M-D supports. In many applications it may be difficult or impossible to impose constraints on the shape of the support, so the more freedom the better. Furthermore, in some applications the shape may not only be completely arbitrary, but the support may even be discontiguous, or contain gaps, etc. Therefore, to be applicable to the largest range of problems, it would be desirable to have a representation that provides complete freedom in the possible ROS's, i.e. one that is applicable to any arbitrary support. While a representation that is applicable to all possible arbitrary supports may not provide high quality for all supports, even the capability to apply it without requiring a "check" on the support may lead to important simplifications.

We have briefly discussed some of the properties that we would like the representations to provide. In the next two sections we present two possible wavelet representations, each providing a different set of properties as shown in Figure 4.1. The first representation is theoretically interesting, but has some practical limitations, while the second representation appears to have practical significance. These representations provide two very different sets of properties. Of course, there may also exist other representations that provide other sets of useful properties.

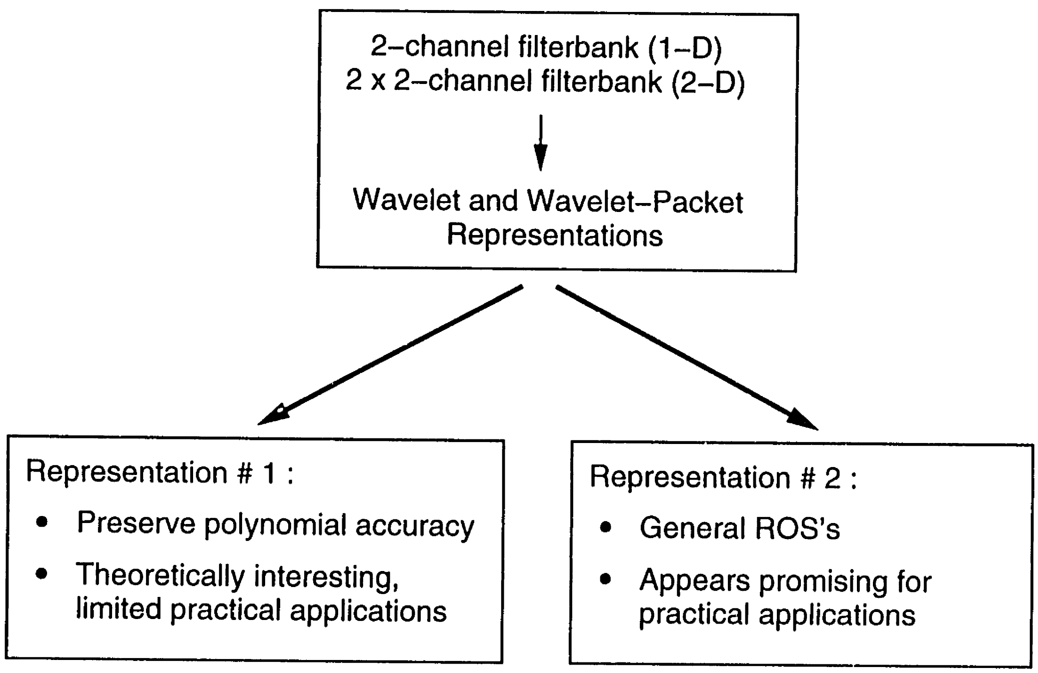

Figure 4.1: Synopsis of Sections 4.2 and 4.3: The focus is on creating a single-level wavelet transform for an A-ROS signal that provides certain properties. Once a single level can be processed, the approach can be recursively applied to any of the subbands to create wavelet or wavelet packet type decompositions. Section 4.2 presents representation # 1 which is designed to preserve polynomial accuracy. Section 4.3 presents representation # 2 which is designed to provide a basis over any arbitrary ROS. Many other potentially useful representations may also exist.

4.2 Representation # 1: Preserving Polynomial Accuracy

Polynomial accuracy is a fundamental approximation property of the wavelet transform that leads to many important theoretical and practical applications. Therefore it is natural, and also very important, to examine if the polynomial accuracy property can be preserved with the proposed approach. In this section we design a representation which preserves the polynomial accuracy of the wavelet transform. We begin by developing a simple condition on the selection of vectors such that the resulting basis preserves the polynomial accuracy. We then examine the consequences of this condition on representing 1-D and 2-D signals.

The notion of polynomial accuracy relates to the ability to exactly reproduce polynomials up to a certain order. In the case of the wavelet transform of an infinite-length 1-D signal, polynomial accuracy can be defined in the following manner.

Polynomial Accuracy If the LP/HP filter pair have polynomial accuracy of order p, then:
(1) Any polynomial signal of order less than p will have all its information concentrated in the lowpass subband, and all the other subbands will be zero [1].
(2) The lowpass subband will also be a polynomial of the same order.

To be specific, consider a single-level wavelet transform of an infinite-length 1-D signal. For Daubechies orthogonal L-tap filters, the LP vectors can exactly reproduce any polynomial of order less than $p = \frac{L}{2}$ and the HP vectors have p vanishing moments. If the signal to be transformed is any polynomial of order less than p, all the HP subband coefficients will be zero. All of the signal information is therefore concentrated in the LP coefficients, and the polynomial signal can be exactly reconstructed using only the LP subband. In addition, the amplitudes of the LP subband coefficients correspond to another polynomial of the same order. Therefore, computing another level of the wavelet transform (iterating a 2-channel filterbank on the LP subband) will further concentrate the signal information: the LP subband will contain all the information, while the bandpass and highpass subbands will be zero. This follows for any number of levels of the transform. As a side note, the typical definition of polynomial accuracy only refers to property (1) above. Property (2) follows for wavelet transforms of infinite-length signals.

The key to the polynomial accuracy property is that the LP synthesis vectors can exactly reproduce polynomial signals, i.e. the space of polynomial signals of order less than p is within the span of the LP synthesis vectors. Therefore to retain polynomial accuracy in our new repre-

[1] p is the order of the first (lowest order) polynomial which can not be exactly reproduced using only the lowpass subband.

sentation, the selected LP vectors must be able to exactly reproduce polynomial signals [2]. Since the proposed representation corresponds to simply selecting or discarding vectors, polynomial accuracy can be preserved by selecting and including in the basis all the LP vectors that overlap the signal's ROS. The resulting selection algorithm can be summarized as follows:

Selection Algorithm for Preserving Polynomial Accuracy

1. Select all the interior vectors (both lowpass and highpass)
2. Select all the boundary lowpass vectors
3. Select an appropriate set of boundary highpass vectors to complete the basis

The proposed representation can preserve the polynomial accuracy for a signal with a given ROS if it is possible to include in the representation (basis) all the LP vectors that overlap the signal's ROS. As will be discussed shortly, for some supports there are more LP vectors that overlap the support and must be selected then there are samples in the ROS – thus preventing the creation of a critically sampled representation. Therefore this representation is not applicable for every possible support. It is applicable for (contiguous) arbitrary-length 1-D signal (of a minimum length) and 2-D/M-D signals with (contiguous) arbitrarily shaped regions of support with a constraint on the curvature (local shape) of support. The representation is also applicable to a signal whose support consists of a number of separate regions, as long as the regions are sufficiently far apart.

This selection algorithm leads to a number of important consequences. Consider the case of a finite-length 1-D signal and assume that it is possible to retain all the LP vectors that overlap the signal's ROS (this point will be examined shortly). If the signal is a polynomial (up to order $p-1$) then it can be exactly reconstructed from the LP vectors, and the LP coefficients will be a polynomial of similar order. For example, if the signal is a constant (zero-order polynomial), then the LP coefficients will also be constant. Also, if the signal is a cropped version of an infinite-length polynomial signal, then the LP coefficients for our representation will have the same amplitudes as those for a conventional wavelet transform of the infinite-length signal. In addition, consider synthesizing a signal from a representation: beginning with a LP subband that is a polynomial and all other subbands that are zero, the corresponding signal will be a polynomial of similar order. These properties illustrate some of the distinguishing features of our proposed approach, and they will be revisited shortly when we compare this approach to previous approaches.

[2] To be precise, the set of vectors that correspond to the "lowpass synthesis vectors" for the representation must be able to exactly reproduce polynomials. For example, a HP boundary vector may be selected and used as a "LP vector".

A second consequence relates to what conditions are necessary for polynomial accuracy to be preserved. Notice that all of the LP vectors are used to reconstruct a polynomial signal, and that there is a unique linear combination of the LP vectors that reconstructs any polynomial signal. As a result, if any one of the LP vectors that overlaps the signal's ROS is not selected, the resulting representation will not preserve the polynomial accuracy. That is, if any of these LP vectors are not included in the basis, then the remaining LP vectors (in the basis) do not span the space of polynomials, and the LP vectors can not exactly reproduce those polynomial signals. Therefore, a simple necessary and sufficient condition for the representation to preserve polynomial accuracy and be critically sampled is for all the LP vectors that overlap the signal's ROS to be selected, and that those vectors be linearly independent over the ROS. If the LP vectors are not linearly independent over the signal's ROS, then too many vectors must be selected to span the signal space and the resulting representation will be overcomplete.

A final interesting consequence is that the resulting HP analysis vectors for the proposed representation will have p vanishing moments. This follows by biorthogonality, since if the synthesis LP vectors span the space of polynomials, the analysis HP vectors must be orthogonal to this space. This may also be viewed from another point of view: if the analysis HP vectors were not orthogonal to this space, their coefficients would be nonzero, leading to the synthesis HP producing a nonzero output. Throughout this work, it is important to consider the properties of the biorthogonal representations that we design. By selecting the synthesis vectors for the representation, one is explicitly designing the synthesis operation. The analysis operation follows only implicitly (from the inverse of the synthesis), thus it is not immediately apparent that the analysis vectors will exhibit any particularly attractive properties. However, by proper choice of the synthesis vectors, the analysis vectors can also provide desirable properties.

4.2.1 Representing 1-D Arbitrary-Length Signals

The proposed representation corresponds to retaining all the interior vectors, all the boundary lowpass vectors, and possibly some boundary highpass vectors to complete the basis. An interesting consequence of this approach, is that for the problem of arbitrary-length 1-D signals and Daubechies orthogonal filters, retaining the interior vectors and only the lowpass boundary filters is sufficient to guarantee a basis. That is, no highpass boundary filters are required to complete the basis. This leads to a very simple selection scheme and potentially reduced computational requirements.

Consider one level of the decomposition. Once we show that the proposed representation provides a basis for any single level, the case of multiple levels follows straightforwardly.

Theorem 4.2.1. *Given a 2-band filterbank composed of Daubechies L-tap orthogonal filters, and an arbitrary-length 1-D signal of minimum length L-1, selecting all the interior vectors and the lowpass boundary vectors will provide a basis for the signal. Furthermore, this basis will preserve the polynomial accuracy property, and for a polynomial signal of order less than* $\frac{L}{2}$ *the LP subband coefficients will also be a polynomial of the same order.*

Overview of proof: The proof is straightforward, though detailed, and therefore is given in the appendix. However, a brief overview is presented to provide intuition for this approach. To begin with, all the interior vectors must be selected. At each boundary, half of the boundary vectors should be selected. Each pair of lowpass and highpass Daubechies vectors are exactly aligned, therefore selecting all the lowpass boundary vectors corresponds to selecting half of the total number of boundary vectors. The lowpass boundary vectors are linearly independent of the interior vectors (actually orthogonal) and also linearly independent of each other since they have staggered (offset) supports. The interior vectors and the boundary lowpass vectors therefore provide a basis for the arbitrary-length signal. The polynomial accuracy property is preserved since all the LP vectors that overlap the signal's support have been included in the basis. The minimum length constraint results from two issues. The first is that the signal length must be $\geq L - 2$ or else a vector may be a boundary vector for both boundaries, i.e. the boundaries are no longer decoupled. A more stringent issue is that a shorter length signal may have more LP vectors that overlap the signal (and therefore must be selected) than samples in the signal, thereby producing an overcomplete representation. This second issue forces the length constraint. This completes the overview of the proof.

As a side note, a slightly shorter length signal can also be represented with the proposed framework, assuming freedom in the placement of the signal with respect to the phase (placement) of the vectors. However, while placing the signal in an advantageous position with respect to the vectors is a reasonable option for 1-D signals, it is probably impractical or impossible for most 2-D or M-D arbitrarily shaped supports. Therefore, we will not consider the option of choosing the relative placement of the signal and vectors. Our goal is to determine an approach that will work for all possible placements.

The proposed framework leads to a basis for any arbitrary-length 1-D signal with length $\geq (L - 1)$. This includes even- and odd-length signals, as opposed to some of the previous approaches that are restricted to even-length signals. The signal can also be of infinite-length, either one-sided (causal or anti-causal) or two-sided. In the case of two-sided infinite signal, the proposed approach corresponds to the conventional approach for processing infinite-length signals.

The representation can be straightforwardly extended to multiple levels: one constructs another level simply by applying the same decomposition to the lowpass subband. The only constraint is that the number of lowpass coefficients is $\geq (L-1)$. In general, any subband can be further decomposed as long as the number of coefficients in that subband is $\geq (L-1)$. This enables the creation of wavelet-packet type representations for arbitrary-length 1-D signals.

Corollary 4.2.2. *Given a 2-band filterbank composed of Daubechies L-tap orthogonal filters and an arbitrary-length 1-D signal of minimum length L-1, the signal can be recursively decomposed by the 2-band filterbank as long as the number of lowpass coefficients at each scale is at least L-1. In addition, any subband can be decomposed as long as the number of coefficients in the subband is at least L-1. In this manner wavelet and wavelet-packet representations may be constructed. These representations are critically sampled perfect reconstruction schemes that preserve the polynomial accuracy property.*

An interesting point is that the existing wavelet methods for arbitrary-length 1-D signals do not fully preserve the polynomial representation properties. For example, while Cohen, Daubechies, and Vial's (CDV) preconditioned representation [29] compacts all the polynomial information into the lowpass subband, the subband is not a polynomial of similar order, i.e. there are some "wiggles" at each boundary. This will result in a (smooth) polynomial signal having a (non-smooth) non-polynomial set of lowpass coefficients, as illustrated for a constant and a linear signal in Figure 4.2. In effect, a smooth signal will have a non-smooth representation. However, the current proposal does fully preserve these properties, e.g. a constant signal will lead to a constant lowpass subband with all other subbands equal to zero. These effects are illustrated in Figure 4.2. Therefore, with the proposed representation, a smooth signal will have a smooth representation, and a smooth representation will correspond to a smooth signal. These properties are beneficial for a number of areas, including interpretation and compression.

The proposed approach produces more lowpass coefficients at any given level than the other approaches. CDV's approach is similar to the conventional wavelet transform in that it decomposes the signal into an equal number of lowpass and highpass basis vectors (and coefficients) at each level. By selecting all the lowpass vectors possible, our approach produces an unequal number of lowpass and highpass vectors at each level. This makes the decomposition in a sense "unbalanced", and may slightly increase the difficulty of indexing. However, it is also advantageous in that there are fewer restrictions on the signal length – approaches which split a signal into an equal number of LP and HP coefficients are limited to signals whose length is a multiple of 2^k, where k is the desired number of levels in the decomposition. Our approach can be applied to any signal length $\geq N-1$, e.g. even or odd length.

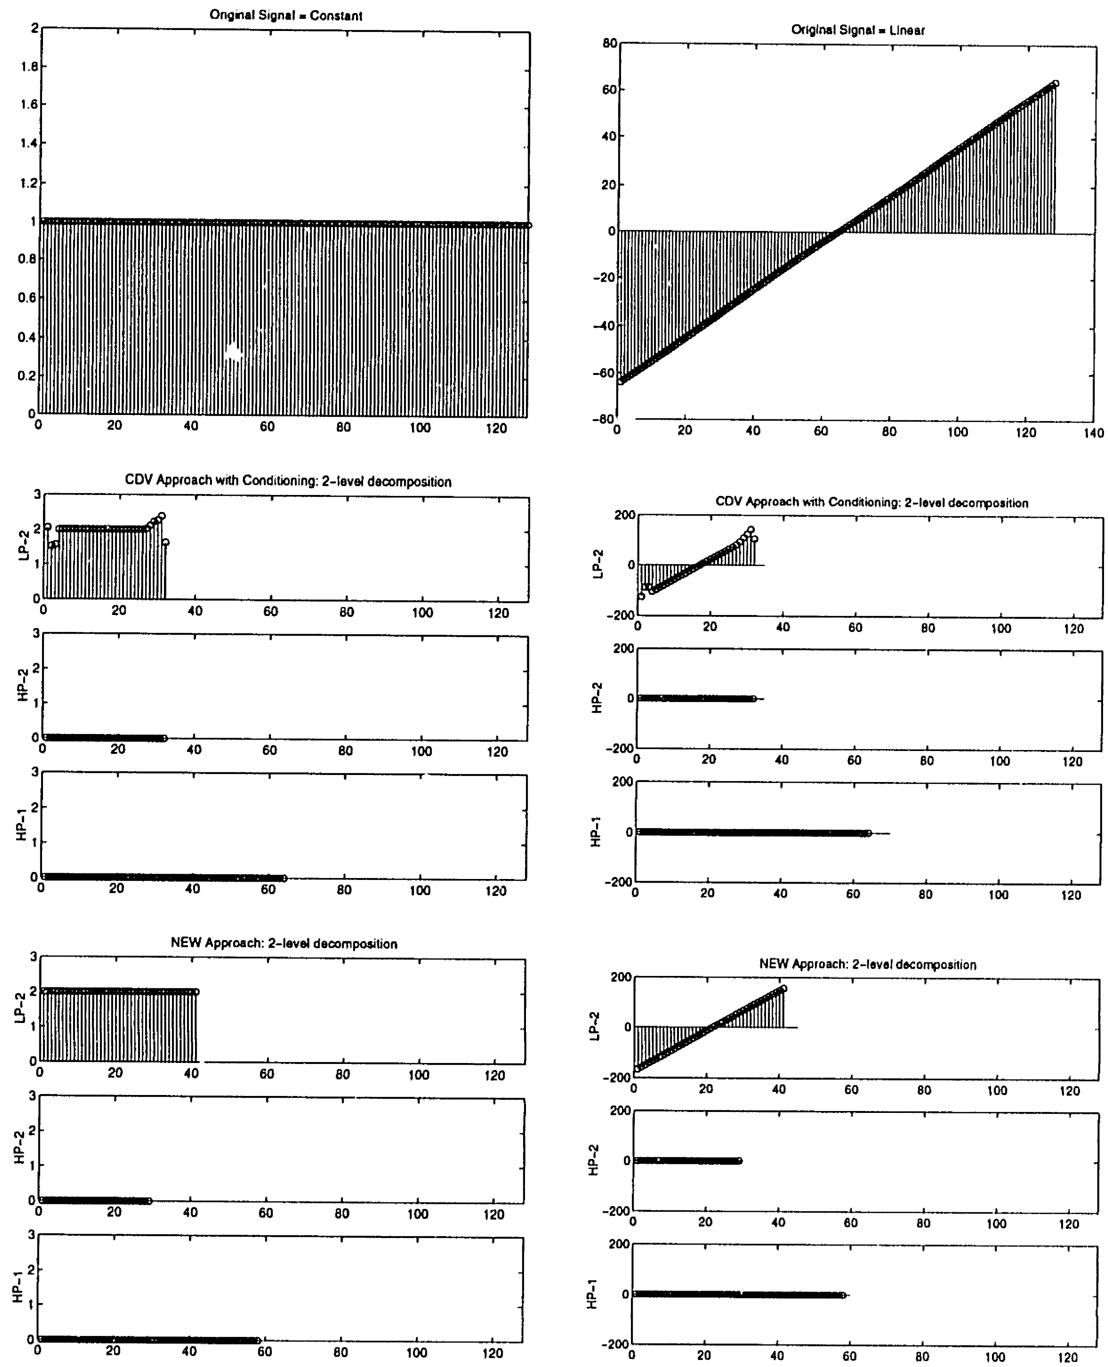

Figure 4.2: The constant signal (left column) and linear signal (right column) are represented using CDV with conditioning (middle row), and the proposed approach (bottom row). A 2-level decomposition (i.e. lowpass, bandpass, and highpass subbands) is created in each case using Daubechies 14-tap orthogonal (closest to linear phase) filters. Both CDV with conditioning and the proposed approach concentrate all the signal information into the LP subband. However, even though the signals correspond to smooth polynomials, the lowpass coefficients of CDV with conditioning are non-smooth near the boundaries, while the proposed approach has smooth (constant and linear) lowpass coefficient amplitudes.

A possible additional benefit of selecting all the LP vectors is that it may potentially lead to a smoother extrapolation for smooth signals. Extrapolation is a basic problem in signal processing and many applications involving signals with arbitrary supports may require an appropriate extrapolation. The proposed approach may perform better than the previous approaches in terms of forming a smooth extrapolation for smooth (polynomial) signals. An intuitive explanation is as follows. In all the proposals, the boundary vectors in the representation are a linear combination of the original boundary basis vectors (in CDV's proposal, an orthogonalization relates the two, and in our proposal there is simply a binary selection process). A simple and natural approach for performing the extrapolation is to "attach the tails back on" to the boundary vectors, that is, reconnect their portion outside the signal's ROS. In the case of a smooth polynomial signal, the transformed signal only has nonzero lowpass coefficients. In CDV's approaches, the LP boundary vectors correspond to a linear combination of both LP and HP vectors. Therefore, the smooth signal is represented at the boundary as a linear combination of smooth LP vectors and not-as-smooth HP vectors and attaching the tails for each potentially leads to a non-smooth extrapolation. In our approach, a polynomial signal is represented using only LP vectors at the boundary. Therefore, attaching the tails extrapolates the LP signal at the boundary using only LP vectors. This may potentially lead to a smoother extrapolation.

This proposed representation has two disadvantages. First, the representation becomes increasingly ill-conditioned with longer filter lengths. Specifically, as the filter length increases, and therefore the boundary vector tails used to express the signal near the boundary become smaller in amplitude, the representation will become increasingly ill-conditioned [3]. A small change in the signal amplitude near the boundary can therefore produce a large change in the coefficient amplitudes, i.e. similar signals can have very different coefficient amplitudes. This may limit the wavelet filters used, as well as limit the potential applications. A second disadvantage relates to the possible ROS's that can be represented. This proposed representation is limited to 1-D signals with contiguous, finite-length supports of minimum length of $(L-1)$, where L is the filter length. The minimum length constraint arises because otherwise the number of lowpass vectors that overlap the signal and are selected exceed the number of samples in the signal. It is natural to filter a signal with a filter that is shorter than the signal length, therefore the constraint in 1-D is highly reasonable. However, the extension of this constraint to M-D is much more restrictive.

To summarize the 1-D results, the proposed representation provides a simple method for creating a critically sampled perfect reconstruction wavelet transform for arbitrary-length 1-D signals. In the case of Daubechies filters, selecting all the interior vectors and only the lowpass boundary vectors both guarantees a basis and preserves the polynomial accuracy property.

[3] Of course, the ill-conditioning results from both the angles between the vectors and the lengths of the vectors.

4.2.2 2-D/M-D Signals with Arbitrarily Shaped Regions of Support

The proposed approach can in principle be extended straightforwardly to 2-D and general M-D problems. Specifically, a basis can be determined for a 2-D or M-D signal with arbitrarily shaped region of support such that the polynomial accuracy is preserved. The discussion here will be limited to 2-D signals, however the analysis extends straightforwardly to general M-D signals.

We focus on the case where the initial 2-D transform (over the superset space) is given by a separable application of 1-D transforms. The initial 2-D basis then corresponds to a tensor product of 1-D bases, and therefore inherits the properties of the 1-D bases oriented along horizontal and vertical directions. In particular, if the 1-D bases have polynomial accuracy p, then a 2-D basis will have polynomial accuracy (p, p). That is, the p^2 2-D polynomials

$$n_1^r n_2^s, \quad 0 \leq r \leq p-1 \text{ and } 0 \leq s \leq p-1 \tag{4.3}$$

defined over the entire 2-D plane can be perfectly reconstructed using only the 2-D lowpass coefficients. All the information (energy) for all polynomials of order (p, p) will be concentrated into the lowpass coefficients, and the lowpass coefficient amplitudes will be a polynomial of the same order.

One level of a separable wavelet-transform for a 2-D signal will lead to four types of wavelet vectors: lowpass-horizontal/lowpass-vertical (LL), lowpass-horizontal/highpass-vertical (LH), highpass-horizontal/lowpass-vertical (HL), and highpass-horizontal/highpass-vertical (HH). To preserve the polynomial accuracy property in 2-D, all the LL vectors that overlap the signal's support must be selected. The selection algorithm therefore corresponds to (1) select all the interior vectors, (2) select all the LL boundary vectors, and (3) select an appropriate number of LH, HL, and HH vectors to complete the basis.

For a finite-length 1-D signal with Daubechies filters, selecting all the boundary lowpass vectors is sufficient to both preserve polynomial accuracy and provide a basis. However, in 2-D/M-D, selecting all the lowpass boundary vectors is sufficient to preserve polynomial accuracy, but it does not provide a basis. Some intuition for this can be built in a straightforward manner. Given a 2-D signal with support in the right-half-side of the 2-D plane, half of the total number of boundary vectors must be retained to provide a basis (this is a simple extension of the 1-D problem). If the boundary region corresponds to a corner (e.g. of a signal with square support) then less then half of the boundary vectors for the corner must be selected. If the boundary region corresponds to the knee of a signal with "L" support, then more than half of the vectors must be selected at the knee. Intuitively, about half of the boundary vectors must be selected over each portion of the boundary to provide a basis. In 1-D, selecting all the lowpass boundary vectors corresponds to selecting half of the boundary vectors. In 2-D, since the signal is decomposed into four subbands, selecting all the LL boundary vectors corresponds to selecting only one-quarter of the total number of boundary vectors. Therefore, additional boundary vectors must be selected to complete the basis.

The problem of selecting additional boundary vectors leads to two issues: (1) which subset of vectors is the best to select, i.e. what is an appropriate performance criteria for the selection, and (2) how do we guarantee that the selection provides a basis.

There may be many possible selections of the additional boundary vectors that would lead to a basis, as well as many possible criteria for performing the selection. Possible criteria for selecting the additional boundary vectors include: conditioning of the resulting basis; symmetry issues; effects on extrapolation; in addition, there may be a preference for selecting from one subband over another, e.g. one may want to consider the shape (ROS) of the resulting subbands (which correspond roughly to downsampled versions of the original signal's ROS).

The problem of proving that the selected boundary vectors are linearly independent is considerably more difficult in 2-D than in the 1-D case. For example, in 1-D all the lowpass boundary vectors were linearly independent since they had staggered supports. In 2-D, the lowpass boundary vectors do not necessarily have staggered supports. Furthermore, in 2-D the supports of the additional boundary vectors that must be selected may exactly align with the supports of the lowpass vectors, e.g. the Daubechies LL, LH, HL, and HH filters exactly overlap each other. These issues greatly complicate the problem of determining if a subset of vectors are linear independent and provide a basis.

We have developed a selection algorithm that appears to lead to a basis for Daubechies orthogonal filters and for all applicable supports. The issue of which supports are applicable will be discussed shortly, but first we present the approach. A pleasing aspect of this selection is that it only involves separately examining the signal's ROS over each local area. For example, each set of four $L \times L$-sample Daubechies vectors LL, LH, HL, and HH exactly overlap each other and cover an $L \times L$ sample area. Examining the signal's support in each $L \times L$ area appears to be sufficient to determine the vectors to select. Consider Φ to be the $L \times L$ matrix that describes the overlap between the signal's support and the $L \times L$ support of a specific group of four vectors. Specifically, Φ is a bitmap where each entry is either 1 or 0, depending on whether or not the signals's support overlaps that sample. Let k describe the total amount of overlap between the signal's support and the vectors, given by the sum of all the entries of Φ. The selection algorithm can then be summarized as:

1. Complete overlap ($k = L^2$): Select all four vectors
2. No overlap ($k = 0$): Select none of the vectors
3. Partial overlap ($0 < k < L^2$): Select a portion of the vectors
   (a) Select the LL vector
   (b) If the L samples in the first column of Φ *or* the L samples in the last column are occupied: Select the LH vector
   (c) If the L samples in the first row of Φ *or* the L samples in the last row are occupied: Select the HL vector

This heuristic selection algorithm appears to lead to a basis for all applicable supports. An example is given in Figure 4.3. It is an interesting open question to determine if this algorithm always leads to a basis for all Daubechies filters and for all applicable supports. However, this proposed representation has a number of disadvantages in 2-D and especially for general M-D problems that may limit its usefulness except under special circumstances. Because of these disadvantages, we have diverted our efforts in a more practically promising direction that leads to the second class of representations which we present in the next section [4].

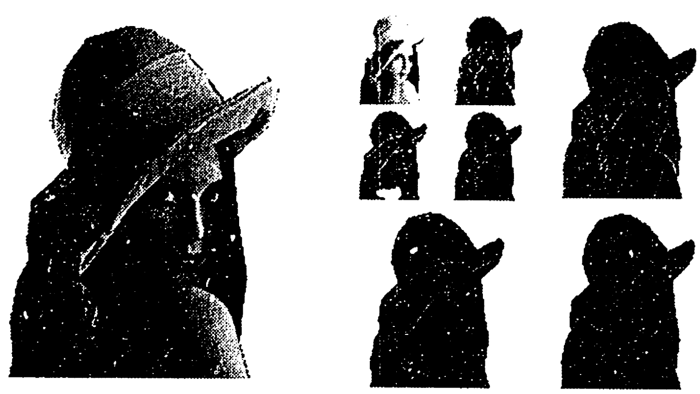

Figure 4.3: A two level decomposition of lenna using representation #1 and Haar filters. The use of Haar filters avoids the difficulties of longer filters. The signal can have any arbitrarily shaped ROS, and in fact can have any arbitrary ROS – this is not true for the use of longer filters. In addition, the resulting representation is well-conditioned. Longer filters lead to constraints on the shape of the ROS and the representation typically becomes increasingly ill-conditioned with the filter length.

[4]Some of the ideas and formulations developed as part of the second class of representation may be useful in analyzing the above selection algorithm.

This proposed representation has two disadvantages. First, the representation becomes increasingly ill-conditioned with longer filter lengths. Intuitively, this occurs because some of the selected boundary vectors correspond to very short truncated "tails" of the original LP boundary vectors. A small change in the signal amplitude near the boundary can therefore produce a very large change in the coefficient amplitudes, i.e. similar signals can have vastly different coefficient amplitudes. This issue is much worse for 2-D then 1-D since the 2-D condition number is approximately the square of the 1-D condition number [5]. This may limit its usefulness to special applications.

A second disadvantage is that the set of possible ROS's that can be represented is severely limited. For example, in 1-D the signal must have a minimum length of $(L-1)$; otherwise the number of lowpass vectors that overlap the signal and are selected exceed the number of samples in the signal. A similar concept of minimum length exists in 2-D: any features on the boundary with a horizontal or vertical size of less than $(L-1)$ (e.g. a bump or divet) will result in too many LP boundary vectors. It is natural to filter a signal with a filter that is shorter than the signal length, therefore the constraint in 1-D is reasonable. However, the constraint in 2-D or general M-D places a severe restriction on the possible ROS's that may be processed. Therefore, while the polynomial accuracy property is theoretically appealing, the disadvantages suggest that this representation may have limited practical usefulness.

4.3 Representation # 2: Arbitrary ROS's

Creating a representation that preserves the wavelet's polynomial accuracy is theoretically appealing. However the usefulness of polynomial accuracy for general signal/image processing is less evident. The ability to accurately represent a constant signal using only the LP subband (no DC-leakage into HP subbands) is very important, but higher order polynomial accuracy may be of limited value. More important properties include good conditioning, good frequency responses for the boundary filters, and possibly most important, flexibility in the possible ROS's that may be processed. In this section we propose a representation that provides these capabilities.

Our representation is based on the idea of *selecting the vector centered at each sample in the signal's ROS*. To begin, let us clarify what is meant by the vector centered at each sample. Consider the basis vectors of a single-level wavelet transform for a 1-D signal. For an odd-length vector, the center of the vector is straightforward. If both the LP and HP vectors have odd-length and their centers are located at adjacent samples (e.g. the typical biorthogonal linear phase filterbank) then

[5]Note that the condition number describes the worst case change [68].

there is a simple mapping between each sample in the signal's ROS and the vector centered at that sample, as shown in Figure 4.4. An even-length vector has two central taps to choose from, and the center can be taken to be the largest of the two taps [6]. If the LP and HP vectors are related by the alternating flip property (e.g. orthogonal and semiorthogonal filters) this works out very conveniently, as the LP vector is centered at one sample, and the HP at the other. In general, it is beneficial for the chosen center tap to be the largest tap of the filter. For typical filters, our notion of centering as described above provides this property. This follows since most filters, especially in image processing, have their energy concentrated toward the center of the filter. For filters where the largest tap is not located in the center (e.g. minimum phase filters) it may be desirable to center on the largest tap.

In all of the above cases, there is a straightforward and unique mapping between the signal's ROS and the vectors to select. For example, assume the LP vectors are centered at even sample locations and the HP vectors are centered at odd sample locations. Then for every even sample in the signal's ROS, the LP vector centered at that sample will be selected, while for every odd sample in the signal's ROS, the HP vector centered at that sample will be selected. This selection scheme extends straightforwardly to separable wavelet bases defined over 2-D or general M-D spaces.

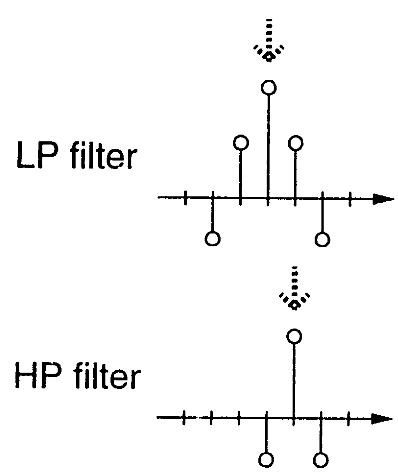

Figure 4.4: The center taps for the 5/3-tap biorthogonal lowpass/highpass filter pair are given by the arrows.

This proposed approach for determining a basis selects for each sample in the signal's ROS the corresponding vector centered at that sample. In this manner we are guaranteed to select the correct total (global) number of vectors required (equal to the total number of samples in the ROS). In general, it is relatively difficult to determine the number of vectors required in any local boundary area. However, this approach provides a reasonable estimate of the number of vectors

[6] As will be shown shortly, choosing the largest tap leads to a better-conditioned representation.

required over any local area. In fact, for an appropriate choice of the initial vectors, all the selected vectors are linearly independent — providing a basis over the ROS. Furthermore, as will be shown, for an appropriate choice of the initial vectors, this approach provides complete freedom in the possible ROS's that may be represented.

Wavelet and filterbank design is often favorably tackled in the frequency domain. However, working with arbitrary supports (similar to an arbitrary sampling problem) makes a time-domain analysis more natural. Our problem can be expressed quite advantageously in matrix form. Consider the 1-D case to simplify the discussion and notation. Assume we are given a signal with an arbitrary region of support, and an initial transform/subband representation defined over a larger space that contains the given signal. Suppose the signal's support contains a total of M points, and the initial transform/subband representation is defined over a space of N points, where $N \geq M$. Let the transform/subband (synthesis) representation be expressed by the $N \times N$ matrix T where the columns of T contain the basis vectors.

The conventional transform/subband synthesis of a signal is given by

$$T \cdot c = f \tag{4.4}$$

where

$$
\begin{aligned}
T &= \text{Transform synthesis matrix} \quad (N \times N) && (4.5)\\
c &= \text{Coefficients} \quad (N \times 1) && (4.6)\\
f &= \text{Signal} \quad (N \times 1). && (4.7)
\end{aligned}
$$

The matrix T is invertible, and its columns contain the basis vectors of the representation. If the signal f has full support ($M = N$), then a linear combination c of all the basis vectors of T will reconstruct the signal. If the signal has smaller support ($M \leq N$), then only a subset of M out of the N vectors is required to reconstruct the signal. To make this explicit, we examine the matrices and submatrices that arise.

The matrix equation in Equation 4.4 has the following form

$$\begin{bmatrix} \quad\quad\quad\quad\quad\quad \\ \\ \\ \\ \end{bmatrix}_{N\times N} \begin{bmatrix} \quad \\ \\ \\ \\ \end{bmatrix}_{N\times 1} = \begin{bmatrix} \quad \\ \\ \\ \\ \end{bmatrix}_{N\times 1} . \tag{4.8}$$

If the signal f is only defined over $M \leq N$ points, then it can be collapsed from an $N \times 1$ to an $M \times 1$ column vector as follows (the undefined points are marked $\times$):

$$\underset{N\times 1}{\begin{bmatrix} \\ \\ \\ \\ \\ \end{bmatrix}} \longrightarrow \underset{N\times 1}{\begin{bmatrix} \times \\ a \\ b \\ \times \\ c \end{bmatrix}} \longrightarrow \underset{M\times 1}{\begin{bmatrix} a \\ b \\ c \end{bmatrix}} \tag{4.9}$$

Only the portion of each vector that overlaps the signal's support is relevant for representing the signal, i.e. for every sample in the signal's support, only the corresponding row of T is relevant. This may be signified by retaining the rows of T that correspond to the support of the signal and discarding the remaining rows, leading to an $M \times N$ linear system. The transform matrix can be collapsed accordingly, leading to a new matrix equation of the following form:

$$\underset{N\times N}{\begin{bmatrix} \times & \times & \times & \times & \times \\ & & & & \\ & & & & \\ \times & \times & \times & \times & \times \\ & & & & \end{bmatrix}} \underset{N\times 1}{\begin{bmatrix} \\ \\ \\ \\ \\ \end{bmatrix}} = \underset{N\times 1}{\begin{bmatrix} \times \\ \\ \\ \times \\ \\ \end{bmatrix}} \longrightarrow \underset{M\times N}{\begin{bmatrix} & & & & \\ & & & & \\ & & & & \end{bmatrix}} \underset{N\times 1}{\begin{bmatrix} \\ \\ \\ \\ \\ \end{bmatrix}} = \underset{M\times 1}{\begin{bmatrix} \\ \\ \\ \end{bmatrix}} \tag{4.10}$$

The resulting system $T_{M\times N} \cdot c_N = f_M$ is underdetermined, i.e. there are more transform coefficients (vectors) than signal samples. A critically sampled solution would require T and c to be of size $M \times M$ and $M \times 1$, resulting in the matrix structure shown below.

$$\underset{M\times M}{\begin{bmatrix} & & \\ & & \\ & & \end{bmatrix}} \underset{M\times 1}{\begin{bmatrix} \\ \\ \\ \end{bmatrix}} = \underset{M\times 1}{\begin{bmatrix} \\ \\ \\ \end{bmatrix}} \tag{4.11}$$

Therefore, determining a basis over the ROS is equivalent to determining M linearly independent columns out of the N columns of $T_{M\times N}$, such that the resulting $M \times M$ submatrix is nonsingular.

The problem can therefore be summarized in the following manner. Beginning with an invertible $N \times N$ matrix T, the support of the signal dictates that M rows of the matrix are retained while the others are discarded, leading to an $M \times N$ matrix of rank M. The goal is then to select M linearly independent columns of the $M \times N$ matrix so that the resulting $M \times M$ submatrix is nonsingular. There exists at least one subset of M columns that lead to a basis, otherwise the original T would not be invertible (would not enable perfect reconstruction). Typically, there are many subsets of M columns that lead to a basis. The goal is an algorithm which given any support *simply selects* a subset of vectors that is *guaranteed to be a basis over the support*. The algorithm should not require any check of the linear independence of the vectors. Furthermore, the basis should be of "high" quality in order for it to be useful for applications.

For a 2-channel filterbank (one level of a wavelet transform), the matrix T corresponds to a 2×2 block Toeplitz matrix. The analysis can be simplified considerably by shifting the filters such that the center tap of the LP filter's impulse response occurs at time sample zero and the HP filter at time sample one. Shifting the filters in this manner results in the vectors being aligned such that the i^{th} vector is centered at sample i and the $(i+1)^{th}$ vector is centered at sample $i+1$, etc. While this is an unconventional expression of the matrix for a 2-channel filterbank, it does not alter the problem in any way, and it does lead to the following important consequence: the selected/retained rows and columns will have the same indices. This results since the proposed representation selects the vectors centered at each sample in the signal's support (alternatively, it discards the vectors centered at the samples that are not in the signal's support). If the i^{th} sample is in the signal's support then the i^{th} row of T will be retained. Since the vector centered at each sample will be selected, the i^{th} vector will be selected which corresponds to the i^{th} column of T being retained. Thus, if the i^{th} row is retained then the i^{th} column is also retained. This is explicitly shown in the matrices below. First we begin with the initial $N \times N$ transform matrix. The signal's support determines which rows are relevant; the rows that are not relevant are discarded (denoted by $\times$). Our proposed representation then selects columns with the same indices as the retained rows (given once again by the indices of the signal support). Finally, the matrix is collapsed into an $M \times M$ submatrix. The remainder of this section focuses on proving the nonsingularity of the resulting $M \times M$ submatrix.

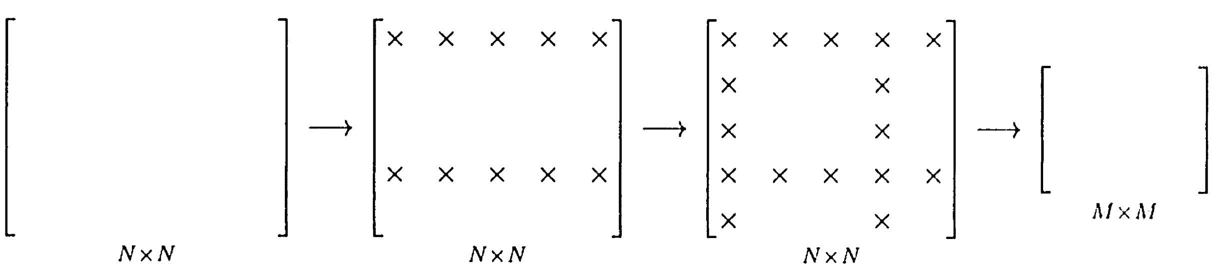

The resulting $M \times M$ submatrix of T corresponds to a *principal submatrix* of T. It is called a principal submatrix because the indices of the retained rows and columns are the same, and the diagonal elements of the submatrix correspond to principal diagonal elements of the original matrix T [68].

The question of whether the proposed representation leads to a basis is therefore equivalent to the problem of whether the resulting principal submatrix is invertible. This is an extremely useful formulation, because it relates the given problem to a problem concerning the nonsingularity of the principal submatrices of a matrix, and there already exists a sizable amount of knowledge in this area. For example, the question *does our representation provide a basis for all possible arbitrary 1-D supports* is equivalent to the question *are all principal submatrices of T nonsingular.* This property extends directly to 2-D and general M-D problems, where the corresponding T is equivalent to the M-times tensor product of the 1-D aligned T discussed above. Establishing that a given transform provides this property is the central theme of this section, and we therefore stress this property:

Proposition 4.3.1. *The proposed representation will provide a basis for all possible arbitrary supports if and only if all principal submatrices of the corresponding matrix T are nonsingular.*

4.3.1 Representing 1-D Signals with Arbitrary Supports

The proposed representation will lead to a basis for 1-D signals with any arbitrary support if and only if all principal submatrices of the matrix T are nonsingular. Therefore, an important question is which 2-band filterbanks correspond to a matrix T whose principal submatrices are all nonsingular. For point of reference, the question of what class of matrices have all principal submatrices nonsingular is an important and open question in the mathematics community. Similarly, the questions of which Toeplitz or block Toeplitz matrices have all principal submatrices nonsingular also appear to be open.

There are a number of simple matrices where all of their principal submatrices are invertible. For example, diagonal matrices (with all diagonal elements nonzero), triangular matrices, and diagonally dominant matrices have all principal submatrices nonsingular. This follows since the original (full size) matrix is nonsingular and every one of its principal submatrices also has the same structure and is therefore also nonsingular. For example, all triangular matrices are nonsingular, and every principal submatrix of a triangular matrix is another triangular matrix and is therefore also nonsingular. Symmetric positive and negative definite matrices also have all principal submatrices nonsingular [68]. All of the above matrix structures are sufficient conditions on a matrix so that all of its principal submatrices are nonsingular. In our case, T is a nonsymmetric block Toeplitz matrix. The above matrix properties are not directly applicable to T because they correspond to severe constraints on the filters of the 2-channel filterbank.

In theory, given a particular filterbank one can examine all the principal submatrices to determine their invertibility. However this is impractical for reasonable filter and signal lengths. A more useful approach would be to develop a condition on the filter taps such that all the principal submatrices are nonsingular. In the following, we propose a number of very simple conditions on the filter taps which are sufficient to guarantee that all principal submatrices are nonsingular.

These conditions are satisfied by a very broad class of filters, including most of the popular filters currently used.

Before proposing our condition, it is useful to briefly consider some necessary conditions on T in order for all of its principal submatrices to be nonsingular. To begin with, T must be nonsingular, otherwise it will not work for a full support signal. All the diagonal elements t_{ii} must be nonzero, or else it will not work for some 1-point signals. In addition, $t_{ii} \cdot t_{jj} \neq t_{ij} \cdot t_{ji} \; \forall i, j$, or else it will not work for some 2-point signals, i.e. there exist some ROS's with 2 samples such that the corresponding principal submatrix is singular. Of course, similar necessary conditions can be constructed for the principal submatrices corresponding to signals with supports of 3, 4, etc., samples, however the complexity of constructing and checking all of these conditions generally precludes their usefulness. It is instructive to note that in the case where T is invertible and the center tap of each vector is the largest amplitude tap of the vector, then the first three conditions are satisfied.

The appropriate T to be used for finite-size 1-D signals can be formed in either of two ways. Beginning with the infinite matrix corresponding to applying the 2-channel filterbank to an infinite-length signal, one can extract an appropriate $N \times N$ principal submatrix. Alternatively, a circulant matrix may be used where the matrix is defined over a support sufficiently larger than the largest possible signal ROS, such that the wrap-arounds do not overlap the signal ROS. The vectors should be aligned, as previously discussed, so that the centers of the vectors lie along the main (principal) diagonal of the matrix. The resulting block Toeplitz matrix T_{toep} and one possible block circulant matrix T_{circ} for signals with supports of up to 8 samples are given below. The matrices correspond to a 2-channel filterbank with 4-tap orthogonal filters, and only the nonzero entries are listed.

$$T_{toep} = \begin{bmatrix} b & -c & & & & & & \\ c & b & a & d & & & & \\ d & -a & b & -c & & & & \\ & & c & b & a & d & & \\ & & d & -a & b & -c & & \\ & & & & c & b & a & d \\ & & & & d & -a & b & -c \\ & & & & & & c & b \end{bmatrix} \tag{4.12}$$

$$T_{circ} = \begin{bmatrix} b & -c & & & & & & & & & d & -a \\ c & b & a & d & & & & & & & & \\ d & -a & b & -c & & & & & & & & \\ & & c & b & a & d & & & & & & \\ & & d & -a & b & -c & & & & & & \\ & & & & c & b & a & d & & & & \\ & & & & d & -a & b & -c & & & & \\ & & & & & & c & b & a & d & & \\ & & & & & & d & -a & b & -c & & \\ & & & & & & & & c & b & a & d \\ & & & & & & & & d & -a & b & -c \\ a & d & & & & & & & & & c & b \end{bmatrix} \tag{4.13}$$

Notice that T_{circ} is larger than T_{toep} and in fact contains T_{toep}. Therefore, if all principal submatrices of T_{circ} are nonsingular that implies that all principal submatrices of T_{toep} are nonsingular. This is useful because under certain cases the analysis of the block circulant T_{circ} leads to the analysis of a corresponding circulant matrix, whose eigenvalues and eigenvectors may be easily computed because of its circulant property. In those situations, various properties of the block Toeplitz matrix T_{toep} may be deduced from a much simpler circulant matrix. Notice that T_{circ} having all principal submatrices nonsingular is a sufficient (but not necessary) condition for all of the principal submatrices of T_{toep} being nonsingular. That is, T_{toep} can have all principal submatrices nonsingular while T_{circ} does not (the wrap-arounds can cause difficulties). Henceforth, T will refer to T_{toep} unless otherwise stated.

For an arbitrary matrix A, a sufficient condition for all principal submatrices to be nonsingular is that the quadratic form

$$x^T \cdot A \cdot x \neq 0, \qquad \forall x \neq 0 \quad . \tag{4.14}$$

This follows since if $x^T A x \neq 0$ for all nonzero x, then the same is true for all principal submatrices, and therefore all principal submatrices are nonsingular. Consider the simple example of the upper left $K \times K$ principal submatrix denoted A_k, produced by retaining the first K rows and columns. To examine the properties of A_k, consider all vectors $\hat{x}$ whose last $N - K$ components are zero, i.e. $\hat{x}^T = [x_k^T \quad 0^T]$ where x_k is a $K \times 1$ column vector and 0 is a $(N - K) \times 1$ column vector. In this case the quadratic form simplifies to

$$\hat{x}^T A \hat{x} = \begin{bmatrix} x_k^T & 0^T \end{bmatrix} \begin{bmatrix} A_k & * \\ * & * \end{bmatrix} \begin{bmatrix} x_k \\ 0 \end{bmatrix} = x_k^T A_k x_k \quad . \tag{4.15}$$

Since $x^T A x \neq 0$ for all nonzero x, then also $\hat{x}^T A \hat{x} \neq 0$ for all nonzero $\hat{x}$, and in particular $x_k^T A_k x_k \neq 0$ for all nonzero x_k. Hence, A_k is nonsingular. In a similar manner one can show that all of the principal submatrices are nonsingular. Therefore, if A satisfies Equation 4.14 then all of its principal submatrices are nonsingular.

The above discussion is familiar for the case of symmetric matrices. If A is a symmetric positive definite matrix the quadratic form is always greater than zero. Therefore all the principal submatrices of a symmetric positive definite matrix are also positive definite, and in particular they are nonsingular. Similarly, a symmetric negative definite matrix has all principal submatrices negative definite, and hence nonsingular. The difficulty is that for virtually all 2-channel filterbanks, the matrix T is a nonsymmetric block Toeplitz matrix.

A key observation is that the above reasoning also applies for *nonsymmetric matrices*. Specifically, the symmetric portion of a nonsymmetric matrix determines the properties of the quadratic form. Any matrix can be decomposed into its symmetric and skew symmetric components

$$T = T_{sym} + T_{skew} \tag{4.16}$$

where

$$T_{sym} = \frac{1}{2} \cdot (T + T^T) \tag{4.17}$$

$$T_{skew} = \frac{1}{2} \cdot (T - T^T) \quad . \tag{4.18}$$

The properties of the quadratic form only depend on the symmetric part of T since

$$x^T \cdot T \cdot x = x^T \cdot (T_{sym} + T_{skew}) \cdot x = x^T \cdot T_{sym} \cdot x \quad . \tag{4.19}$$

Therefore, if T_{sym} is positive or negative definite, then all the principal submatrices of T are nonsingular.

At this point, we provide a roadmap for the remainder of this section. The goal is to determine appropriate conditions on the filter taps which would guarantee that the representation provides a basis for any arbitrary 1-D support. This is equivalent to determining conditions on the filter taps that guarantee that all principal submatrices of the initial matrix are nonsingular.

Directly determining these conditions for general orthogonal or biorthogonal filters is very difficult. Instead, we solve a number of simpler problems which lead to the case of general filters. We begin with the case where the LP and HP filters are related by alternating flip (orthogonal and semiorthogonal filters) and where the center taps have the same sign. This case posseses a very convenient structure that leads to a simple solution. We then examine some of the possible linear pre- and post-processing operations that can be applied to a matrix while preserving the nonsingularity of its principal submatrices. This enables us to extend our conditions on the filter taps to the general case where the LP and HP filters are related by alternating flip (but the center taps do not necessarily have the same sign). Finally, we develop conditions for the general case of a biorthogonal filterbank.

In the important case of orthogonal and semiorthogonal filters, T_{sym} has a very convenient structure. Let the LP and HP filters be denoted by $g_0(n)$ and $g_1(n)$, respectively [7]. For (causal) orthogonal and semiorthogonal filters of length L, the LP and HP filters are related by *alternating flip*, which in the time domain corresponds to

$$g_1(n) = (-1)^n g_0(L-1-n) \tag{4.20}$$

and in the z-domain to

$$G_1(z) = (-z)^{(-L+1)} G_0(-z^{-1}) \quad . \tag{4.21}$$

Therefore, when the filters are related by alternating flip their center taps are equal, except possibly of differing signs (depending on the filter length L). Assume that the center taps are of the same sign, which may require multiplying the HP filter by -1. With the filters appropriately shifted (delayed) the center taps of the LP and HP filters are $g_0(0)$ and $g_1(1)$, where $g_0(0) = g_1(1)$, and the filters are related by

$$g_1(n) = (-1)^{(n+1)} g_0(-n+1) \quad . \tag{4.22}$$

The condition on the signs of the center taps imposes a particular structure on the polyphase components. In addition, this condition coupled with the alignment of the filters results in the principal diagonal of T being constant, as shown in Equation 4.12. The condition on the signs of the center taps will be relaxed shortly, however it leads to an important special case, and also

[7]Notation: Throughout this chapter we focus on creating a basis for synthesizing the signal. Therefore, to stress this point, we have chosen to express the filters in the initial representation using the notation of the synthesis filters from Section 2.1.1. If the goal is to create a synthesis basis, it is natural to begin with a set of good synthesis vectors, however we must remark that irrespective of how the initial representation was formed, the created basis may be used for *either analysis or synthesis*.

greatly simplifies the proof. In this case, while T is block Toeplitz, T_{sym} is symmetric and Toeplitz – the block structure has disappeared. Furthermore, T_{sym} is only a function of the center phase of the filters. This structure is probably most evident in the polyphase domain. Consider the polyphase matrix corresponding to the center and off-center phases of the filters, where center and off-center are with respect to the center tap of the filter. The polyphase matrix for T in Equation 4.12 is

$$T_{poly} = \begin{bmatrix} b & & & & -c & & & \\ d & b & & & -a & -c & & \\ & d & b & & & -a & -c & \\ & & d & b & & & -a & -c \\ c & a & & & b & d & & \\ & c & a & & & b & d & \\ & & c & a & & & b & d \\ & & & c & & & & b \end{bmatrix} \tag{4.23}$$

and $T_{poly-sym}$ is

$$T_{poly-sym} = \frac{1}{2} \cdot \begin{bmatrix} 2b & d & & & & & & \\ d & 2b & d & & & & & \\ & d & 2b & d & & & & \\ & & d & 2b & & & & \\ & & & & 2b & d & & \\ & & & & d & 2b & d & \\ & & & & & d & 2b & d \\ & & & & & & d & 2b \end{bmatrix} . \tag{4.24}$$

The general case where the LP and HP filters are related by alternating flip, and their center taps are of the same sign, T can be expressed in polyphase notation as

$$T_{poly} = \begin{bmatrix} G_c & -G_o^T \\ G_o & G_c^T \end{bmatrix} \tag{4.25}$$

where G_c and G_o are Toeplitz matrices corresponding to the center and off-center phases, respectively. $T_{poly-sym}$ is then given by

$$T_{poly-sym} = \frac{1}{2} \cdot \begin{bmatrix} G_c + G_c^T & \\ & G_c + G_c^T \end{bmatrix} \tag{4.26}$$

and it is evident that the symmetric component only depends on the center phase of the filters.

The dependence on the center phase is somewhat intuitive, since the centering of the vectors coupled with the relationship between the LP and HP vectors means that the center phase is always (irrespective of the ROS) selected. Since the center taps are of the same sign, the off-center phases are of opposite signs and always cancel each other out.

The dependence on the center phase is also very intriguing because it describes some of the freedom that exists in the filters, while still providing a basis over any arbitrary 1-D support. Specifically, the problem of if the representation will lead to a basis for any arbitrary 1-D support only depends on the center phase of the filters, and does not depend on the off-center phase. The freedom in the off-center phase may be exploited to design improved filters for specific applications while still guaranteeing a basis for any arbitrary 1-D supports.

Since T_{sym} and $T_{poly-sym}$ are related by permutations of their rows and columns (equivalent to pre and post multiplying by an orthogonal matrix and its transpose), they have the same eigenvalues. Therefore, either matrix can be examined for possible positive or negative definiteness. While it is sufficient to only examine T_{sym} in this case, we also examine $T_{poly-sym}$ since the insights that arise are directly applicable to problems of more general filters. A matrix is positive (negative) definite if all of its eigenvalues are positive (negative). Therefore, we desire conditions on the filter taps such that all of the eigenvalues are either positive or negative. T_{sym} and $T_{poly-sym}$ are both symmetric and Toeplitz, therefore they can each be described by a fundamental vector that appears along their rows and columns. If $g(n)$ and $g_c(n)$ correspond to the LP filter and its center polyphase component, respectively, then the fundamental vectors denoted by $s(n)$ and $s_{poly}(n)$ are given by

$$s(n) = \frac{1}{2}(g(n) + g(-n)) \tag{4.27}$$

$$s_{poly}(n) = \frac{1}{2}(g_c(n) + g_c(-n)) \quad . \tag{4.28}$$

Excluding truncation at the matrix boundaries, every column of T_{sym} (and every row since T_{sym} is symmetric) corresponds to a shifted version of the fundamental vector $s(n)$. The shift is such that for every column $s(0)$ lies on the principal diagonal. The same properties exist for $T_{poly-sym}$ and $s_{poly}(n)$. $s_{poly}(n)$ is equal to $s(n)$ downsampled by 2.

Theorem 4.3.2. *If the LP and HP filters in a 2-channel filterbank are related by alternating flip, and their center taps are of the same sign, then a sufficient condition for the proposed representation to lead to a basis for all arbitrary 1-D supports is*

$$|s(0)| > \sum_{n \neq 0} |s(n)| \tag{4.29}$$

which is equivalent to

$$|s_{poly}(0)| > \sum_{n \neq 0} |s_{poly}(n)| \quad . \tag{4.30}$$

Proof: The goal is to show that the matrices are positive or negative definite. The inequalities establish the diagonal dominance of T_{sym} and $T_{poly-sym}$. Assume the center filter tap is positive, then the matrices have (constant) positive main diagonal. If the matrices are diagonally dominant and the diagonals are positive, then by Gershgorin's Theorem [8] on the location of the eigenvalues, all the eigenvalues are in the right half plane ($Re(\lambda_i) > 0$). Since the matrices are symmetric all the eigenvalues are real. Therefore, all the eigenvalues are positive real numbers and the matrices are positive definite. Similarly, if the matrices satisfy the above inequalities and the center tap is negative, then the matrices are negative definite ■

A somewhat broader condition can be produced by beginning with T_{circ} instead of T_{toep}. The symmetric portion of T_{circ}, denoted by $T_{circ-sym}$, has similar properties to the above, and in addition, is circulant. Its eigenvalues can therefore be computed by the Discrete Fourier Transform (DFT) of its first column, once again focusing attention on $s(n)$ and $s_{poly}(n)$. Since $s(n)$ and $s_{poly}(n)$ are symmetric, their DFT's (and all the eigenvalues) are real. A sufficient condition for $T_{circ-sym}$ to be positive definite is for the Discrete Time Fourier Transform (DTFT) of $s(n)$ to be positive for

[8] Gershgorin's Theorem [68, 69]: The eigenvalues of the $N \times N$ matrix A with entries a_{ij} lie in the union of the N circles

$$|z - a_{ii}| = \sum_{j \neq i} |a_{ij}| \quad 1 \leq i \leq N \tag{4.31}$$

in the complex plane. If the matrix A is strictly diagonally dominant

$$|a_{ii}| > \sum_{j \neq i} |a_{ij}| \quad 1 \leq i \leq N \tag{4.32}$$

then all the circles are away from the origin, hence all the eigenvalues are nonzero and therefore A is nonsingular.

all frequencies. Similarly, $T_{circ-sym}$ will be negative definite if the DTFT of $s(n)$ is negative for all frequencies. This leads to the following corollary of the previous Theorem.

Corollary 4.3.3. *If the LP and HP filters in a 2-channel filterbank are related by alternating flip, and their center taps are of the same sign, then sufficient conditions for the proposed representation to lead to a basis for all arbitrary 1-D supports are*

$$S(e^{j\omega}) > 0, \quad \omega \in [0, \pi] \tag{4.33}$$

or

$$S(e^{j\omega}) < 0, \quad \omega \in [0, \pi] \tag{4.34}$$

where

$$S(e^{j\omega}) = \sum_{n=-L/2}^{L/2} s[n]e^{-j\omega n} = s(0) + \sum_{n=1}^{L/2} 2s[n]\cos(\omega n). \tag{4.35}$$

Proof: The conditions follow straightforwardly from the fact that the eigenvalues of a circulant matrix are given by the DFT of its first column. Since $s(n)$ is real and symmetric its DTFT $S(e^{j\omega})$ is also real and symmetric. Therefore $S(e^{j\omega})$ can be written in terms of $\cos(\omega n)$, and it is only necessary to examine the frequency range $\omega \in [0, \pi]$ ■

In the general case of orthogonal and semiorthogonal filters where the LP and HP filters are related by alternating flip, but the center taps are not necessarily of the same sign, much of the structure exploited above no longer exists. Specifically, $T_{toep-sym}$ is not necessarily Toeplitz and $T_{circ-sym}$ is not necessarily circulant. However, the following simple lemma is very helpful.

Lemma 4.3.4. *Let A and B be $N \times N$ matrices related by*

$$A = D_1 \cdot B \cdot D_2 \tag{4.36}$$

where D_1 and D_2 are any $N \times N$ nonsingular diagonal matrices. Then A has all principal submatrices nonsingular if and only if B has all principal submatrices nonsingular.

Proof: Let A_{sub}, B_{sub}, D_{1-sub}, and D_{2-sub} be corresponding principal submatrices of A, B, D_1 and D_2, respectively, i.e. all of them are formed from their original matrices by retaining the exact same set of rows and columns. The principal submatrices are related by

$$A_{sub} = D_{1-sub} \cdot B_{sub} \cdot D_{2-sub}. \tag{4.37}$$

D_{1-sub} and D_{2-sub} are always nonsingular since they are principal submatrices of (nonsingular) diagonal matrices. If B has all principal submatrices nonsingular then B_{sub} is always nonsingular, and by Equation 4.37, A_{sub} is nonsingular. Therefore, if B has all principal submatrices nonsingular then A also has all principal submatrices nonsingular. By pre-multiplying Equation 4.36 by D_1^{-1} and post-multiplying by D_2^{-1}, the same argument as above shows that if A has all principal submatrices nonsingular then so does B ■

In essence, this lemma says that if a submatrix is nonsingular, scaling all its rows and/or all its columns by nonzero scaling factors will still preserve its nonsingularity. That is, if a set of vectors are linearly independent, scaling each vector by a non-zero constant will not alter the linear independence. The case where the center taps of the LP and HP filters are of differing signs simply corresponds to the previous T with either the even or odd numbered columns (corresponding to either LP or HP filters) multiplied by -1. Then by the above lemma we get the following important theorem.

Theorem 4.3.5. *If the LP and HP filters in a 2-channel filterbank are related by alternating flip, then a sufficient condition for the proposed representation to lead to a basis for all arbitrary 1-D supports is*

$$|s_{poly}(0)| > \sum_{n \neq 0} |s_{poly}(n)| \tag{4.38}$$

where

$$s_{poly}(n) = \frac{1}{2}(g_c(n) + g_c(-n)) \quad . \tag{4.39}$$

Proof: If the center taps of the LP and HP filters are of the same sign, Theorem 4.3.2 gives sufficient conditions such that all the principal submatrices of the corresponding T, call it $T_{same-sign}$, are nonsingular. The general case of no constraints on the signs of the center taps corresponds to a $T_{no-constraints}$ which is equal to $T_{same-sign}$ with some of the columns multiplied by -1. Therefore, by Lemma 4.3.4, if $T_{same-sign}$ has all principal submatrices nonsingular then so does $T_{no-constraints}$. The sufficient conditions for the general case are equivalent to the sufficient conditions of the special case given in Theorem 4.3.2. Specifically, the conditions only depend on the center phase ■

This theorem provides a very simple sufficient condition on the filter taps such that the proposed representation is guaranteed to provide a basis for all arbitrary 1-D supports. This condition is satisfied by a surprisingly large number of orthogonal and semiorthogonal filters. Some of the filters which satisfy this condition include Daubechies 2-16 tap filters (closest-to-linear phase for 8-16 taps), Smith and Barnwell's 8-tap, and Simoncelli's 9-tap.

It is beneficial to briefly reflect on the approach of this proof. The difficult problem of the general orthogonal and semiorthogonal filterbank was solved by beginning with a simple special case which could be easily analyzed, and then bootstrapping off of it via Lemma 4.3.4 to provide a simple proof for the general problem. Determining if the proposed representation leads to a basis for a given set of filters is often rather difficult, since T_{sym} is typically not positive definite. However, the proof points to a general approach that may be useful for converting a number of difficult problems into much simpler problems. Specifically, assume that we begin with some T for which T_{sym} is not positive definite. If there exists a scaling of the rows and/or columns of T to produce a $\mathcal{T}$ such that $\mathcal{T}_{sym}$ is positive definite, then all the principal submatrices of the original T are nonsingular. In addition, this helps to identify what are the fundamental issues that determine if a pair of filters will or will not work. For example, this theorem shows that irrespective of if the center taps have the same signs or not (if the off-center phases cancel or not), the nonsingularity of the principal submatrices *only depends on the center phase*. This knowledge may be useful in a number of areas, such as the design of filters with improved characteristics.

In the general case where the LP and HP filters are not related by alternating flip, there is even less structure in T_{sym} to exploit. This case includes the very important class of biorthogonal linear phase filters. However, the ideas developed above based on positive definiteness and diagonal dominance are still applicable, and they apply for a surprisingly large number of filters.

Let the LP and HP filters for a general two-channel filterbank be denoted by $g_0(n)$ and $g_1(n)$, and assume they are shifted (or indexed) so that their center taps are at $g_0(0)$ and $g_1(1)$, respectively. Their center and off-center phases, denoted by $g_{0c}(n)$, $g_{0o}(n)$, $g_{1c}(n)$, and $g_{1o}(n)$, respectively, are given by

$$g_{0c}(n) = g_0(2n) \tag{4.40}$$

$$g_{0o}(n) = g_0(2n+1) \tag{4.41}$$

$$g_{1c}(n) = g_1(2n+1) \tag{4.42}$$

$$g_{1o}(n) = g_1(2n+2) \tag{4.43}$$

and $g_{0c}(0)$ and $g_{1c}(0)$ contain the center taps of the LP and HP filter. To aid in clarification, the LP and HP filters are related to their polyphase components in the z-transform domain by

$$G_0(z) = G_{0c}(z^2) + z^{-1}G_{0o}(z^2) \tag{4.44}$$

$$G_1(z) = z^2 G_{1o}(z^2) + z^{-1}G_{1c}(z^2) \quad . \tag{4.45}$$

The matrix T formed by this 2-channel filterbank is 2×2 block Toeplitz. That is, going along the principal diagonal of T, the columns are periodic with period of two (even columns correspond to the LP filters and the odd columns to the HP filters). In the special case discussed earlier, when the LP and HP filters were related by alternating flip and their center taps were of the same signs, cancellations occurred and T_{sym} was symmetric and Toeplitz. In the general case no cancellations occur and T_{sym} is symmetric but 2×2 block Toeplitz. Therefore, instead of specifying T_{sym} via a single fundamental vector, two vectors are needed. All the even numbered columns of T_{sym} are equal (to within a shift and truncation at the matrix boundary). Similarly, all the odd columns of T_{sym} are equal. T_{sym} can therefore be characterized by two vectors, $s_{even}(n)$ and $s_{odd}(n)$, centered on the principal diagonal of the even and odd numbered columns. These vectors are given by

$$s_{even}(n) = \begin{cases} a_c(\frac{n}{2}) & \text{if } n \text{ is even,} \\ a_o(\frac{n-1}{2}) & \text{if } n \text{ is odd} \end{cases} \tag{4.46}$$

where

$$a_c(n) = \frac{1}{2}(g_{0c}(n) + g_{0c}(-n)) \tag{4.47}$$

$$a_o(n) = \frac{1}{2}(g_{0o}(n) + g_{1o}(-n)) \tag{4.48}$$

and

$$s_{odd}(n) = \begin{cases} b_c(\frac{n}{2}) & \text{if } n \text{ is even,} \\ b_o(\frac{n-1}{2}) & \text{if } n \text{ is odd} \end{cases} \tag{4.49}$$

where

$$b_c(n) = \frac{1}{2}(g_{1c}(n) + g_{1c}(-n)) \tag{4.50}$$

$$b_o(n) = \frac{1}{2}(g_{1o}(n) + g_{0o}(-n)) \quad . \tag{4.51}$$

Note that $b_o(n) = a_o(-n)$. As an aside, if A_c, A_o, B_c, and B_o are Toeplitz matrices corresponding to a_c, a_o, b_c, and b_o respectively, then $T_{poly-sym}$ is given by

$$T_{poly-sym} = \begin{bmatrix} A_c & B_o \\ A_o & B_c \end{bmatrix} = \begin{bmatrix} A_c & A_o^T \\ A_o & B_c \end{bmatrix} \quad . \tag{4.52}$$

The above formulation for a general 2-channel filterbank leads to the following theorem.

Theorem 4.3.6. *For a 2-channel filterbank, a sufficient condition for the proposed representation to lead to a basis for all arbitrary 1-D supports is for the following two conditions to hold*

$$|s_{even}(0)| > \sum_{n \neq 0} |s_{even}(n)| \tag{4.53}$$

and

$$|s_{odd}(0)| > \sum_{n \neq 0} |s_{odd}(n)| \tag{4.54}$$

where the variables are defined as above.

Proof: If the inequalities are satisfied then T_{sym} is diagonally dominant. Assume that the inequalities are satisfied. If the center taps of both the LP and HP filter are positive, then all the eigenvalues of T_{sym} are positive and T_{sym} is positive definite. If the center tap of either or both the LP and HP filter are negative, then the columns of T corresponding to the negative center taps can each be scaled by -1 (equivalent to post-multiplying by a diagonal matrix with 1's and -1's on the principal diagonal) resulting in a new T_{new} whose symmetric part is positive definite. Therefore, T_{new} has all principal submatrices nonsingular, and by Lemma 4.3.4, T also has all principal submatrices nonsingular ■

This simple condition is a generalization of Theorem 4.3.5 and is applicable to any 2-channel filterbank. Specifically, the filterbank may be orthogonal, semiorthogonal, biorthogonal, or even a non-perfect reconstruction filterbank [9]. This condition is satisfied by a surprisingly large number of filter sets – in fact, every conventional LP/HP filter pair tested satisfied this condition. Further details are given in Section 4.3.3.

That this is a generalization of our previous results follows straightforwardly. For example, in the special case of orthogonal and semiorthogonal filters where the center taps have the same sign, this theorem simplifies to Theorem 4.3.2. Specifically, if the LP and HP filters are related by alternating flip and their center taps have the same sign, then

$$g_1(n) = (-1)^{(n+1)} g_0(-n+1) \tag{4.55}$$

and therefore

$$g_{1c}(n) = g_{0c}(-n) \tag{4.56}$$

$$g_{1o}(n) = -g_{0o}(-n) \tag{4.57}$$

which leads to

$$a_c(n) = b_c(n) \tag{4.58}$$

$$a_o(n) = b_o(n) = 0 \tag{4.59}$$

and

$$s_{even}(n) = s_{odd}(n) \quad . \tag{4.60}$$

Therefore, under these conditions, Theorem 4.3.6 is equivalent to Theorem 4.3.2 as expected. In addition, $A_c = B_c = \frac{1}{2}(G_c + G_c^T)$ and $A_o = B_o = 0$, leading back to Equation 4.26.

Remarks The discussion has focused on the selection of a set of vectors for synthesizing the signal, that is providing a synthesis basis. However, the results also apply for the selection of a set of vectors for analyzing the signal, that is for choosing a set of vectors such that their inner products with the signal are linearly independent and enable us to reconstruct the signal. Of course, this follows since once a basis is determined, it can be applied for either analysis or synthesis. This can also be seen as follows. For synthesis we focus on the matrix T, for analysis we focus on the matrix T^T. In both cases all the principal submatrices must be nonsingular. The principal submatrix for

[9]Typically, a filterbank is referred to as perfect reconstruction (PR) or non-PR based on if it can achieve PR given some constraints on the analysis and synthesis structure, e.g. FIR filters. However, a set of filters may be "non-PR" but may still provide a useful basis for our approach.

the analysis of a given signal is the transpose of the principal submatrix examined for synthesis of the signal. Therefore, if one is nonsingular then so is the other. This also follows since the proofs examine the symmetric portion of the synthesis matrix T, given by $\frac{1}{2}(T + T^T)$, which is equivalent to the symmetric portion of the analysis matrix T^T.

This approach may also be applied to filterbanks where the filters are not conveniently centered, e.g. for minimum phase filters where the largest amplitude taps are near the beginning of the filter (such as in Daubechies original construction of orthogonal filters). Most of the popular 2-channel filters have their energy concentrated over a subset of their taps, and typically have a single tap of largest magnitude. The largest magnitude tap can therefore be taken to be the "center tap" of the filter. This is a natural approach, and is equivalent to choosing the exact center tap for typical linear phase, odd length filters. The LP and HP filters can be advanced or delayed versus each other so that the corresponding "center taps" will be adjacent to each other, leading once again to a simple selection based on the ROS. Choosing the largest magnitude taps as the center taps makes the principal diagonal of T larger than the off-diagonal elements [10]. Thus T_{sym} tends to be diagonally dominant, and this also improves the conditioning of the resulting representation.

4.3.2 Representing 2-D Signals with Arbitrary Supports

General 2-D/M-D problems are considerably more complex than their 1-D counterparts, and the same situation exists here. To begin, we extend the matrix formulation to 2-D. The 2-D signal to be represented can have any support, and the initial transform is defined over a superset space with convenient support. The separable transform of a 2-D signal X with a square $N \times N$ sample ROS is given by

$$Y = T_v^{\mathrm{T}} \cdot X \cdot T_h \tag{4.61}$$

where T_h and T_v are the (possibly different) transforms applied along the horizontal and vertical directions, and Y is the transformed signal. The signal can be reconstructed by inverse transforming the coefficients

$$X = T_v \cdot Y \cdot T_h^{\mathrm{T}} \tag{4.62}$$

This may be expressed as a single matrix-vector product via the use of Kronecker products.

[10]Specifically, each diagonal element is larger than all the other elements in its column, and depending on the relation between LP and HP filters, it may also be larger than all the other elements in its row.

The Kronecker product (also called direct product or tensor product) of two $N \times N$ matrices B and C is defined by

$$A = B \otimes C = \begin{bmatrix} b_{11}C & b_{12}C & \dots & b_{1n}C \\ b_{21}C & b_{22}C & \dots & b_{2n}C \\ \vdots & \vdots & \ddots & \vdots \\ b_{n1}C & b_{n2}C & \dots & b_{nn}C \end{bmatrix} \quad . \tag{4.63}$$

Thus, A is an $N^2 \times N^2$ matrix composed of $N \times N$ blocks, where the (i, j) block is $b_{ij}C$.

The matrix equation 4.62 can then be expressed as a single matrix-vector product via the use of the Kronecker product and by columnwise scanning of the $N \times N$ matrices X and Y into the $N^2 \times 1$ vectors x and y:

$$X = T_v \cdot Y \cdot T_h^{\mathrm{T}} \Longleftrightarrow x = (T_h \otimes T_v)y \tag{4.64}$$

In this manner the 2-D problem is converted into a linear system. In typical applications, the same transform is applied along both horizontal and vertical directions, i.e. $T_h = T_v$. In the following we assume that $T = T_h = T_v$. To simplify our notation, the Kronecker product matrix for the linear system will be denoted by

$$\Theta = T \otimes T \quad . \tag{4.65}$$

For this proposed representation, the set of the selected vectors is equal to the support of the signal. For instance, the set of coefficients of Y that are selected for representing a signal X is equivalent to the ROS of the signal X. This simple relationship provides quite a number of benefits. For instance, the question of if the representation will provide a basis for all possible arbitrary 2-D supports is equivalent to the question of if all principal submatrices of Θ are nonsingular. This is a simple extension from 1-D, and follows once again since we are selecting the vector centered at each sample in the ROS.

Determining the class of filters for which the proposed representation leads to a basis for any arbitrary support is much more intricate in 2-D than in 1-D. For example, a set of filters must work in 1-D if they are to work in 2-D. However, if a filter set works in 1-D it does not necessarily work in 2-D or M-D. The Haar filters provide a very simple example of this. The Haar matrix for transforming a four sample 1-D signal is given by

$$T = \begin{bmatrix} 1 & -1 & & \\ 1 & 1 & & \\ & & 1 & -1 \\ & & 1 & 1 \end{bmatrix} \tag{4.66}$$

and Θ for a 4×4 sample 2-D signal is given by the 16×16 matrix

$$\begin{bmatrix}
1 & -1 & & & -1 & 1 & & & & & & & & & & \\
1 & 1 & & & -1 & -1 & & & & & & & & & & \\
& & 1 & -1 & & & -1 & 1 & & & & & & & & \\
& & 1 & 1 & & & -1 & -1 & & & & & & & & \\
1 & -1 & & & 1 & -1 & & & & & & & & & & \\
1 & 1 & & & 1 & 1 & & & & & & & & & & \\
& & 1 & -1 & & & 1 & -1 & & & & & & & & \\
& & 1 & 1 & & & 1 & 1 & & & & & & & & \\
& & & & & & & & 1 & -1 & & & -1 & 1 & & \\
& & & & & & & & 1 & 1 & & & -1 & -1 & & \\
& & & & & & & & & & 1 & -1 & & & -1 & 1 \\
& & & & & & & & & & 1 & 1 & & & -1 & -1 \\
& & & & & & & & 1 & -1 & & & 1 & -1 & & \\
& & & & & & & & 1 & 1 & & & 1 & 1 & & \\
& & & & & & & & & & 1 & -1 & & & 1 & -1 \\
& & & & & & & & & & 1 & 1 & & & 1 & 1
\end{bmatrix} . \tag{4.67}$$

In this case there are a number of principal submatrices that are singular. For example, the 2×2 principal submatrix corresponding to retaining the first and sixth rows and columns is singular. Therefore, the proposed representation will not provide a basis for the ROS corresponding to this principal submatrix. This example is also useful for illustrating how the matrix Θ grows with the square of the number of samples in the support of the initial transform. Therefore, as the size of the signal support increases, the matrix Θ grows very quickly, and it is beneficial to avoid ever creating it. This problem becomes even more prohibitive for M-D transforms.

We would like to design a simple test or condition on the filter taps that determines whether or not all the principal submatrices of Θ are nonsingular. This would be analogous to our simple and successful test in 1-D. Part of the difficulty in 2-D arises because even if the 1-D transform T

has T_{sym} positive definite, the 2-D transform Θ does not necessarily have Θ_{sym} positive definite. This follows since if $T = T_{sym} + T_{skew}$ then

$$\begin{aligned} \Theta &= T \otimes T && (4.68) \\ &= (T_{sym} + T_{skew}) \otimes (T_{sym} + T_{skew}) && (4.69) \\ &= T_{sym} \otimes T_{sym} + T_{sym} \otimes T_{skew} + T_{skew} \otimes T_{sym} + T_{skew} \otimes T_{skew} && (4.70) \end{aligned}$$

and

$$\begin{aligned} \Theta_{sym} &= \frac{1}{2}(\Theta + \Theta^{\mathsf{T}}) && (4.71) \\ &= \frac{1}{2}(T \otimes T + (T \otimes T)^{\mathsf{T}}) && (4.72) \\ &= T_{sym} \otimes T_{sym} + T_{skew} \otimes T_{skew} && (4.73) \end{aligned}$$

thus Θ_{sym} is a function of both T_{sym} and T_{skew}. If T_{sym} is positive definite, then $T_{sym} \otimes T_{sym}$ is also positive definite. However, the second term may result in Θ_{sym} not being positive definite.

These issues make it difficult to prove that the proposed representation will work for a given filter set in 2-D or general M-D. The 1-D analysis and conditions can be extended to 2-D and M-D, but a straightforward extension leads to relatively narrow sufficient conditions, i.e. conditions that cover only a small number of the filter sets that empirically appear to work. Alternatively, one may individually examine each filter to determine if all the principal submatrices of Θ are nonsingular. For instance, as is discussed in Section 4.3.3, Daubechies 4, 6, and 8-tap orthogonal filters, Daubechies 9/7-tap biorthogonal filters, and the Binary filters lead to a positive definite Θ_{sym} in 2-D. Therefore, the proposed representation provides a basis for any arbitrary 2-D support using these filters. However, as is discussed in Section 4.3.3, both empirical evidence and theoretical considerations suggest that a much larger group of filters provide this property. Therefore, an important avenue of future research involves determining what class of filters provide a basis for any 2-D support with the proposed representation, as well as determining what simple conditions on the filter taps guarantee that a given filter is within this class.

4.3.3 Choice of Filters

For specific choices of LP/HP filter pairs, the proposed representation is guaranteed to provide a basis for any arbitrary ROS. In the following, we specifically consider the question of which filters can be used with the proposed representation such that the representation is guaranteed to provide a basis for any arbitrary ROS.

A filter pair will be referred to as admissible (for a particular dimension, e.g. 1-D or 2-D) if the proposed representation provides a basis for any arbitrary ROS (in that dimension). Specifically, we apply the ideas and conditions that we have developed for this representation to a number of the popular filter pairs to determine if they are admissible. The filters tested included orthogonal filters (Daubechies [12], Smith-Barnwell [70]), pseudo-semiorthogonal filters (Simoncelli [71]) and linear-phase biorthogonal filters (Daubechies 9/7-tap [12], Strang and Sweldens "Binary" 9/7-tap [2], and Le Gall 3/5-tap [72]).

The conditions that we have developed make the 1-D analysis quite simple. Specifically, every filter tested is an admissible filter for the representation in 1-D – the representation provides a basis for any arbitrary 1-D support when any of the tested filters are used. The results are tabulated in Table 4.3.3. All the filters can be confirmed to be admissible in 1-D by simply examining the filter taps and using Theorems 4.3.5 and 4.3.6 and Lemma 4.3.4.

The analysis in 2-D and general M-D is somewhat more complicated. A set of simple and useful conditions on the filter taps that ensure that the representation provides a basis for every arbitrary 2-D or M-D ROS is currently unknown. The conditions currently known are either extremely complex or highly restrictive such that no meaningful filters (for signal/image processing) satisfy the conditions. As will be discussed shortly, theoretical considerations and empirical evidence suggests that a very large class of filters are admissible in 2-D and general M-D. An interesting avenue of future research involves finding a simple description of this class of filters, as well as finding a simple condition on the filter taps to determine if a given filter is within this class.

Determining whether a given filter pair is admissible is a one-time process. Therefore, one possible approach for 2-D/M-D is to examine individually each filter under consideration in order to determine if it is admissible. That is, for each filter under consideration examine Θ and Θ_{sym} to determine if all the principle submatrices of Θ are nonsingular. In the simplest case, this may be achieved by showing that Θ_{sym} is positive definite. This approach was used to show that Daubechies 4-tap, 6-tap, and 8-tap orthogonal filters, Daubechies 9/7 tap biorthogonal filters, and the Binary filters work for 2-D. For the case where Θ_{sym} is not positive definite, one may use

Filters	1-D	2-D	3-D	4-D
Daubechies 2-tap (Haar)	Y	N	N	N
Daubechies 4-tap	Y	Y	e	e
Daubechies 6-tap	Y	Y	e	e
Daubechies 8-tap	Y	Y	e	e
Daubechies 10-tap	Y	e	e	e
Daubechies 12-tap	Y	e	e	e
Daubechies 14-tap	Y	e	e	e
Daubechies 16-tap	Y	e	e	e
Daubechies BO 9/7-tap	Y	Y	e	e
Smith-Barnwell 8-tap	Y	e	e	e
Binary (Strang's 9/7-tap)	Y	Y	e	e
Simoncelli 9-tap	Y	e	e	e
Le Gall 3/5-tap	Y	e	e	e

Table 4.1: The proposed representation leads to a basis for any arbitrary region of support for specific choices of filters. This table summarizes the properties of the proposed representation for a variety of different filters and for 1-D, 2-D, 3-D, and 4-D arbitrary regions of support. "Y": yes, provides a basis for any arbitrary ROS. "N": no, there exist supports for which it does not provide a basis. "e": empirical evidence suggests that it provides a basis for a wide range of possible supports (provided a basis for every support tested).

Lemma 4.3.4 as well as explore other possible methods to show that all principal submatrices are nonsingular.

A large number of tests were performed in order to check if the various filters work for 2-D, 3-D, and 4-D arbitrary ROS's. In particular, for each filter a large number of randomly generated arbitrary 2-D, 3-D, and 4-D supports were created, and in each case the representation was checked to see if it provided a basis [11]. For every filter pair except Haar, the representation provided a basis for every support tested. The Haar filters were examined in Section 4.3.2 and it was shown that there exist some 2-D supports such that the proposed representation with Haar filters does not provide a basis. Since the Haar filters are not admissible for 2-D, they are also not admissible for any higher dimension. However, the key point is that *for every filter except Haar, the representation provided a basis for every single tested support.* This is very encouraging because it shows that at the very least the proposed representation can be used with a broad range of possible filters to provide a basis for a large variety of possible arbitrary supports. For example, any of the filters tested (even if not shown to be admissible) will probably provide a basis for any given 2-D arbitrary support.

In a separate experiment, the condition number was computed for each filter and for a large number of randomly generated 2-D supports (details given in Section 4.3.4). Typically the worst condition number measured for a given filter over all the tested supports was reasonably low, roughly about 10 to 20. This suggests that not only does the proposed representation provide a basis in a wide range of situations, but also that the bases created are reasonably well-conditioned. This is also encouraging because if any of the resulting bases were ill-conditioned, the representation would be close to singular, and this would suggest that there may exist a support that could make it singular.

When this proposed representation was introduced, we related the issue of whether the representation provides a basis for all possible arbitrary supports to the issue of whether all the principal submatrices of an initial matrix are nonsingular. This relationship provides considerable insight into the success of the representation. For example, a random matrix (composed of entires with random amplitudes) is very likely to be nonsingular, i.e. a singular matrix is a very special type of matrix. In addition, a principal submatrix of a random matrix is probably nonsingular, as are all of its possible principal submatrices. A block Toeplitz matrix created by two random vectors is also very likely to be nonsingular, and probably have nonsingular principal submatrices. It is a highly structured pair of vectors that lead to singularity. While the filters designed for various applications are highly structured, their structure typically prevents the singularity. For instance, most filters have a single largest tap. Centering the filters at that tap produces matrices with a

[11] The tests in 2-D were quite extensive, but the tests in 3-D and especially in 4-D were much more limited because of memory and computational constraints.

strong diagonal component, thus increasing the likelihood of nonsingularity (and also typically improving the conditioning). Therefore it makes sense that the proposed approach should provide a basis for all arbitrary supports for a large group of filters.

To summarize, we have shown that in 1-D using any of the tested filters with the proposed representation will provide a basis for any arbitrary 1-D support. In 2-D and general M-D, we have shown that we have complete freedom in the ROS for a small group of filters. However, the above empirical evidence and theoretical considerations strongly suggest that the representation leads to a basis for any arbitrary support for a much larger group of filters then we have explicitly shown. An interesting open question is to identify what fundamental property, within the context of the proposed representation, enables some filters to lead to a basis for any arbitrary support.

4.3.4 Properties of Representation # 2

Wavelet representations have a number of basic properties that have been crucial to their success throughout a wide range of applications. Section 4.1.1 briefly reviewed some of the important properties of a conventional wavelet transform for infinite-length signals. In this section, we examine the degree to which the proposed representation preserves these important properties. We also consider some additional properties that are important for representing signals with arbitrary supports.

Flexibility in the Possible ROS's One of the most important properties of a representation for signals with arbitrary supports is flexibility in the range of possible supports that may be represented. As an example of a typical 2-D signal to be analyzed, consider the Lenna signal shown in Figure 4.5. The Lenna signal has a relatively simple arbitrarily shaped ROS that contains 7632 pixels. A two-level wavelet decomposition of the Lenna signal using Daubechies 8-tap filters is also shown in Figure 4.5. The representation provides many of the features that we are accustomed to from wavelet transforms of signals with convenient, square ROS's.

A very powerful feature of this proposed representation is that it provides great flexibility in the possible supports that may be analyzed. For example, the Star signal shown in Figure 4.6 has a very complex support – the support has discontiguous regions, regions with holes, and isolated samples. There are a total of 3143 pixels in the support. The two-level wavelet decomposition of the signal using Daubechies 8-tap filters with the proposed approach is also shown in Figure 4.6.

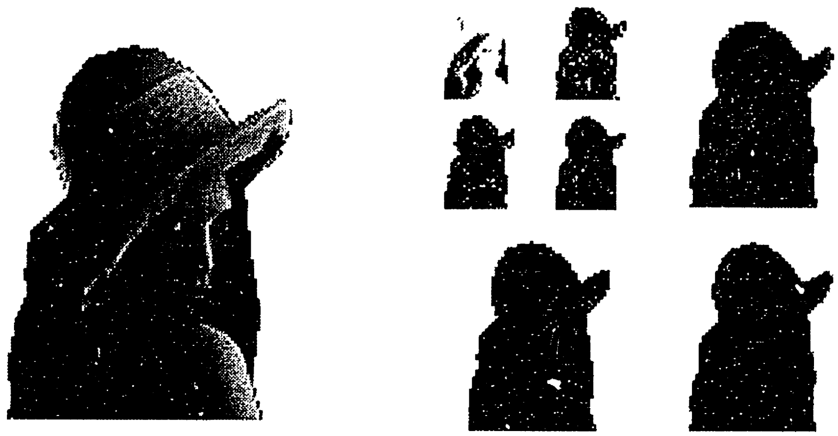

Figure 4.5: The 2-D Lenna signal which has an arbitrarily shaped region of support, and its two-level wavelet representation using representation #2. To aid in interpretation, the lowpass subband has been scaled down by a factor of 2 and the other subbands have been offset to gray.

Critical Sampling and Perfect Reconstruction These representations are critically sampled and provide the perfect reconstruction property. That is, the number of coefficients in the representation is equal to the number of samples in the original signal, and the signal can be perfectly reconstructed from the coefficients.

As a simple comparison, a conventional two-level wavelet transform over a circumscribing square would have 10200 and 7803 nonzero coefficients, respectively, for the two signals. This results in a 34% and 148% increase in the number of coefficients in the representations as compared to pixels in the original signals. This is summarized in Table 4.3.4. As will be discussed in Chapter 5, in many compression applications it may be advantageous to have a critically sampled approach. The increase in the number of coefficients in an overcomplete representation may potentially lead to an inefficiency in compression due to the extra overhead that may be needed to specify which coefficients are coded.

Signal	Proposed Representation (Critically Sampled)	Transform over Square (Overcomplete)
Lenna	7632	10200 (+ 34 %)
Star	3143	7803 (+ 148 %)

Multiresolution The representation is multiresolution in that the signal is decomposed at a number of scales and also in that the representation is computed by solving for the decomposition at a succession of coarser scales. An important point is that the representation enables a natural

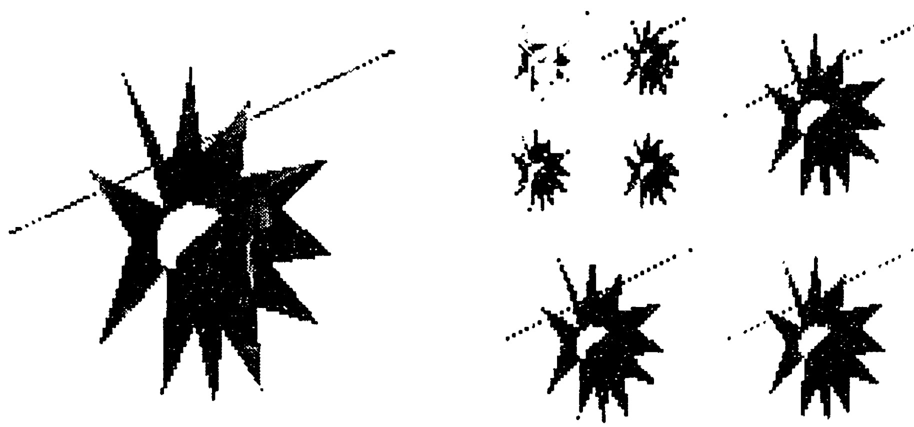

Figure 4.6: The 2-D Star signal with a very complex ROS and its two-level wavelet representation using representation #2. To aid in interpretation, the lowpass subband has been scaled down by a factor of 2 and the other subbands have been offset to gray.

reconstruction of the A-ROS signal at different scales. This property is illustrated in Figure 4.7, where a coarse approximation of the signal is reconstructed by using only the lowest frequency subband, and finer approximations are reconstructed by also using higher frequency subbands. The multiresolution property is an important feature of wavelet representations of infinite-length signals, and it is equally significant in wavelet representations of A-ROS signals.

An additional benefit of this representation is that the ROS can also be reconstructed at a number of different scales. Specifically, the shape of the object in Figure 4.7 is expressed in a natural manner at a number of different scales. Therefore, this representation enables one to express both the amplitude and the shape (support) of the signal at a number of different scales.

DC-leakage into Highpass Subbands This representation does not preserve the polynomial accuracy property. In many signal and image processing applications, preserving the polynomial accuracy property is not critical. However, concentrating the DC-energy into the lowpass subband is quite important. For example, in image compression it is desirable for the DC-energy to be concentrated in the lowpass subband and to have minimal leakage of DC-energy into the highpass subbands.

If the selected LP vectors can exactly represent a constant signal, then all the DC-energy is concentrated in the LP subband. This follows from the biorthogonality of the representation: the HP analysis vectors are orthogonal to the LP synthesis vectors, so if the span of the LP synthesis vectors includes the constant signal, then the HP analysis vectors are orthogonal to the constant

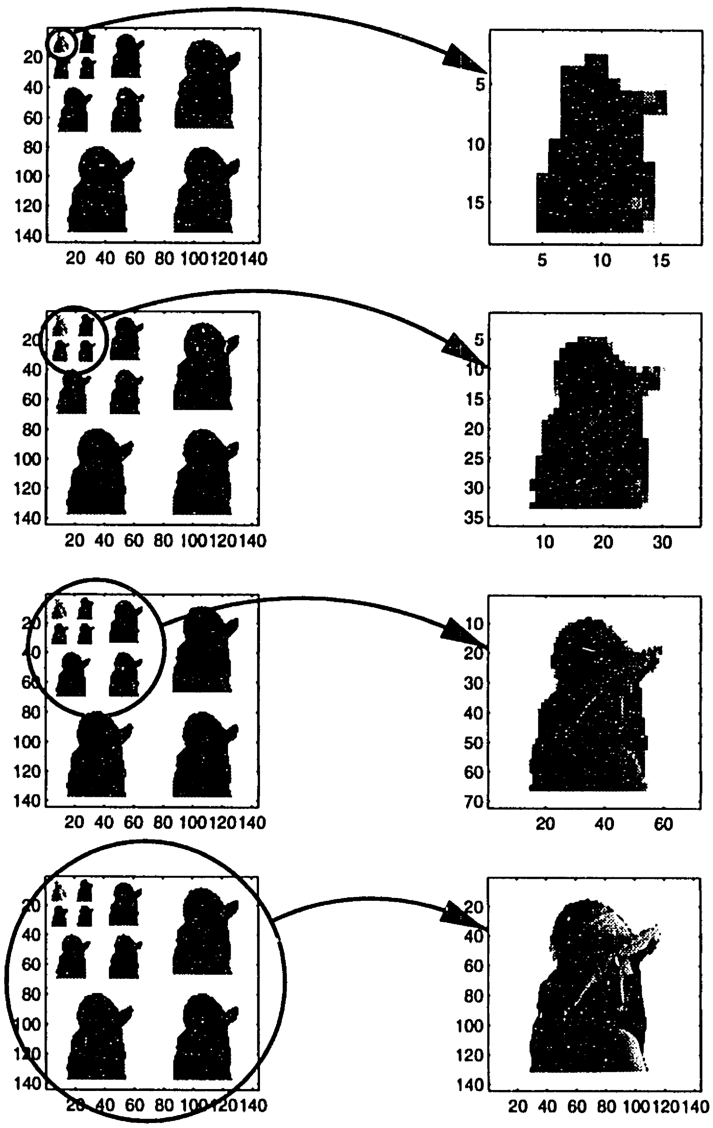

Figure 4.7: Reconstruction of the lenna signal with an arbitrarily shaped support at progressively finer scales by using more and more subbands. Note that both the amplitude of the signal and the signal's support (shape) can be reconstructed at a number of different scales.

signal and therefore have zero DC-gain. If the LP vectors can not exactly represent the constant signal then some HP vectors are required to reconstruct the constant signal; thus there will be DC-leakage into the HP subband.

In general, the selected LP vectors can not exactly represent a constant signal, and therefore the representation does not completely concentrate the DC-energy into the LP subband. The amount of DC-leakage into highpass subbands depends on the specific filters used. The DC-leakage for six filter pairs is shown in Figure 4.8. The original signal was a 48-point unit-amplitude signal which was aligned with respect to the initial 128-point representation so that the left boundary lines up with a LP vector at every level while the right boundary lines up with a HP vector at every level. All of the coefficients for the 128-point initial representation are displayed, however only 48 coefficients are used for representing the signal. The plots show that the left side of the signal has significantly better leakage performance than the right side. This is explained by the fact that the LP vectors at the left ends act as anchors for approximating the constant signal.

An important point in image compression is that it is not necessary to have zero DC-leakage into the highpass subbands – as long as the leakage is relatively small, then the adverse effects will be minor. When designing new filters for a particular application, one should incorporate some measure of either the LP vectors' ability to represent a constant signal or the amount of DC-leakage into the highpass subbands.

Stability It is desirable for the representation to be stable to perturbations in the signal. Specifically, similar signals should have similar representations, and similar representations should correspond to similar signals.

This proposed representation appears to be stable to perturbations in the signal and its support. A change in the signal's support requires a new basis. However, the change in the basis for this representation is exactly specified by the change in the support. Therefore, a change in the support will lead to a change of the same "size" in the basis.

A change in the signal's amplitude will give rise to a change in the coefficient amplitudes. The relative change in the coefficient amplitudes that results from a relative change in the signal amplitude is bounded by the condition number of the representation. The conditioning depends on the signal's support, the specific filters used, and the choice of the center tap. In the following, the center tap is chosen as the largest tap of the filter [12].

A number of experiments were performed to gain an understanding of the conditioning

[12] For all filters except Daubechies 6-tap and 14-tap, the "center tap" and the largest tap are equivalent.

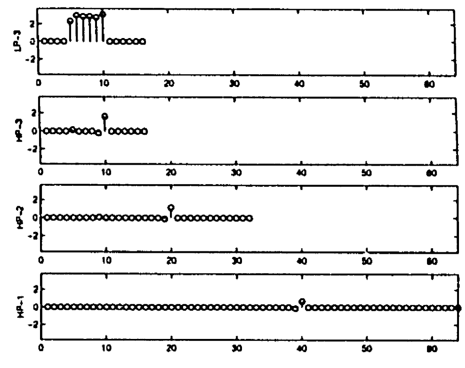

(a) Daubechies 8-tap Filters

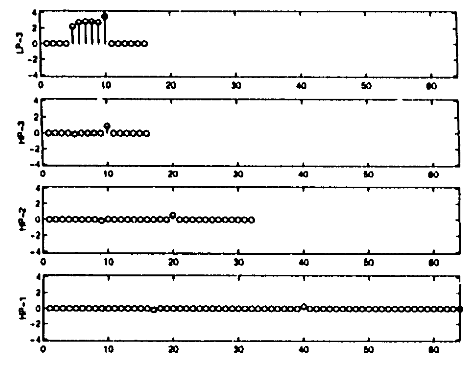

(b) Daubechies 10-tap Filters

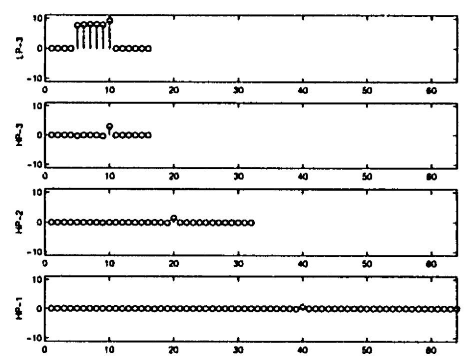

(c) Binary 9/7-tap Filters

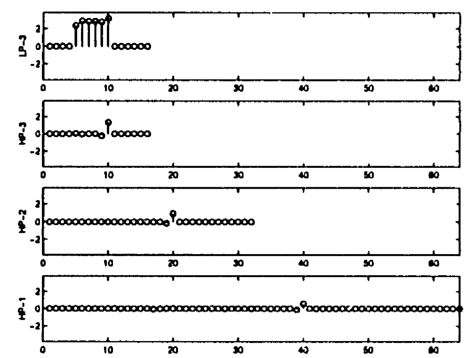

(d) Simoncelli 9-tap Filters

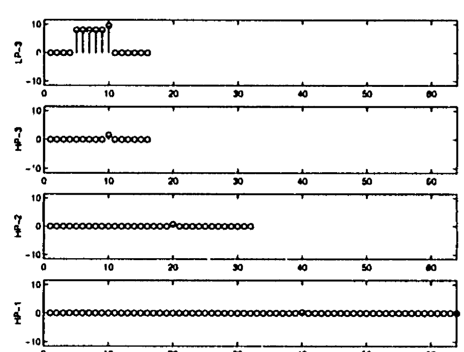

(e) Le Gall 3/5-tap Filters

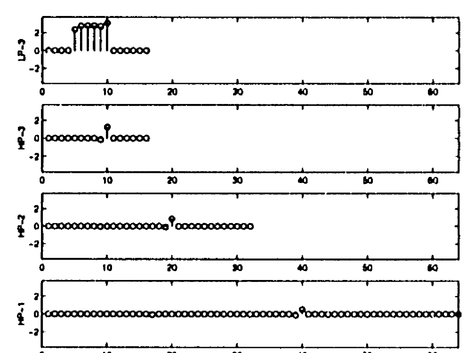

(f) Daubechies 9/7-tap Filters

Figure 4.8: Examining the DC-leakage for six different filter pairs. In each case a 3-level wavelet decomposition is computed and the resulting four subbands are plotted with the lowpass subband at the top and the highpass subband at the bottom. The original 48-point unit-amplitude signal is placed with respect to the initial representation so that the left boundary lines up with a LP vector at every level while the right boundary lines up with a HP vector at every level.

Filters	1-D Test Supports				
	Sig1	Sig2	Sig3	Sig4	Random
Daubechies 2-tap (Haar)	1.41	1.00	1.41	1.41	1.41
Daubechies 4-tap	1.15	1.15	1.04	1.15	1.40
Daubechies 6-tap	1.15	1.07	1.15	1.15	1.41
Daubechies 8-tap	1.17	1.17	1.06	1.17	1.41
Daubechies 10-tap	1.37	1.02	1.37	1.37	1.41
Daubechies 12-tap	1.18	1.18	1.07	1.18	1.41
Daubechies 14-tap	1.21	1.07	1.21	1.21	1.45
Daubechies 16-tap	1.18	1.18	1.08	1.18	1.41
Daubechies BO 9/7-tap	1.35	1.35	1.35	1.32	1.63
Smith-Barnwell 8-tap	1.19	1.19	1.06	1.19	1.42
Binary (Strang's 9/7-tap)	2.00	2.00	2.00	1.99	2.00
Simoncelli 9-tap	1.12	1.12	1.12	1.12	1.44
Le Gall 3/5-tap	3.99	3.99	3.99	3.99	3.98

Table 4.2: The condition number of the representation as a function of chosen filters, for four fixed 1-D arbitrary length supports as well as the largest (worst) condition number over 10,000 randomly generated 1-D arbitrary supports.

that the representation provides. The conditioning of the representation in 1-D and 2-D was examined for a number of different supports and all the filters previously considered. In each case the conditioning was computed for a single-level of the representation. In 1-D the conditioning was computed for four fixed finite-length 1-D supports as well as for a large number of randomly generated 1-D supports, and the results are given in Table 4.3.4. The four finite-length supports correspond to the four possible situations with respect to having either a LP or a HP vector serving as the anchor at each end of the finite-length support. Specifically, Sig1 has LP vectors at both the start and end of the support. Sig2 has a LP vector at the start and a HP vector at the end. Sig3 has a HP vector at the start and a LP vector at the end, and Sig4 has HP vectors at both the start and end. The Random result for each filter corresponds to the largest (worst) condition number measured over 10,000 randomly generated 1-D supports. Each of these randomly generated supports was 64-sample long where each sample has a probability of .5 of being within the support and a probability of .5 of being outside the support.

It is evident from Table 4.3.4 that most of the filters lead to condition numbers close to one. This is very encouraging, especially when considering that the largest condition number over the 10,000 randomly generated 1-D arbitrary supports for each filter is also reasonably close to one. An interesting observation is that the condition number for all of the orthogonal filters appears to be upper bounded by approximately $\sqrt{2}$. Understanding this intriguing result may shed some further light on the structure of the principal submatrices that result for orthogonal filters.

To examine the conditioning of the representation in 2-D, three fixed 2-D supports were considered as well as 1000 randomly generated 2-D arbitrary supports. The 2-D supports examined are shown in Figure 4.9 as well as one example of a randomly generated 2-D support. The random 2-D supports correspond to 20×20 sample regions where each sample in the region has a probability of .5 of being within the support and a probability of .5 of being outside the support. The same set of 1000 randomly generated supports was examined for each filter. For each filter, Table 4.3.4 lists the conditioning of the representation for each of the three fixed 2-D supports as well as the largest (worst) condition number out of 1000 randomly generated 2-D supports. The conditioning for each filter and the 1000 randomly generated supports is shown in Figures 4.10 and 4.11.

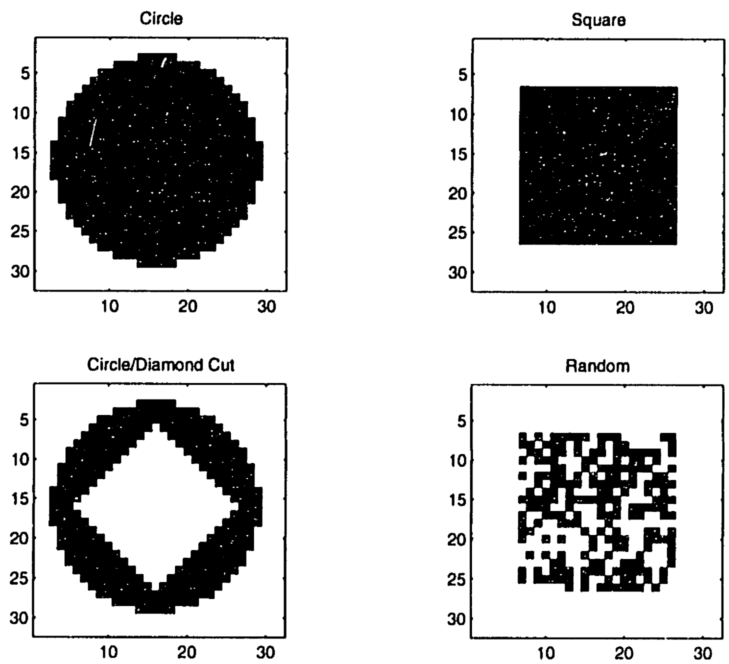

Figure 4.9: The 2-D supports used for examining the conditioning in 2-D: Three fixed 2-D supports, and an example of a randomly generated 2-D arbitrary support. The samples included within the support are shown in black and the samples not in the support are white.

The results of the 2-D conditioning test are similar to the results from the 1-D tests. To begin with, most of the condition numbers are relatively small. This is important for practical applications, and also lends further support to the hypothesis that the proposed representation provides a basis for a large class of filters and for arbitrary supports. For example, if the representation did not provide a basis for some choice of filter and some support, then the representation would be singular with an infinite condition number. If there was a case where the representation was ill-condition (very large condition number), that would correspond to a representation that was

Filters	2-D Test Supports			
	Circle	Square	Circle Diamond Cut	Random
Daubechies 2-tap (Haar)	–	–	–	–
Daubechies 4-tap	1.40	1.33	1.46	5.05
Daubechies 6-tap	1.48	1.15	1.55	8.32
Daubechies 8-tap	1.49	1.37	1.56	9.32
Daubechies 10-tap	1.94	1.05	2.04	40.13
Daubechies 12-tap	1.55	1.39	1.63	13.20
Daubechies 14-tap	1.62	1.15	1.71	29.29
Daubechies 16-tap	1.59	1.40	1.68	16.71
Daubechies BO 9/7-tap	1.95	1.80	1.82	9.76
Smith-Barnwell 8-tap	1.55	1.41	1.63	13.52
Binary (Strang's 9/7-tap)	3.93	3.90	3.47	9.72
Simoncelli 9-tap	1.50	1.24	1.57	15.32
Le Gall 3/5-tap	15.72	15.69	14.05	25.31

Table 4.3: The condition number of the representation as a function of chosen filters, for three fixed 2-D supports as well as the largest (worst) condition number over 1000 randomly generated 2-D arbitrary supports. The condition number was not computed for Haar, since there exist 2-D supports for which the representation using Haar does not provide a basis.

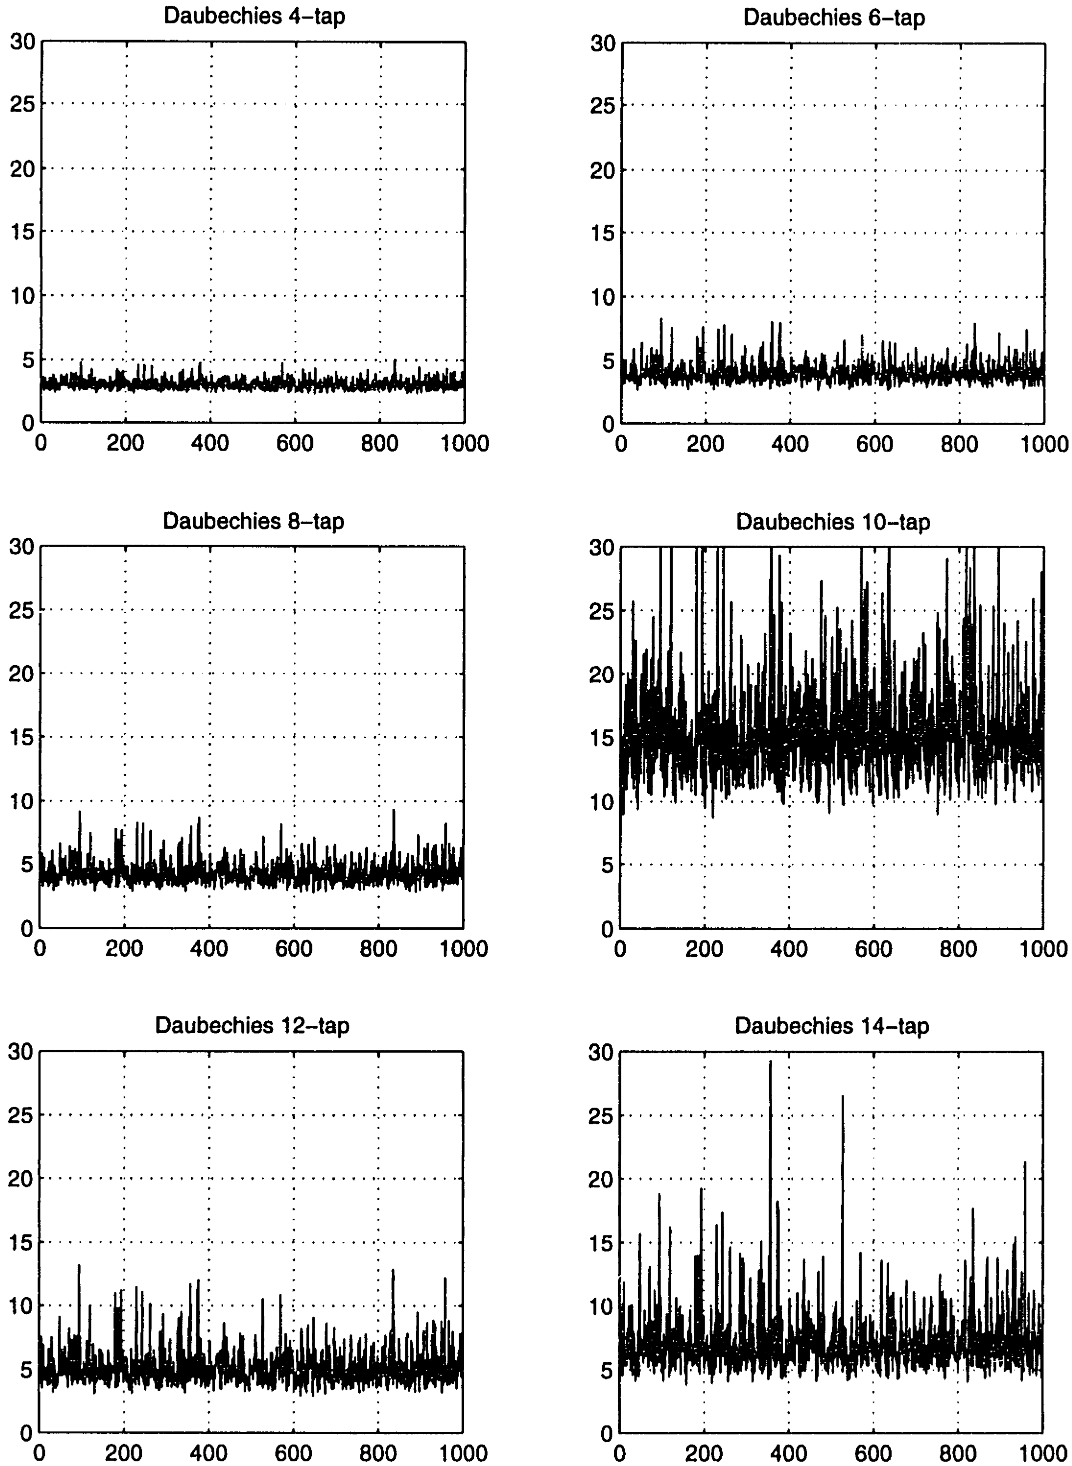

Figure 4.10: The condition number of the representation when using the first 6 filters from Table 4.3.4 and for the 1000 randomly generated 2-D arbitrary supports. The condition number is plotted along the y-axis for each of the 2-D arbitrary supports (x-axis). The plots display a maximum conditioning of 30 to improve the visualization. Note that Daubechies 10-tap filters lead to a condition number greater than 30 for 12 of the 1000 supports, with a maximum of 40.13.

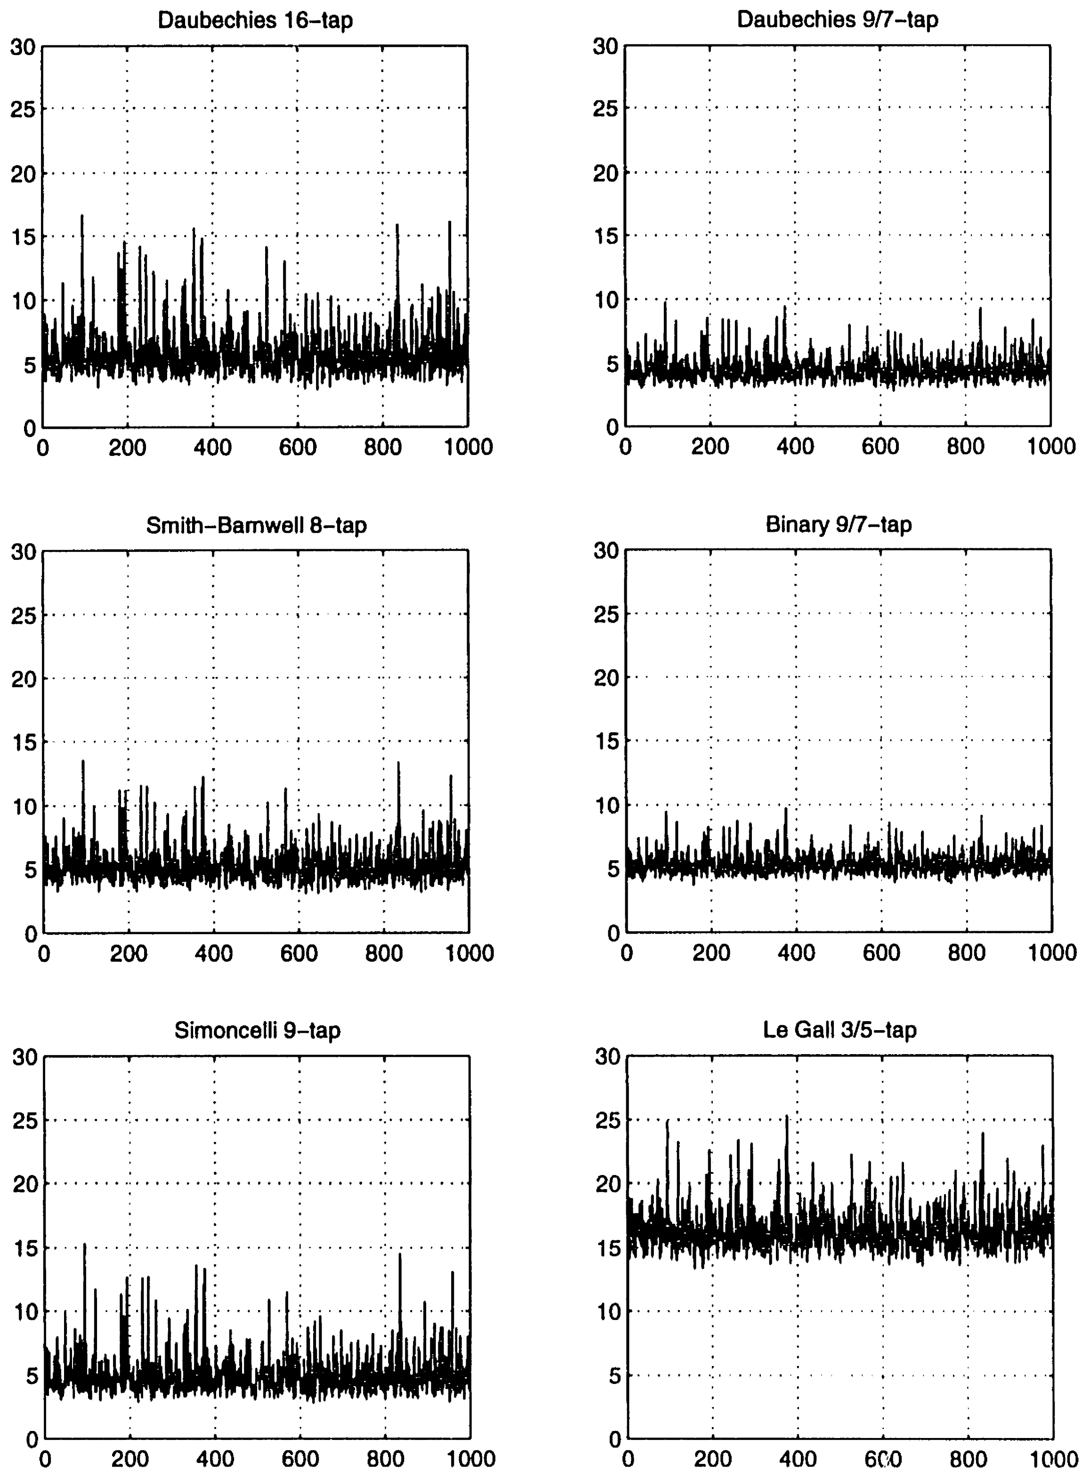

Figure 4.11: The condition number of the representation when using the last 6 filters from Table 4.3.4 and for the 1000 randomly generated 2-D arbitrary supports. The condition number is plotted along the y-axis for each of the 2-D arbitrary supports (x-axis).

close to singular, suggesting that there may exist a support that could make it singular. However, for every filter except Haar (which was shown not to be admissible) the representation provided a basis for every support tested, and the basis was reasonably well-conditioned in every case (maximum condition number measured was about 40). As an interesting sidenote, the Haar filter did not lead to a basis for any of the randomly generated supports tested.

A highly encouraging result from a practical point of view, is that for most of the filters, the condition number for the three fixed 2-D supports (circle, square, and circle/diamond cut) is less than two. These supports have similar structure to many of the 2-D signals with arbitrary supports that the representation is designed to represent, e.g. the 2-D objects or regions within an image or video. This suggests that the representation will be well-conditioned for representing these types of signals.

The results in Figures 4.10 and 4.11 provide an indication of the distribution of the condition numbers for a given choice of filters, as a function of the possible supports. An analysis of this form is beneficial for choicing an appropriate filter for a given application. For example, given everything else equal, one would prefer the filter with the lowest expected condition number as well as the lowest maximum possible condition number.

The results presented correspond to the conditioning for a single-level representation. A multiple-level representation will have a condition number that may be considerably larger. Of course, the condition number provides an upper bound on the relative change, therefore the relative change that will occur for a given support and a typical perturbation will probably be far less than these upper bounds.

A thorough analysis of the conditioning would include bounding the worst possible conditioning over all possible arbitrary ROS's. This is equivalent to determining the worst possible conditioning over all possible principal submatrices. This analysis should be performed for each filter of interest. A bound on the worst possible conditioning for a general filter is currently unavailable, and this is an interesting avenue for future research. Another important issue to examine is the relative change in the representation that results from a change in both the signal's amplitude and its support. The discussion here has shown that the change in the basis is stable to changes in the support, and that for a fixed support, the representation is well-conditioned for changes in the signal's amplitude. A thorough analysis of the stability of the representation to perturbations in both the signal's amplitude and support is another interesting avenue for future research.

4.3.5 Summary and Discussion for Representation # 2

The goal of this work was to create a critically sampled perfect reconstruction representation for signals with any arbitrary region of support. The underlying idea is to determine a basis for a given support by selecting the vector centered at each sample in the signal's support. That is, for each sample in the signal's support one selects the vector centered at that sample. This is a very simple selection method which for a large class of filters appears to provide a basis for completely arbitrary supports.

By centering (aligning) the vectors within the representation, we related the problem of whether the representation provides a basis for all possible supports to the problem of whether all the principal submatrices of an initial matrix are nonsingular. This relationship provided considerable insight into the success of the representation, and aided in establishing a number of conditions on the admissibility of a given filter pair. That is, conditions were determined for the filters such that if they are satisfied the representation will provide a basis for any arbitrary support when using those filters. In particular, for 1-D we have established a simple sufficient condition on the filter taps such that if it is satisfied the filter is admissible for all possible supports. Every filter that has been examined is admissible in 1-D. Thus, the representation can be used with any of the tested filters and guarantees a basis for any arbitrary 1-D support. The case of 2-D and general M-D is more complex than 1-D, and the admissibility of a filter in 1-D does not guarantee admissibility in 2-D or M-D, e.g. the Haar filter is only admissible in 1-D. Currently, the conditions for admissibility in 2-D/M-D have not been adequately simplifed to provide a simple test on the filter taps, and each filter must therefore be examined individually. Specifically, for 2-D/M-D we have only shown that we have complete freedom in the ROS for a small group of filters. However, both empirical evidence and theoretical considerations suggest that this is true for a much larger group of filters. An interesting avenue of future research is to identify what class of filters provides complete freedom in the possible supports that may be represented, and to develop a simple test on the filter taps to determine if a given filter is within this class.

Additional insights into the success of the proposed representation may be gained by considering the somewhat analogous problem of data interpolation in multidimensional space. For example, given a set of N data points (N amplitudes at N locations in 2-D space) the goal is to determine a linear combination of N vectors such that they pass through (interpolate) all of the data points. Splines are often used as the vectors in 1-D and radial basis functions in M-D (e.g. an M-D Gaussian kernel). The vectors are often designed so that if they are centered at the locations of the given data points, it is guaranteed that there exists a unique solution to the interpolation problem. For example, all the vectors are typically simple translates of a basic vector (e.g. a Gaussian kernel), and the basic vector is constructed so that all possible interpolation matrices are diagonally dominant and hence nonsingular. This construction guarantees that a unique solution always exists. Our problem differs in a number of ways. First, there is more than a single basic vector that is used for interpolation – a single-level wavelet transform in 1-D involves two basic vectors, in 2-D there are four basic vectors, and in M-D there are 2^M basic vectors. More importantly, the vectors in a wavelet transform are designed to have a variety of properties ranging from polynomial accuracy to good frequency selectivity, but no consideration is made for the nonsingularity of any interpolation matrices. Therefore, there is no apparent reason to suggest that the conventional wavelet vectors would be applicable/successful for this type of problem. Nevertheless, our proposed representation is analogous to an M-D interpolation problem over a discrete domain, and we have shown that wavelet vectors perform surprisingly well in this case.

In this representation, the set of selected vectors is equal to the region of support of the signal. More precisely, let M_{ROS} be a mask of 1's and 0's corresponding to the samples within and outside the signal's ROS, respectively, and let M_{select} be the mask of 1's and 0's describing the vectors to select and not select, respectively. With the vectors aligned (centered), the selection is given by $M_{select} = M_{ROS}$. For instance in the case of a 2-D signal (Equation 4.62), the set of coefficients of Y that are selected for representing a signal X is equivalent to the ROS of the signal X. This simple mapping between the support for a given signal and the vectors selected to create a basis for the signal holds true for 1-D, 2-D and general M-D signals. This simple relationship may prove to be useful in a number of different applications.

The proposed representation provides a number of desirable properties. For an appropriate choice of filters, it provides a basis for a signal with any arbitrary support. For example, the support may consist of a single contiguous arbitrarily shaped region – corresponding to the canonical arbitrarily shaped object/region in an image/video. In addition, the support can contain holes, it can be discontiguous, and it can include isolated samples. Therefore, for an appropriate choice of filters, the representation provides complete freedom in the possible supports that may be analyzed. Furthermore, the basis that the representation provides appears to be of high quality for representing the objects/regions in an image or video, which was the primary motivation for this work. For typical objects/regions from an image, most of the signal information is concentrated into a small fraction of the total number of coefficients, and by only retaining this small fraction of the coefficients one can reconstruct an accurate approximation of the original signal. This is analogous to the famous energy compaction property [1] for orthogonal transforms, which is the enabling property for most image/video [73] and audio [74] compression algorithms today. In addition, the representation expresses the signal at a number of different scales, and enables a natural reconstruction of the signal and its support (shape) at each of these different scales. The DC-leakage into the highpass subbands occurs only around the boundary areas, and for an appropriate choice of filters the leakage appears reasonably small (this partially explains the concentration of information into a small fraction of the coefficients). The representation also appears to be quite stable to perturbations in both the signal's amplitude and its support. Overall, this representation appears quite promising for representing signals with arbitrary supports, and in particular, for representing the objects/regions within an image or video.

4.4 Computational Issues

The complexity of the proposed representations depends on the complexity of three operations:

1. Determining the basis (selecting the vectors) given the signal's ROS.
2. Computing the coefficient amplitudes (of the selected vectors) to represent the signal.
3. Reconstructing the signal from the representation.

Typical applications, such as compression, require all three operations to be performed. However some applications, such as signal analysis, may only require the first two. Consider the general problem where the signal is analyzed, processed, and then synthesized, as illustrated in Figure 4.12. The analysis begins by determining a basis for the signal given the signal's support. Once the basis is found, the analysis computes the coefficient amplitudes to represent the signal. The coefficients are typically processed in some manner (e.g. compression) and then sent with a description of the signal's support (shape) to the synthesis. The synthesis begins by determining the basis for the signal given the signal's support. The synthesis then uses the basis and the coefficient amplitudes to reconstruct the signal.

An important consequence of our choice of a multiresolution representation is that the original problem is decomposed into a succession of smaller problems, each of which is solved as shown in Figure 4.12. To be precise, the operations shown in Figure 4.12 are performed at each level of the wavelet representation. In the following discussion, we consider how to analyze and synthesize a single level. While the following discussion is applicable to both of the proposed wavelet-based representations, we will focus our discussion on representation # 2 because it appears to be the most promising for practical applications.

Selecting the Vectors/Determining the Basis We select the vector centered at each sample in the signal's support. Thus, the vectors to select are precisely given by the support of the signal. This selection algorithm leads to an extremely simple, automatic, and unique selection of the vectors

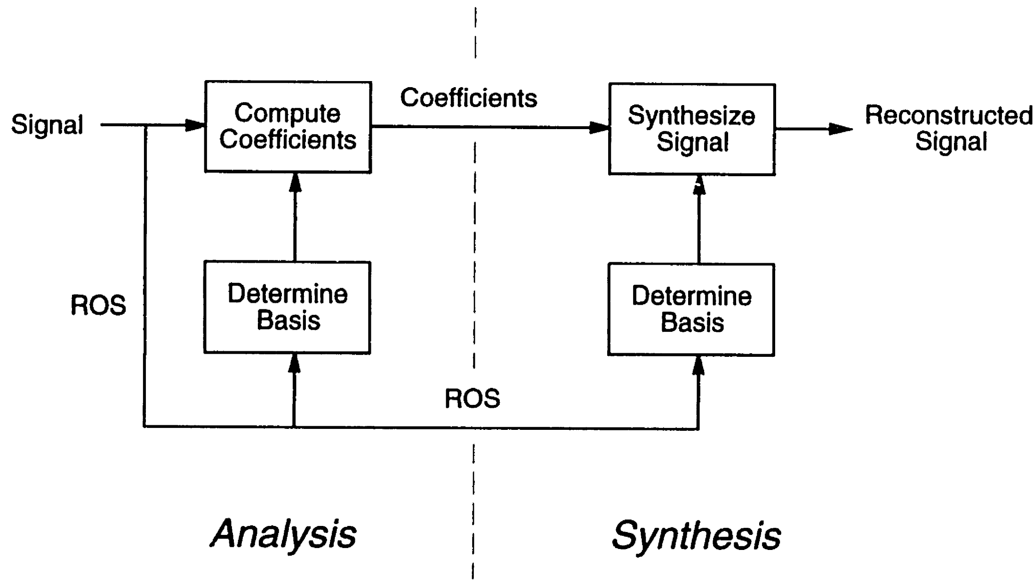

Figure 4.12: Overview of the operations for both the analysis and the synthesis.

given the signal's ROS. The simplicity of the selection is extremely important, because it must be performed at both the analysis and synthesis, e.g. both the encoder and decoder for a compression application.

Computing the Coefficient Amplitudes The coefficients can be computed with a number of approaches: (1) construct the dual vectors and compute their inner products with the signal, (2) set up and solve the linear system, (3) use iterative methods which exploit the structured initial basis, and (4) what we refer to as the "peel-away" method.

Constructing the dual vectors is equivalent to finding the inverse of the representation. This approach is typically prohibitively complex for any large signal. Setting up and solving a linear system is typically a complex and cumbersome process. Our choice of a wavelet representation with FIR filters corresponds to the linear system being highly sparse, which helps to simplify the problem. The special structure of the representation appears to be especially useful for an iterative approach to compute the coefficients [75, 76, 77]. In general, the primary difficulty with iterative approaches is that they require a large number of iterations to reach a reasonable estimate of the answer, and they are also somewhat unsatisfying since we never get the exact answer. However, if an acceptable level of accuracy in the solution can be achieved with low computational requirements, then this may be practical.

The "peel-away" method is easiest to visualize for an arbitrary-length 1-D signal, orthogonal filters, and a selection of only the LP boundary vectors as in representation # 1. First, we compute the interior coefficients in the conventional filterbank manner. We then reconstruct and subtract out the "interior" from the signal. The residual is non-zero only around the boundaries and it can be expressed using the selected boundary vectors. Since the selected boundary vectors have staggered supports, we can compute their amplitudes sequentially. For instance, the amplitude of the innermost boundary vector (longest overlap with the signal) can be computed independently of the other boundary vectors. In this way, we solve for the amplitude of each boundary vector and then subtract or "peel-away" its contribution from the residual. Then we continue with the next boundary vector. In essence there is a simple triangular linear system to be solved. The "peel-away" method may be beneficial for 2-D or M-D signals with large contiguous ROS's. The "peel-away" method would then solve for the interior, subtract it out, and then go around the boundary, spiraling or peeling-away outward. Depending on the specific ROS and the choice of filters, the peel-away method may be used in conjunction with the previously discussed approaches to potentially provide an efficient method for computing the coefficient amplitudes.

The appropriate computational approach for a given problem will depend on a number of factors such as the problem's dimensionality (1-D, 2-D, etc.), the character of the region of support (e.g. single large contiguous region or many small regions), properties of the filters (e.g. orthogonality, staggered supports), required accuracy in the solution, and computational and memory constraints.

Synthesis The synthesis stage determines which coefficients were selected (performs the selection based on the ROS) and reconstructs the signal from the coefficients. The selection is very simple as previously discussed. The inverse transform is a simple, conventional, one-level inverse wavelet transform followed by cropping the reconstructed signal to the original ROS. Thus decoding a 2-D signal with an arbitrary ROS is almost as simple as decoding a 2-D signal with rectangular support.

Approximate Computation of the Coefficient Amplitudes In a number of applications, such as compression, it is unnecessary to determine the exact solution. That is, an approximate answer may be sufficient. This may be exploited in a variety of ways. For example, an appropriate extrapolation of the signal can be computed, and then a conventional one-level filterbank can be applied. This will provide a reasonable initial estimate of the coefficient amplitudes. Alternatively, for a large contiguous support one can skip the extrapolation and instead use the interior coefficient amplitudes to predict the boundary coefficients. A prediction of this form should work reasonably well for smooth signals such as images. Both of these approaches would probably give reasonable estimates of the coefficient amplitudes. Probably only the coefficients at the boundary would need to be improved, and this may potentially be achieved by applying a few iterations of an appropriate iterative scheme. Since coefficient refinement is only required in the boundary areas, it may be possible to only apply computations in that area, corresponding to "partial iterations" requiring lower complexity. As a final interesting note, an error in a boundary coefficient will have a smaller impact on the reconstructed MSE than an error of the same size on an interior coefficient. Therefore the boundary coefficients do not need to be estimated with the same accuracy as the interior coefficients.

4.5 Summary

In this chapter we explored the possibility of creating wavelet-type representations for signals with arbitrary supports within the context of the general framework presented in Chapter 3. In particular, we introduced two novel wavelet-type representations for signals with arbitrary supports.

The first representation was designed to preserve the wavelet's polynomial accuracy property. By retaining all the lowpass vectors that overlap the signal, the representation inherits all of the polynomial accuracy properties of the wavelet representation for infinite-length signals, i.e. the information in a polynomial signal is concentrated into the lowpass subband and the lowpass subband is a polynomial of the same order. This leads to a simple necessary and sufficient condition on the set of selected vectors such that the representation preserves the polynomial accuracy: all the LP vectors that overlap the signal's support must be selected. This condition leads to constraints on the possible supports that may be represented: 1-D finite-length signals with a minimum length, and 2-D or M-D signals with arbitrarily shaped supports with constraints on the local boundary shape. Two difficulties with this representation are that for long filters (1) the range of possible 2-D and M-D supports is highly constrained, and (2) the representation can be highly ill-conditioned. These difficulties may limit the potential applications of this representation.

The second representation was designed for signals with completely arbitrary supports. This approach provides a number of advantages: complete freedom in the possible support that may be analyzed (for appropriate choice of filters), a very simple selection of the vectors to provide the basis given the support, a potentially large class of filters to choose from, a natural multiresolution representation of the signal and its support, good stability to perturbations in the signal's amplitude and its support, and high concentration of information into a small fraction of the coefficients for typical signals. Overall, this representation appears very promising as a general purpose wavelet-type representation for signals with arbitrary ROS's, and appears particularly promising for representing the objects/regions within an image or video.

These representations provide, in a sense, opposite extremes in terms of some of the desirable properties. The first representation fully preserves the polynomial accuracy property, at the expense of severe constraints on the possible supports that may be represented. The second representation provides complete freedom over the supports that can be represented, but does not preserve polynomial accuracy. An interesting observation along similar lines is that while the Haar filters are arguably the "worst" filters to use with the second representation for a 2-D signal, they are also arguably the "best" filters to use with the first representation for the same 2-D signal.

Finally, we remark that within the general framework presented in Chaper 3 there most likely exists a number of potentially useful wavelet representations. In particular, one may design a different representation which provides a different set of desired properties. As a very simple example, even the second representation can lead to a number of variations. Specifically, there are a number of ways to choose the "center" of the filter that lead to admissible filters. We have focused on centering at the largest tap, since it leads to a natural notion of the appropriate vectors to select and also typically provides a very well-conditioned representation. However, other choices of the center may potentially lead to very useful properties, e.g. complete freedom in the possible supports and partially preserving the polynomial accuracy. Therefore, depending on the application, a yet undiscovered representation may achieve better performance than the two representations presented in this chapter.

Chapter 5

Application to Image/Video Compression

In this chapter we explore the possible application of the representations developed in this thesis to image/video processing and specifically toward image/video compression. This chapter begins by providing a high level overview of image/video compression and the conventional compression algorithms. The notion of object/region-based compression is then presented and its potential benefits are briefly discussed. There are a number of fundamental issues that must be addressed for the success of object/region-based compression, and we focus on the problem of representing the interior of the arbitrarily shaped 2-D objects/regions in an image or video scene. The work in this chapter is preliminary in the sense that developing a high quality image/video compression system is a complete system problem, and any work on one element of the system, no matter how detailed, is preliminary unless incorporated and examined in the context of the entire system. Nevertheless, this chapter highlights some of the distinguishing features of the representation developed within this thesis and discusses its potential usefulness for compression.

5.1 Overview

One goal of source coding or compression is to reduce the *redundancy* that exists in the source in order to transmit or store the signal at a lower bit rate. The redundancy in a video signal is evident in consecutive frames of video, which are quite alike. Also within any single frame there are large regions (objects or background areas) that have similar characteristics. This is redundant information. It would be wasteful of channel bandwidth to transmit the same information repeatedly. An important issue in image/video compression that is not inherent to all source coding schemes is the question of what is *perceptually relevant* and what is not. Since the criterion for quality of the video signal is the human visual system, rather than any analytical metric such as mean square error, what the human visual system finds relevant or irrelevant is extremely important. For example, human spatial acuity is much less for randomly moving imagery than for still imagery. This tradeoff of spatial and temporal resolution can be used to great advantage in video compression systems. Therefore, the goal of source coding as applied to image/video is to reduce the redundancy and irrelevancy, or equivalently, to transmit only the essential and relevant information.

A video signal is a continuous function of time, space, and wavelength, which is acquired and discretized for digital processing and transmission. This intrinsically involves sampling and quantization of each of the video dimensions. Temporal sampling is performed when creating the individual frames of video. Spatial sampling of these frames results in image samples, often referred to as picture elements or *pixels*. The temporal and spatial sampling structures may be independent, or they may be coupled together as in interlaced scanning used in conventional television. To represent color, video is usually modeled as the additive combination of three primary colors: red, green, and blue. Each image sample is composed of three color components with a finite number of (often 8) bits of accuracy each.

5.1.1 Redundancy and Irrelevancy

The goal of video compression is to reduce the redundancy and the irrelevancy inherent in the video signal [78, 1]. The sources of redundancy include:

- Temporal: Most frames are highly correlated with their neighbors.
- Spatial: Nearby pixels are correlated with each other.
- Color space: RGB components are correlated among themselves.

A high-performance video compression system identifies and exploits each of the redundancies within the video signal.

The redundancies are relatively easy to identify and exploit for compression. However, the question of what is relevant or what we can see is much harder to quantify. The human visual system (HVS) is a complex biological process that does not lend itself easily to analytical modeling. The various forms of relevancy can be decomposed into areas based on our temporal perception, spatial perception, and color perception. Also, the spatial and temporal masking phenomena of the HVS are important because they can be used to direct our attention away from a particular distortion within the video, thereby hiding or masking it. The important perceptual elements include sensitivity to motion, spatial frequency, color, and brightness. Irrelevancy can be viewed as a form of *perceptual redundancy*, where an element is represented with more resolution than is perceptually required. Conventional compression algorithms incorporate relatively basic spatiotemporal models of the HVS in an attempt to minimize the perceived distortion in the reconstructed video signal.

5.1.2 Principles of Video Compression

A digital compression system can be expressed as the combination of three distinct, though interrelated, operations: *Representation, Quantization,* and *Codeword Assignment.* These operations are depicted in a general digital compression system shown in Figure 5.1. In the first stage, the signal is expressed in a more efficient representation which facilitates the process of compression. The representation may contain more pieces of information to describe the signal than the signal itself, but most of the important information will be concentrated in only a small fraction of this description. In an efficient representation, only this small fraction of the data must be transmitted for an appropriate reconstruction of the signal. The primary approaches taken toward creating an efficient and compact representation may be classified broadly into the categories of *predictive processing, transform/subband filtering,* and *model-based processing.* These approaches may be intermixed and applied along the temporal, spatial, and color space dimensions of the video signal. The second operation, quantization, performs the discretization of the representation information in order to enable transmission over the digital channel. The information may be quantized one parameter at a time, *scalar quantization,* or a group or vector of parameters may be quantized jointly, *vector quantization.* The third operation takes the quantized parameters and assigns to each an appropriate codeword for transmission. In the simplest case *fixed-length codeword assignment* may be performed, or more complex *entropy coding* (e.g. Huffman or arithmetic coding) may be applied to exploit the statistical redundancy of the quantized parameters and thereby reduce the average bit rate.

Each operation may be designed to exploit both the statistical redundancy and the perceptual irrelevancy inherent in the video signal. The first and third operations may be performed in a lossless manner, with any loss of information being localized solely within the quantization operation. By isolating the potential loss of information in a single operation, a simpler design process and fine tuning of the system are possible. The three operations are so closely related that highly optimizing one may potentially reduce the gain from the others. Maximum performance is of course achieved by jointly optimizing all three operations. However, complexity issues typically result in each operation being optimized independently of the others.

5.1.3 Conventional Compression Algorithms

The following contains a very brief overview of conventional image and video compression algorithms. The basics of a gray-scale image coder are discussed first, followed by a description of the additional processing needed for coding a color image and a video signal. More detailed

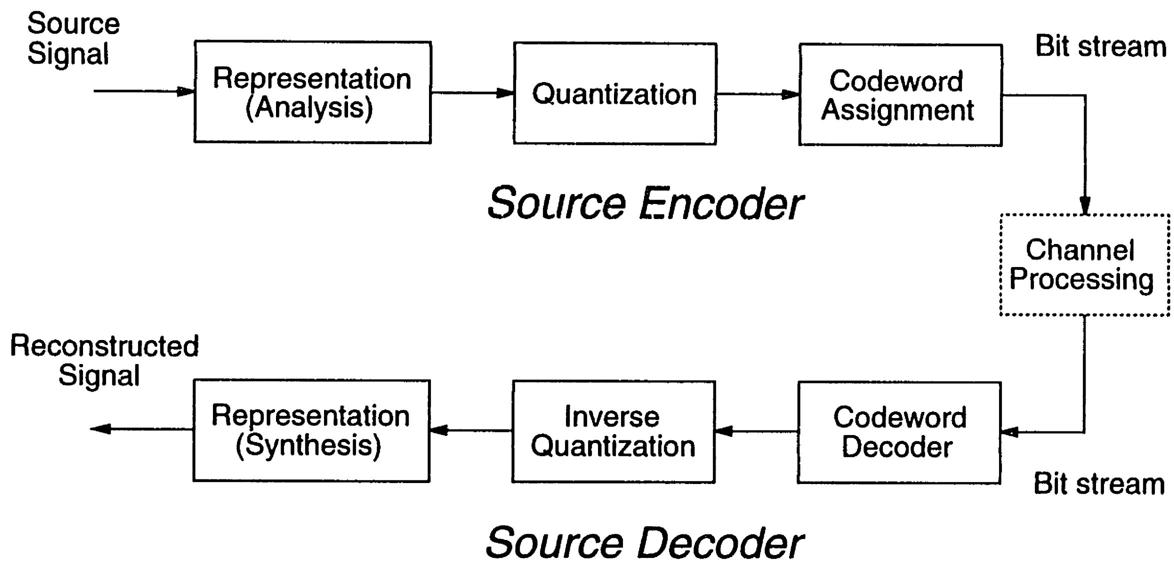

Figure 5.1: An overview of a digital compression system showing the three basic operations at both the encoder and the decoder. The operations at the decoder are the inverse of those at the encoder.

descriptions and analysis can be found in a number of sources, e.g. [1, 73, 79]. These ideas form the basis for a number of international standards for compression of images and video, such as the JPEG image compression standard and the MPEG-1, MPEG-2, CCITT H.261, and ITU H.263 video compression standards [1].

A conventional gray-scale image coder begins by transforming the image to another domain where most of the signal energy is concentrated in a small fraction of the coefficients. The most popular approach is to partition the image into separate 8×8 pixel blocks and then compute the 2-D DCT of each block. This approach is often referred to as the 8×8 Block DCT. Other popular spatial transforms include lapped transforms and wavelet transforms. Once the image has been transformed, the transform coefficients are quantized and entropy coded. To be precise, conventional coders scale each coefficient by an appropriate factor to account for the HVS and the available bitrate or buffer fullness, and then pass each coefficient through a uniform quantizer. Since the great majority of the coefficients are quantized to zero, only the locations and amplitudes of the non-zero quantized coefficients are coded. For example, the coefficients are typically scanned in some order, and the number (runlength) of consecutive zero coefficients before a non-zero coefficient is coded, followed by the non-zero coefficient amplitude. Finally, the various parameters

[1]JPEG is the Joint Photographic Experts Group. MPEG is the Moving Pictures Expert Group. CCITT is the International Telegraph and Telephone Consultative Committee, which has recently been renamed to the International Telecommunication Union (ITU).

are typically Huffman coded to exploit their statistical properties and reduce the average bit rate.

The coding of a color image typically begins by transforming from the red, green, and blue (RGB) color space to the YIQ or YUV color space. This transformation reduces the correlation among the three color components and also enables subsequent processing to exploit the differing HVS response to the luminance and chrominance components. Each of the luminance and chrominance components is then coded via an appropriately-tuned gray-scale image coder as described above.

A video signal corresponds to a sequence of images acquired and displayed in rapid succession to give the impression of continuous motion. The high frame rate necessary to achieve proper motion rendition usually results in considerable temporal redundancy among the consecutive images. Specifically, consecutive frames typically contain the same information, possibly with movement between frames. Therefore, predictive processing approaches are typically applied, where they use the previous frame to form a prediction of the current frame to be encoded, and then only code the prediction error or residual. The accuracy of the prediction is greatly improved by accounting for the motion between consecutive frames. The motion between the previous frame and the current frame is estimated (motion estimation, ME), and then the prediction is formed while compensating for the presence of motion (motion-compensated prediction, MC-P). The error in the MC-prediction, referred to as the MC-residual, is then coded via an image coder appropriately-tuned to the characteristics of the MC-residual signal. ME/MC-P is typically applied to a video signal by partitioning each frame into square blocks and applying the ME/MC-P to each block. The first frame in a video signal is coded as a normal image (separate from any of the other frames) and is used as an initialization to apply the predictive processing to the subsequent frames. If the MC-residual for a frame is more difficult to code than the original frame, then the MC-processing is turned off and the original frame is coded.

5.1.4 Object/Region-based Compression Algorithms

Most of the conventional image/video compression algorithms involve block-based or overlapped-block-based schemes. The image is partioned into blocks (possibly overlapping) and each block is processed independently of the others. All conventional Block-DCT, lapped transform, and wavelet image/video compression algorithms can be viewed as block-based or overlapped block-based schemes. Block-based processing achieves good performance while allowing architectural simplicity. However, these block-based schemes do not exploit (and in fact neglect) the actual content of the image or video. In effect, these approaches implicitly assume a source model of moving square blocks. However, a typical video scene is not composed of moving square blocks.

Therefore, block-based schemes impose an artificial structure on the video signal and then try to encode this structure, as opposed to recognizing the structure inherent to a particular video scene and attempting to exploit it.

An improved source model may be developed by identifying and efficiently representing the structure that exists within a video scene. For example, since real scenes contain objects, a promising source model is two or three dimensional moving objects. This approach may provide a much closer match to the structure in a video scene than the aforementioned block-based schemes. This source model is illustrated in Figure 5.2.

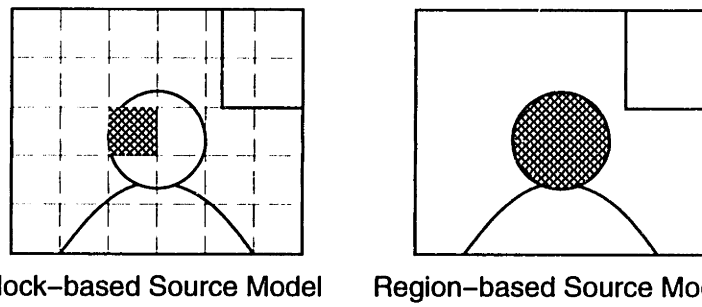

Figure 5.2: Block-based versus object- or region-based source models.

An image can be decomposed into 2-D objects or regions. A video signal can be decomposed into 3-D objects or regions in a number of different ways. The simplest approach conceptually, is to identify the various 3-D objects in the scene. However, this is typically extremely complex. An alternatively approach is to identify the 2-D objects or regions in a single frame, and track their evolution with time. The resulting signal may be thought of as a 3-D signal with a (somewhat) arbitrary support. However, a more practical viewpoint would be a 2-D signal that evolves with time, specifically both the intensity within the region and the region shape can change with time.

An improved source model of this form leads to a more natural signal representation which may lead to improved compression and possibly more importantly may facilitate higher-level processing. For example, the content in the scene may be identified (either automatically or by human intervention) and the scene may be processed based on its content, providing a form of *content-based processing*. This content-based image/video representation may facilitate novel algorithms and applications. The individual objects/regions in an image or video may be individually extracted and manipulated, with different processing for each.

An object/region-based coder may potentially provide improved compression. For example, an accurate estimate of each object's shape and motion within the sequence can lead to a significant gain in MC-prediction performance. If the characteristics of the interior of an object is homogeneous, this may be exploited for coding the interior. Also, the artifacts that result in an object/region-based coder will probably be much more benign than the highly structured and artificial blocking and mosquito noise artifacts that afflict conventional block-based coders. In addition, an object/region-based coder also enables direct access and manipulation/processing of the objects from the compressed bitstream, thus providing a simple and natural object-based scalability at the bitstream level. Currently considerable research is being performed toward developing object/region-based compression algorithms both within the general research community and within the MPEG-4 standardization process [80].

Object/region-based image/video processing appears very promising, however two fundamental issues must be addressed before it will be successful:

1. The video scene must be segmentation into objects
2. Each object must be processed appropriately

The type of processing of course depends on the particular application. In the case of video compression, the required processing includes

1. Efficient tracking and motion-compensated prediction of the various objects
2. Efficient coding of the segmentation information (object shape or ROS)
3. Efficient coding of the object interiors (amplitude of the signal)

A high performance and computationally efficient transform/subband approach for representing signals with arbitrary supports may greatly assist the coding of the object interiors. Furthermore, a representation of this form may facilitate many other types of object/region-based processing such as restoration, extrapolation, texture analysis, etc.

5.2 Representing the Interior of Objects/Regions

The underlying goal of a representation developed for image/video compression is to reduce the redundancy by concentrating the perceptually important information into a small fraction of the parameters. This small fraction may be further processed for transmission while the remaining elements may be simply discarded. The chosen representation, as well as the amount of required compression, are dependent upon the specific application and its implementation constraints. In this section we consider the problem of representing the interior (or so-called "texture") of the objects/regions within an image or video scene. This is equivalent to representing the amplitude of a 2-D signal with an arbitrary support.

5.2.1 Previous Approaches

The problem of representing the signal amplitude of an arbitrarily shaped object or region has been discussed for many years, and a multitude of different approaches have been proposed. All of the approaches developed for representing an image with square or rectangular support have also been considered for representing the arbitrarily shaped objects or regions. In addition, each of these approaches have been modified in a variety of ways to account for the arbitrarily shaped supports. These approaches were discussed in Section 2.4, and here we will briefly summarize the various approaches and highlight some of their attributes.

The earliest approaches approximated the signal's amplitude as a low-order polynomial [5]. These approaches did not provide perfect reconstruction capability. Another set of approaches is based on the idea of actually constructing a basis for the given signal ROS. For example, one can begin with a set of vectors that spans the signal space, and apply an orthogonalization procedure to construct an orthogonal basis over the signal space [37], or model the signal as an appropriate stochastic process, and compute the KLT over the ROS (see, e.g. [35, 36]). A difficulty with these approaches is that they are quite complex, with complexity at both the encoder and decoder. Low complexity approaches based on 1-D transforms have also been proposed [32, 34], however they possess certain undesirable properties.

A natural approach for representing a 2-D signal with an arbitrarily shaped support is to use a transform/subband scheme defined over a circumscribing square. For example, the signal can be extrapolated to fill the circumscribing square, and a conventional transform/subband representation can be computed over the square. A number of approaches have been proposed along these lines, and they differ in how they use the superset basis to represent the signal, how the extrapolation is performed, how the important coefficients are identified, and how their amplitudes are computed [60, 62, 61, 48, 47, 65, 45]. These overcomplete approaches guarantee perfect reconstruction (if all of the coefficients are retained), they are in principle applicable to any ROS, and they facilitate the use of the separable/fast properties of the underlying superset basis. However, their overcomplete nature may make processing more complex, and the greater number of possible parameters to encode may lead to a potential inefficiency in compression.

5.2.2 Proposed Approach

The general framework presented in Chapter 4 can be used to create a number of different representations that may be useful for image/video compression. However, in this chapter our discussion will focus on the second wavelet representation presented in Section 4.3. This representation appears to be promising for representing the objects/regions in image/video.

5.2.3 Comments on Coding the Interior of Objects/Regions

The goal of image/video compression is to achieve the highest visual quality for the available bit rate. A number of comments are in order before evaluating the various possible approaches. First of all, as previously mentioned, there exists no completely satisfactory metric for visual quality [2]. To facilitate the development and analysis of compression algorithms, researchers have focused on mathematically tractable quality metrics such as MSE or weighted MSE (which incorporates the human visual system's spatial, temporal, and color sensitivities and signal-dependent masking effects). Developing improved visual quality metrics is an important area of current research. To simplify the discussion, we restrict ourselves to MSE as the quality metric.

Another issue relates to determining the optimal solution for a given signal and an available bit rate. Given a compression algorithm and some tractable quality metric such as MSE, determining the optimal method to compress a given signal is an extremely difficult problem. That is, determining the optimal method to use the available bits to maximize the quality metric is extremely difficult. This results since the typical compression problem corresponds to an (unstructured) discrete optimization problem, and requires exhaustive search to determine the optimal solution. Efficient algorithms have been developed for specific subproblems (e.g. [81, 82]), however determining the optimal solution for a general compression system is still impractical.

The above two issues help illustrate the art of image/video compression. Specifically, an optimal measure of visual quality is unknown, and even given some measure of quality, determining the optimal method of allocating the available bits for a given signal is impractical. Therefore researchers attempt to create algorithms that achieve high-quality performance with low complexity. This is an important point since the practical significance of our work in the context of compression depends on the development of high-quality, low-complexity algorithms for compression.

[2] One may argue that there will never be a completely satisfactory metric for visual quality, for the simple reason that different people have different preferences, and no metric will be satisfactory for everyone.

Orthogonal vs. Biorthogonal Representations One of the most successful algorithms in image/video coding is based on the notion of energy compaction and the idea of threshold coding. Consider the case of an orthonormal transform/subband scheme. Typical images have most of their energy concentrated in a small fraction of the coefficients. A high-quality approximation of the signal can therefore be obtained by selecting the most energetic coefficients and discarding the rest. This simple approach is optimal in terms of determining the appropriate subset of coefficients for minimizing the MSE in the reconstructed signal. However, it does not address the question of minimizing the MSE for a given bit budget. Nevertheless, this approach performs surprisingly well. Specifically, selecting and coding the most energetic coefficients performs surprisingly close to the optimal solution. This is especially striking considering the simplicity of the approach.

It would be advantageous if threshold coding could also be successful applied for coding the interior of the arbitrarily shaped objects in an image or video scene. However, virtually all the proposed representations for the arbitrarily shaped objects are biorthogonal instead of orthogonal. This can lead to a number of complications. In the following we briefly compare the properties of orthogonal and biorthogonal representations.

An orthonormal transform/subband representation greatly simplifies the analysis. Consider the problem of representing an N-point signal using only k coefficients ($k < N$). Specifically, the goal is to choose k coefficients and their amplitudes in order to minimize the MSE of the reconstructed signal. An orthonormal representation provides the following important properties.

1. The k most important coefficients are given by the k most energetic (largest magnitude) coefficients.
2. The optimal amplitudes for the k coefficients is their original amplitudes.
3. The energy of the reconstruction error is equal to the energy of the discarded coefficients.

All of these properties follow because the representation is orthonormal, and we have the equivalent of Parseval's theorem: the energy in the coefficient domain is equal to the energy in the signal domain. The contribution of each coefficient for representing the signal is given by its energy.

For a biorthogonal representation, the situation is much more complex. Assume given a biorthogonal transform/subband representation of a signal, and the goal is once again to choose k coefficients and their amplitudes in order to minimize the MSE of the reconstructed signal.

1. The k most important coefficients are *not necessarily* given by the k largest magnitude coefficients.

2. Given a set of k coefficients, their optimal amplitudes for reconstructing the signal are *not* given by their original amplitudes.

3. The energy of the reconstruction error is *not* equal to the energy of the discarded coefficients.

To be more specific, determining the optimal set of k coefficients for representing a given signal is equivalent to determining the k-dim subspace (formed by k vectors from the transform/subband representation) that lies closest to the given signal. This is a combinatorial problem, and in general requires an exhaustive search to find the optimal solution. Given any set of k coefficients, their optimal amplitudes for approximating a given signal is given by the projection of the signal onto the k-dim subspace formed by their vectors. Determining the optimal coefficient amplitudes is therefore equivalent to solving a least-squares problem. For any set of k coefficients and their amplitudes, the energy in the reconstruction error must be determined by actually reconstructing the signal with the k coefficients and then computing the error, i.e. the MSE can not simply be determined by examining the amplitudes of the discarded coefficients.

These issues also arise when quantizing the coefficients: the quantization error (distortion) of the reconstructed signal is not equal to the quantization error of the coefficients. Specifically, the quantization noise may be either *amplified* or *attenuated* when going from the coefficient to the signal domain.

In summary, a biorthogonal representation leads to a number of possible complications over an orthogonal representation. If the biorthogonal representation is close to orthogonal, then one can often process it in the same manner as for an orthogonal representation without a significant loss in compression efficiency. However, if the biorthogonal representation is far from orthogonal, then the required processing may become much more intricate.

Representing the MC-residual for an object/region In video compression, the MC-residual for each object/region may be represented using the proposed approach or another approach specifically designed for representing residual signals. For example in conventional video coding, a representation based on an overcomplete set of 2-D modulated Gaussians appears in some cases to provide better performance than a critically sampled approach [83]. Loosely speaking, an overcomplete representation is sometimes advantageous at very low bit rates. At very low rates only a small number of vectors can be coded, and the overcomplete representation provides more freedom for choosing the few crucial vectors to code – this benefit may outweigh the extra overhead (location information) required to describe which vectors were chosen from the overcomplete representation. When a residual signal has energy concentrated in a few spatial locations, a representation that can provide a vector similar to the error at each of these locations is advantageous, e.g. a modulated Gaussian centered at each pixel location. At higher rates or for a high-energy residual signal where the energy is not well localized, the extra overhead of an overcomplete representation may be detrimental and a critically sampled representation may achieve higher performance [84].

5.2.4 Experimental Results

This section presents a simple comparison of the proposed approach against some alternative approaches with similar structures. Our *proposed approach* begins with an initial transform over a square and subsequently selects a subset of vectors to provide a basis for the given signal. A natural alternative approach is to use the initial transform itself for representing the signal in a conventional manner. Computing the transform in this fashion requires defining the samples outside the signal's support but within the square. The simplest approach to define these samples is to assume they have an amplitude of zero. This is called the *zero-extrapolation approach*. This approach would not be used in an image compression system since the discontinuities that result at the signal boundaries produce artificial high frequency components and degrade the energy compaction. A more compression-friendly approach would be to perform a smooth extrapolation to fill the area outside the signal's support such that there are no discontinuities at the signal boundaries. Specifically, the mean is extrapolated in each local area in a smooth manner, similar in spirit to, for example, [85]. This will be referred to as the *Extrap-AS-ROS approach*.

Thus, both of these approaches use the same vectors at each level as in our proposed approach, however while our approach only selects and uses an appropriate subset of vectors for representing the signal, these approaches use all the vectors. Both of these approaches are overcomplete representations since they use more coefficients to represent the signal then samples in the original signal.

A critically sampled approach using the same filters can be obtained by applying a 1-D two-channel filterbank first to each row, and then to each column. Specifically, each row or column is treated as a finite-length 1-D signal and symmetric extension is applied at each boundary prior to the application of a conventional 1-D two-channel filterbank to the 1-D signal. This approach will be referred to as the *1-D Row/Col approach*. This approach does not provide perfect reconstruction unless every row and every column of the region is of a minimum length dictated by the filter lengths. This minimum length constraint must hold at each level of the decomposition.

The proposed representation, the extrapolation-based representations, and the 1-D Row/Col approach are related in that they all attempt to create wavelet-type representations for the given signal. In all cases, the same lowpass/highpass filter pair is used. The proposed representation is critically sampled and provides perfect reconstruction and is applicable to any 2-D ROS. The extrapolation-based representations are overcomplete, but they provide perfect reconstruction if all of the coefficients are retained and they are applicable to any 2-D ROS. The 1-D Row/Col approach is critically sampled, however it does not provide perfect reconstruction except for very limited ROS's. In particular, its performance degrades quickly for signals outside a limited range of ROS's.

To examine the potential compression performance of each approach, three 2-D signals with varying degrees of arbitrary supports were examined. The signals are "Face", "Lenna", and "Star" and are shown in Figure 5.3. To compare the different approaches, we define quality measures for signals defined over arbitrary regions of support. If $x(n_1, n_2)$ is the original signal and $\widehat{x}(n_1, n_2)$ is the reconstructed signal, then the mean square error (MSE) or distortion of the reconstructed signal is given by

$$\text{MSE} = \frac{1}{|S|} \sum_{(n_1,n_2) \in S} (x(n_1, n_2) - \widehat{x}(n_1, n_2))^2 \tag{5.1}$$

where S defines the region of support of the original signal and $|S|$ is the cardinality of S, that is the number of samples in the region of support of the signal. The peak signal-to-noise ratio (PSNR) is defined as

$$\text{PSNR} = 10 \log(\frac{255^2}{\text{MSE}}). \tag{5.2}$$

Note that the MSE and thus also the PSNR only depend on the reconstruction of the signal over its support.

To compare the performance of the different approaches, the representation of each signal was first computed using each approach. In each case, a three level transform is computed using Le Gall's short symmetric kernal filters [72]. These are biorthogonal linear phase filters with 5/3 taps for the lowpass/highpass analysis filters and 3/5 taps for the lowpass/highpass synthesis filters. The reconstructed quality was then examined for retaining a prescribed number of coefficients and discarding the rest. Specifically, for each representation the k largest amplitude coefficients were retained and used for reconstructing the signal while the remaining coefficients were discarded. There was no modification of the amplitudes of the retained coefficients. The proposed representation was created with the 3/5 tap synthesis filters, and the coefficients were computed by solving the sparse linear system at each level.

The resulting PSNR's for each of the representations as a function of the number of retained

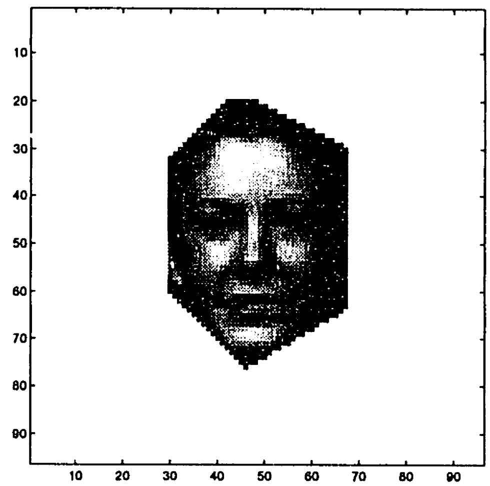

(a) Test region: Face

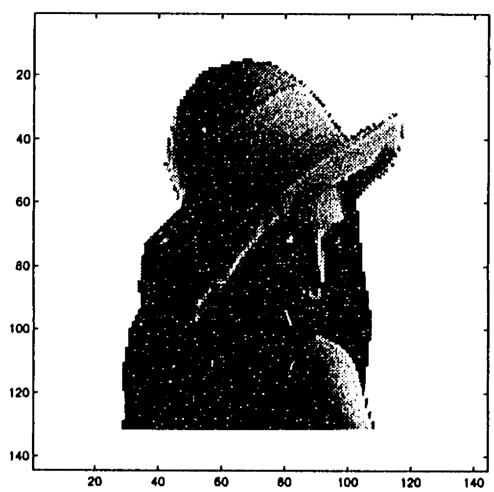

(b) Test region: Lenna

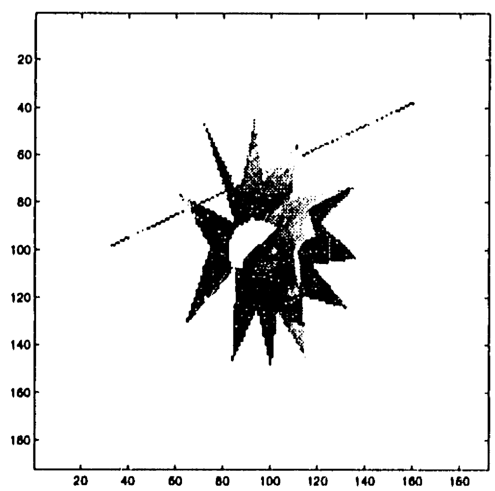

(c) Test region: Star

Figure 5.3: The three 2-D test signals with varying degrees of arbitrary supports.

coefficients is given in Figure 5.4. Note that the same number of coefficient are retained for each approach (critically sampled or overcomplete), however to simplify matters the number is expressed as a fraction of (with respect to) the number of coefficients in a critically sampled representation (equivalent to the number of samples in the ROS). Figure 5.5 illustrates the reconstructed Lenna signals that result for each approach, when retaining a number of coefficients equivalent to 20 % of the samples in the original signal.

A number of observations can be made from the performance plots. First, the simple 1-D Row/Col approach performs well for signals with appropriate ROS's, but performs poorly otherwise. Specifically, if each row and each column of the signal is greater than a minimum length determined by the particular filter, then the representation of that row or column is invertible. If the length is less than that minimum length then the representation is not invertible and there will also be a hit in compression performance. Notice that this minimum length must be satisfied at every level of the decomposition. The 1-D Row/Col approach achieves the highest performance for the Face signal and also for the Lenna signal for up to 8 % of the coefficients being retained. However, it performs badly for the Star signal with its complex support. In addition, it does not provide perfect reconstruction for any of the test signals. It is important to note that the filters used in these tests are extremely short and help to enhance the performance of the 1-D Row/Col approach with respect to the other approaches. If longer filters were applied, such as Daubechies biorthogonal 9/7 tap filters, then the 1-D Row/Col approach would provide significantly worse performance.

The overcomplete representations differ from each other only in the initial extrapolation performed: zero extrapolation and smooth extrapolation. A very interesting consequence is that sometimes the smooth extrapolation works better (Lenna and Star) and other times the zero extrapolation works better (Face). This points to the important issue, discussed in Section 2.4, that the optimal extrapolation will depend on the particular signal under consideration, and that a fixed extrapolation method will not perform well for all signals. This is a natural idea, however it is very surprising that the zero extrapolation would ever work better than a smooth extrapolation. This may be do to the fact that the segmented Face signal is not a completely homogeneous signal, rather the signal amplitude varies considerably near the signal boundaries. Therefore attempting to create a smooth extrapolation is not very successful. It is highly desirable for the representation to be robust – to perform well irrespective of the quality of the segmentation or the character of the given 2-D signal. Therefore, the dependence that these overcomplete representations have on the initial extrapolation is potentially troublesome.

Both the proposed approach and the overcomplete representations enable perfect reconstruction for a 2-D signal with any support. Since there are more potential coefficients to choose

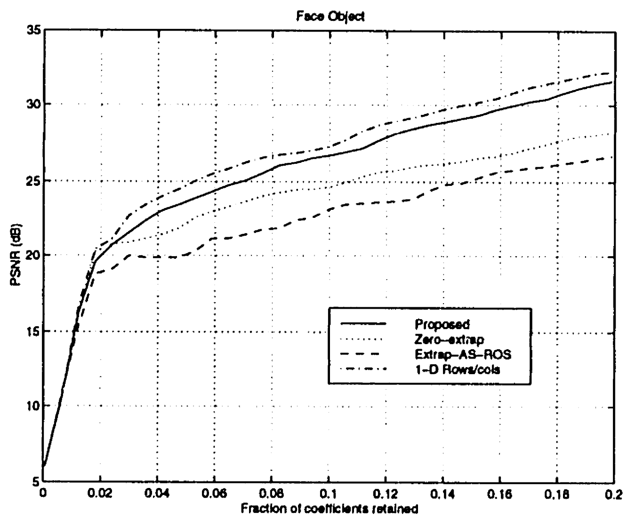

(a) Performance for Face Signal

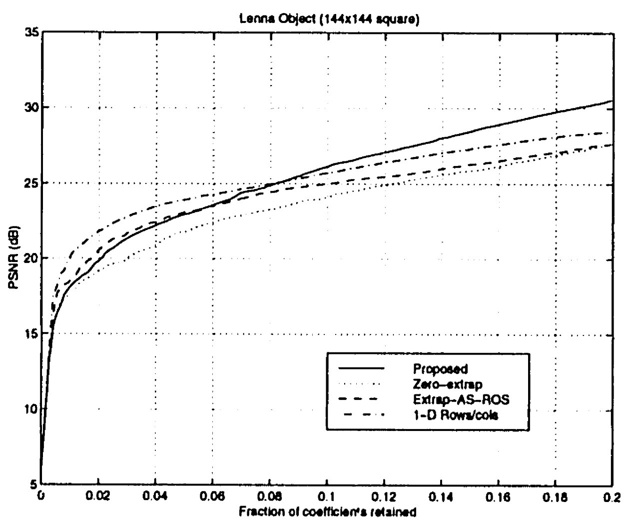

(b) Performance for Lenna Signal

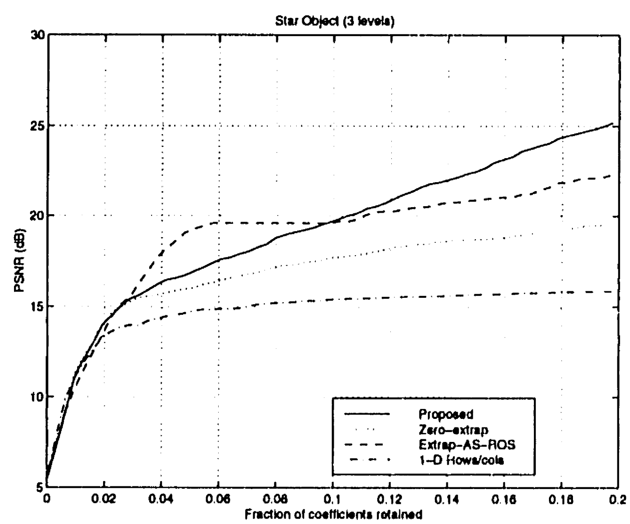

(c) Performance for Star Signal

Figure 5.4: PSNR (y-axis) vs. fraction of coefficients retained (x-axis) for the four approaches and the three test signals.

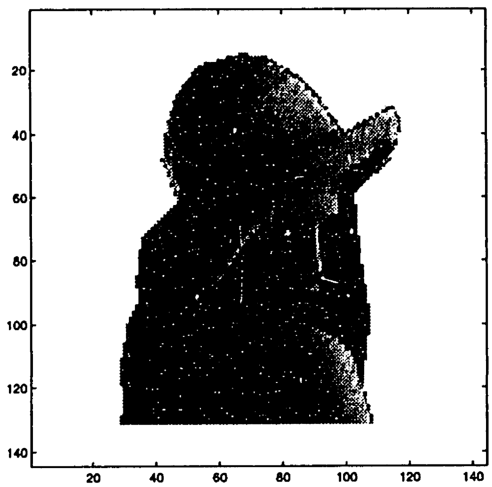

(a) Zero Extrapolation (27.72 dB)

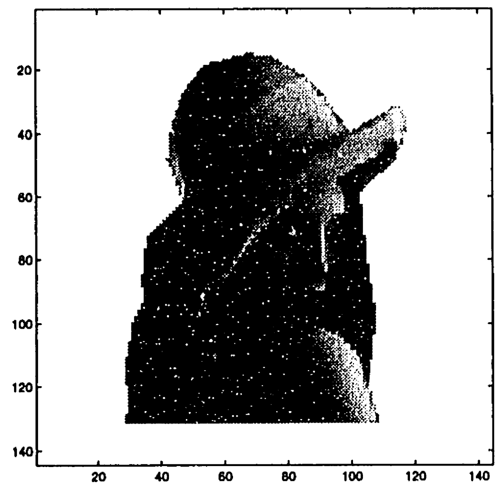

(b) Extrapolation for AS-ROS (27.68 dB)

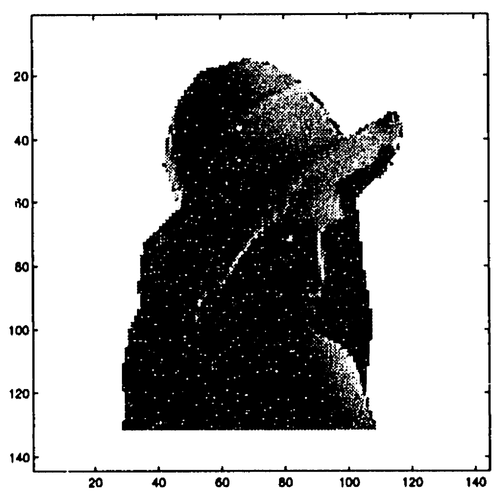

(c) 1-D Row/Col (28.52 dB)

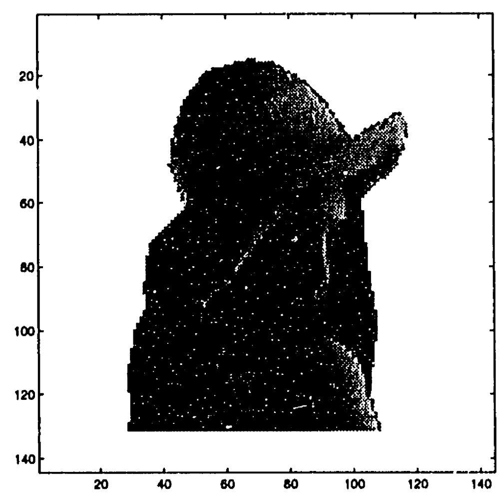

(d) Proposed Approach (30.55 dB)

Figure 5.5: Reconstructed signals for each representation while retaining 20 % of the coefficients with the largest amplitudes.

from in the overcomplete representations, more bits may be required to describe which coefficients have been retained, i.e. the overcomplete representations may require more *location information.* For this reason, the overcomplete representations may have a potential disadvantage in a full compression system. This issue with the location information is not reflected in these tests.

The proposed approach leads to a critically sampled representation and provides perfect reconstruction for a 2-D signal with any support. It also appears to provide good performance under a number of different circumstances. For example, it is applicable to signals with relatively simple supports and signals with very complex supports. It can be applied with short or long filters. The same algorithm is used in all cases – the algorithm does not depend on the signal's amplitude, the signal's support, or the choice of filters. Unlike the 1-D Row/Col approach, the proposed representation is directly applicable to any arbitrary support, and guarantees perfect reconstruction for every support. Unlike the overcomplete representations, no initial extrapolation is required and therefore there is no problem of determining the appropriate initial extrapolation for the given signal. Furthermore, since the proposed representation is critically sampled, it most likely requires fewer bits to code the locations of the non-zero coefficients than the overcomplete representations.

5.3 Summary

This chapter provides a preliminary investigation of the use of the ideas developed in this thesis for image/video compression. We identified some of the issues that must be addressed to develop a high-quality object/region-based compression system, and we examined the potential use of the proposed representation within the context of an object/region-based coder. This study was preliminary in that a full comparison of the proposed approach versus other possible approaches requires the development of an entire coder optimized for each representation. This would require an examination into the issues of quantization, location information, and entropy coding of all the parameters; this is beyond the scope of this thesis. Nevertheless, this chapter highlights some of the distinguishing features of the proposed representation in the context of object/region-based compression and demonstrates its possible application to compression.

In this chapter we have compared the proposed representation to a number of other possible representations. We have shown that the proposed representation provides a critically sampled perfect reconstruction representation for signals with any arbitrary support, and for typical objects/regions it concentrates most of the signal information into a small fraction of the coefficients which provide an accurate reconstruction of the original signal. In addition, the proposed representation avoids many of the undesirable characteristics of the other approaches.

Chapter 6

Concluding Remarks and Future Work

In this thesis we have developed critically sampled perfect reconstruction transform/subband representations for discrete 1-D, 2-D, and general M-D signals defined over arbitrary regions of support. These signals arise in a number of areas, including the arbitrarily shaped objects/regions within an image or video.

We proposed a novel general framework for creating critically sampled perfect reconstruction transform/subband representations for signals with arbitrary regions of support. This general approach determines a basis for a given signal by selecting an appropriate subset of vectors from a basis defined over a convenient superset space. This general approach may be based on any linear transform/subband representation (DFT, DCT, wavelet, etc.) and provides a significant amount of freedom in its creation. An important feature of this approach is that it provides a critically sampled, perfect reconstruction representation over an arbitrary support using tensor product (separable) basis vectors. This leads to a representation where the complexity is localized at either the analysis or the synthesis side, with the other side having a complexity comparable to that of a conventional separable transform. For example, in compression applications the reconstruction of the signal at the decoder can be made extremely simple.

We explored the creation of wavelet-type representations for signals with arbitrary supports within the context of the general framework presented in Chapter 3. In particular, we introduced two novel wavelet-type representations for signals with varying degrees of arbitrary support. The first representation was designed to fully preserve the wavelet's important polynomial accuracy property that we are accustomed to for infinite-length signals. The resulting representation provides some interesting theoretical features relative to previous approaches, however fully preserving polynomial accuracy leads to two disadvantages that may limit its practical usefulness.

The second wavelet representation was designed to represent signals with completely arbitrary supports. The main idea was to select the vector centered at each sample in the signal's support. By centering (aligning) the vectors within the representation, we related the problem of whether the representation provides a basis for all possible supports to the problem of whether all the principal submatrices of an initial matrix are nonsingular. This relationship provided considerable insight into the success of the representation, and aided in establishing a number of conditions on the admissibility of a given filter pair. Specifically, for 1-D, we have determined sufficient conditions on the filter taps such that the proposed representation will provide a basis for any arbitrary 1-D support. These conditions were satisfied by every wavelet filter tested (many of the popular wavelet filters in the literature). The 2-D/M-D problem is more complex, and some of the filters that are admissible in 1-D are not admissible in higher dimensions, e.g. Haar filters. We have shown that a small group of filters will provide a basis for any arbitrary 2-D support, and considerable empirical evidence and theoretical considerations suggest that this holds true for a much larger set of filters.

This representation provides a number of compelling properties. Foremost, for appropriate filters, it provides a basis over any possible arbitrary support, i.e. the ROS can have any shape, it can contain holes, it can be disconnected, and it can include isolated samples. While it may not be beneficial to apply this representation for all possible arbitrary supports, the fact that there are no limitations can greatly simplify a number of applications. In addition, with our proposed method of aligning the filters, the selection of the vectors is simply related to the signal's support, thereby providing an extremely simple approach for determining a basis given a signal's support. Furthermore, for typical signals, the resulting basis preserves many of the general characteristics provided by wavelet representations of signals with convenient supports. For example, the representation provides a natural multiresolution representation of the signal and its support, good stability to perturbations in the signal's amplitude and its support, a potentially large class of filters to choose from, and a high concentration of information into a small fraction of the coefficients for typical signals. Overall, this representation appears very promising as a general purpose wavelet-type representation for signals with arbitrary ROS's, and appears particularly promising for representing the objects/regions within an image or video.

We have conducted a preliminary analysis of the application of the proposed representation to object/region-based image/video compression. Specifically, we examined the problem of representing the amplitudes of the objects/regions in an image or video scene using our second wavelet representation, and made a preliminary comparison with other possible approaches. The proposed representation provides a critically sampled perfect reconstruction representation for signals with any arbitrary support, and for typical objects/regions it concentrates most of the signal information into a small fraction of the coefficients which can be used to provide an accurate reconstruction of the original signal. In addition, the proposed representation avoids many of the undesirable characteristics of the other approaches. While considerable work remains to properly evaluate this representation for object/region-based coding (e.g. incorporation within a complete coder), current results indicate that this is a very promising approach that warrants further investigation.

6.1 Directions for Future Research

There are a number of potentially promising directions for future research. We briefly consider three potential areas: extending the wavelet-based representations developed within this thesis, examining other initial transform/subband representations used within the general approach, and investigating potential applications.

Extending the Wavelet-based Representations The most direct extension of this thesis is to develop an improved understanding and characterization of the two wavelet representations introduced in this thesis. As already mentioned, it would be very interesting to characterize the class of filters for which representation #2 provides a basis for any arbitrary support, and to develop a simple test on the filter taps that would determine the admissibility of a filter for 2-D and general M-D. Identifying the fundamental properties of a filter that lead to its admissibility can enable the design of improved filters for different applications that involve processing of signals with arbitrary supports.

In this thesis we primarily focused on two representations in which we selected the vectors to create synthesis bases. These representations provided a number of benefits including localizing the complexity at the analysis (encoder) side, thus leading to a very simple synthesis (decoder) side. Another important question relates to examining the creation of an analysis basis: what are the possible properties and what is the usefulness of selecting the vectors to create an analysis basis? These questions may be pertinent for some sampling problems. Furthermore, we should reiterate that there may be a large number of potentially-rewarding wavelet representations within this framework that remain to be identified and explored.

A more pressing practical issue is the investigation of fast methods for computing the representation for a given signal. We have speculated on various strategies that may provide a low complexity solution depending on the particular application. A number of factors influence the choice of computational method including the structure of the signal's support, the properties of the filters, the required accuracy, and the memory and computational constraints. A careful study of these issues, both theoretically and by extensive simulations, may have a significant impact on the practical usefulness of the proposed representations.

Examining Other Initial Transform/Subband Representations Another potentially promising area of future research is to expand this thesis in terms of examining other initial representations within the context of our general approach. Specifically, the idea of developing a global DFT-or DCT-based representation appears especially intriguing. In this case, one would begin with a DFT or DCT basis defined over a superset space that contains the given signal. The Fourier representation appears very interesting because Fourier bases arise in a large number of contexts. One may potentially approach the problem as a sampling problem in the Fourier domain, where considerable research already exists, e.g. [86]. This may also provide a useful representation for the dual problem of a signal with an arbitrary support in the frequency domain. Problems of this form arise in a number of areas such as array processing and beam forming. A DCT basis defined over the entire support may potentially be useful for compressing large objects/regions in an image that are highly homogeneous.

A number of issues must be overcome for the practical success of these approaches, including the potentially large complexity and the stability of the representation. Specifically, it is somewhat doubtful that the selection algorithm for determining an appropriate basis for a given support will be as simple for the DFT or DCT as it is for the wavelet. Also, a small perturbation in the signal's support may potentially lead to a relatively larger change in the DFT-based basis than in the wavelet-based basis. These are just two of the many issues that should be investigated for a DFT- or DCT-based representation.

A problem motivated by compression applications is to develop an appropriate representation for the 8×8 DCT. This is more practically feasible than a global DCT taken over the entire signal's support. Many conventional object/region-based compression algorithms apply an 8×8-pixel Block-DCT transform to each object/region. The boundary blocks that are partially inside and partially outside the signal's support are processed by extrapolating the signal to fill the 8×8 pixel block and then computing a conventional DCT of the extrapolated block – leading to an over-complete representation. Developing a critically sampled perfect reconstruction representation for the boundary blocks may lead to a more compression-efficient representation.

Potential Applications Perhaps the most interesting area of future research is to examine what problems or applications can be successfully addressed using the ideas developed within this thesis.

Object/region-based image/video processing and compression are two of the more obvious possible applications. As discussed in Chapter 5, the second wavelet representation appears very promising for representing the objects/regions in an image or video scene. Important questions that must be examined for compression include determining an appropriate coding method for a biorthogonal representation and an appropriate representation of the MC-residual of the objects/regions. Extensive analysis and simulations are required to develop an appropriate coder using this proposed representation and to compare it to other possible approaches.

The ideas and representations developed within this thesis appear to facilitate general object/region-based image/video processing, and so-called content-based processing. Representations of the form presented in this thesis may possibly be used to analyze, enhance, restore, and compress these signals. However, this is a very new and unexplored area, with many questions to be answered as well as identified. The field of medical imaging is another area where we may want to perform similar types of processing on signals with arbitrary supports. A careful study of the relevant problems in medical imaging and the conventional solutions may identify potential application areas for this work.

The numerical solution of partial differential equations typically have extremely stringent performance requirements, e.g. they require polynomial accuracy, good conditioning, and applicability to arbitrary supports. Currently, it is an open question whether wavelet-based methods will provide improved performance over the traditional approaches such as finite elements. The two wavelet representations presented in this thesis do not appear well suited for these applications. However, it may be possible to create other representations with this general approach that provide the necessary properties, perhaps increasing the applicability of wavelet-based methods for certain classes of problems.

Clearly, most of the above ideas are rather speculative, and further research is needed to verify them both theoretically and experimentally. Perhaps one of the most convincing reasons for examining these possible applications is that by determining desirable properties for the representation one may identify some of the most promising directions for expanding/extending this thesis. Overall, these representations for signals with arbitrary supports appear promising for a number of areas and applications and deserve further study.

Appendix A

Proof of Theorem 4.2.1

Theorem 4.2.1 *Given a 2-band filterbank composed of Daubechies L-tap orthogonal filters, and an arbitrary-length 1-D signal of minimum length L-1, selecting all the interior vectors and the lowpass boundary vectors will provide a basis for the signal. Furthermore, this basis will preserve the polynomial accuracy property, and for a polynomial signal of order less than $\frac{L}{2}$ the LP subband coefficients will also be a polynomial of the same order.*

Proof: An overview of the proof for this theorem was given in Section 4.2.1, here we fill in some details. The vectors that overlap the signal's support can be split into two groups: (1) the interior vectors which lie entirely inside the signal's ROS and (2) the boundary vectors which lie partially inside and partially outside the ROS. Let V_{sig} refer to the signal space, V_{int} to the subspace spanned by the interior vectors, and $V_{\text{bnd-trunc}}$ to the subspace spanned by the boundary vectors when truncated to the signal's ROS. The subspaces are related by:

$$V_{\text{sig}} = V_{\text{int}} \oplus V_{\text{bnd-trunc}} .$$

Determining a basis for V_{sig} is equivalent to determining a basis for both V_{int} and $V_{\text{bnd-trunc}}$. Note that V_{int} and $V_{\text{bnd-trunc}}$ are orthogonal. All the interior vectors must be selected to provide a basis for V_{int}. The problem that remains is to determine a basis for $V_{\text{bnd-trunc}}$. The important question are: What is the dimension of $V_{\text{bnd-trunc}}$? And do the selected lowpass boundary vectors provide a basis for $V_{\text{bnd-trunc}}$?

Let us begin by restricting our attention to a causal (one-sided) signal in order to consider only a single boundary. There are two possible phases of the signal boundary relative to the translated basis vectors: (1) the basis vectors line up with the signal boundary (even phase) and (2) the basis vectors are offset by one sample from the signal boundary (odd phase). These two phases lead to two different boundary situations to analyze. Consider the even phase first.

In the case of Daubechies L-tap filters and even phase positioning, the boundary vectors overlap $L-2$ samples of the signal. Each pair of lowpass and highpass filters are aligned, and have signal overlaps of $2, 4, 6, ...,$ and $L-2$ samples. There are a total of $\frac{L-2}{2}$ pairs of filters for a total of $L-2$ boundary vectors: $\frac{L-2}{2}$ lowpass and $\frac{L-2}{2}$ highpass. In addition, it is shown in

[30] that $V_{\text{bnd-trunc}}$ has dimension $\frac{L-2}{2}$. Therefore, at each boundary half of the total number of boundary vectors should be selected.

We now show that retaining only the lowpass boundary vectors is sufficient to form a basis for $V_{\text{bnd-trunc}}$. The lowpass boundary vectors lie in the space of the boundary. The boundary vectors are linearly independent if and only if their projections onto $V_{\text{bnd-trunc}}$ are linearly independent. The projection of the boundary vector onto $V_{\text{bnd-trunc}}$ is equivalent to the portion of that boundary vector within the signal's support. The lowpass boundary vectors are offset from each other, that is they have *staggered* supports. Since the lowpass boundary vectors have staggered supports they are linearly independent of each other. As previously discussed, there are a total of $\frac{L-2}{2}$ lowpass boundary vectors. The lowpass boundary vectors correspond to $\frac{L-2}{2}$ linearly independent vectors in a space of dimension $\frac{L-2}{2}$, therefore they provide a basis for $V_{\text{bnd-trunc}}$. Thus, the interior basis vectors (both highpass and lowpass) and the boundary lowpass vectors provide a basis for the space of one-sided signals where the boundary has even phase relative to the vectors.

The case of odd phase of the signal relative to the vectors follows straightforwardly. The odd phase signal space can be viewed as an even phase signal space increased by one sample. Note that the set of interior vectors and V_{int} are the same for both phases (since the vectors are translated by two samples, the signal's ROS must change by at least two samples to change V_{int}). The dimension of $V_{\text{bnd-trunc}}$ increases by one, to $\frac{L}{2}$. However, there is one extra pair of lowpass/highpass vectors that interacts with the boundary then in the even phase case. Hence there are $\frac{L}{2}$ lowpass boundary vectors. Once again the lowpass boundary vectors have staggered supports and are therefore linearly independent, providing a basis for $V_{\text{bnd-trunc}}$ in the case of odd phase.

We have shown that retaining all the interior vectors and the lowpass boundary vectors provides a basis for the class of one-sided (causal) 1-D signals of even or odd phase with respect to the basis vectors. The case of an anti-causal 1-D signals follows directly: the analysis is exactly the same, the vectors can simply be considered time-reversed versions of the causal case.

The analysis also holds for any arbitrary-length signal (any even or odd length) as long as the signal is of a minimum length. The minimum length arises from two issues. First, no vector should correspond to a boundary vector for both ends simultaneously. Otherwise that vector will have its support clipped at both ends, and the set of lowpass boundary vectors may no longer have staggered supports, and our simple proof for linear independence no longer holds. This problem can be avoided by restricting the signal length to be $\geq (L-2)$. The more stringent issue relates to the goal of selecting all the lowpass vectors that overlap the signal's ROS in order to preserve the polynomial accuracy property. The set of selected lowpass vectors must be linearly independent in order to have any chance of forming a basis. However, if the signal is too short, there may be more lowpass vectors that overlap the signal then there are samples in the signal. For example, given a 2-point signal and 4-tap filters (Daub4 filters), in the case where the signal's ROS has odd phase relative to the filters, then 3 lowpass vectors will overlap the signal. One can show that for an N sample signal and L-tap filters, a maximum of $\lceil \frac{N+L-1}{2} \rceil$ lowpass vectors overlap with the signal. By restricting the signal length to be $\geq (L-1)$, the number of overlapping lowpass vectors will always be less than the number of samples in the signal. Combining the two constraints, the proposed representation is valid for any signal of length $\geq (L-1)$ ■

Bibliography

[1] J. S. Lim, *Two-Dimensional Signal and Image Processing*. Englewood Cliffs, N.J.: Prentice Hall, Inc., 1990.

[2] G. Strang and T. Nguyen, *Wavelets and Filter Banks*. Wellesley, MA: Wellesley-Cambridge Press, 1996.

[3] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*. Englewood Cliffs, N.J.: Prentice Hall, Inc., 1995.

[4] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*. Englewood Cliffs, N.J.: Prentice Hall, 1993.

[5] M. Kunt, A. Ikonomopoulos, and M. Kocher, "Second generation image coding techniques," *Proceedings of the IEEE*, vol. 73, pp. 549–574, Apr. 1985.

[6] H. Musmann, H. Hötter, and J. Ostermann, "Object-oriented analysis-synthesis coding of moving images," *Signal Processing: Image Communications*, vol. 1, pp. 117–138, October 1989.

[7] J. Konrad, A.-R. Mansouri, and E. Dubois, "Towards region-based coding for very low bit rate video services," Tech. Rep. 95-13, INRS-Télécommunications, April 1995.

[8] G. Strang, *Introduction to applied mathematics*. Wellesley, Mass.: Wellesley-Cambridge Press, 1986.

[9] W. Dahmen, A. Kunoth, and R. Schneider, "Operator equations, multiscale concepts and complexity." Preprint, 1996.

[10] A. Cohen, W. Dahmen, and R. DeVore, "Adaptive-tree-structured multiscale approximations & PDE's." MIT Applied Math Colloquium, October 21, 1996.

[11] W. Briggs, *A Multigrid Tutorial*. Philadelphia, PA: SIAM, 1987.

[12] I. Daubechies, *Ten Lectures on Wavelets*. CBMS-NSF Regional Conference Series in Applied Mathematics, SIAM, 1992.

[13] S. Mallat, "From Fourier to wavelet signal processing." Course notes for MIT Course 6.971, Fall 1994. To appear as textbook: Wavelet Signal Processing.

[14] M. Vetterli and C. Herley, "Wavelets and filter banks: Theory and design," *IEEE Transactions on Signal Processing*, vol. 40, pp. 2207–2232, September 1992.

[15] G. Strang, "Wavelets and dilation equations: A brief introduction," *SIAM Rev*, vol. 31, pp. 614–627, December 1989.

[16] S. Mallat, "Multifrequency channel decompositions of images and wavelet models," *IEEE Transactions Acoustics, Speech, and Signal Processing*, vol. 37, pp. 2091–2110, December 1989.

[17] M. Smith and S. Eddins, "Analysis/synthesis techniques for subband image coding," *IEEE Transactions on Signal Processing*, vol. 38, pp. 1446–1456, August 1990.

[18] G. Karlsson and M. Vetterli, "Extension of finite length signals for subband coding," *Signal Processing*, vol. 17, no. 2, pp. 161–168, 1989.

[19] R. Bamberger, S. Eddins, and V. Nuri, "Generalized symmetric extension for size-limited multirate filter banks," *IEEE Transactions on Image Processing*, vol. 3, pp. 82–87, Jan. 1994.

[20] H. Kiya, K. Nishikawa, and M. Iwahashi, "A development of symmetric extension method for subband image coding," *IEEE Transactions on Image Processing*, vol. 3, pp. 78–81, Jan. 1994.

[21] C. Brislawn, "Preservation of subband symmetry in multirate signal coding," *IEEE Transactions on Signal Processing*, vol. 43, pp. 3046–3050, Dec. 1995.

[22] R. de Queiroz and K. Rao, "On reconstruction methods for processing finite-length signals with paraunitary filter banks," *IEEE Transactions on Signal Processing*, vol. 43, pp. 2407–2410, October 1995.

[23] V. Nuri and R. Bamberger, "Size-limited filter banks for subband image compression," *IEEE Transactions on Image Processing*, vol. 4, pp. 1317–1323, September 1995.

[24] C. Brislawn, "Classification of nonexpansive symmetric extension transforms for multirate filter banks," *Preprint*, pp. 1–36, March 5 1996.

[25] J. Williams and K. Amaratunga, "High order wavelet extrapolation schemes for initial value problems and boundary value problems," *MIT IESL Technical Report*, July 25 1994.

[26] J. Williams and K. Amaratunga, "A discrete wavelet transform without edge effects using wavelet extrapolation," *MIT IESL Technical Report*, January 28 1995.

[27] B. Jawerth and W. Sweldens, "An overview of wavelet based multiresolution analysis," *SIAM Review*, vol. 36, no. 3, pp. 377–412, 1994.

[28] Y. Meyer, "Ondelettes sur l'intervalle," *Rev. Mat. Iberoamericana*, vol. 7, pp. 115–133, 1992.

[29] A. Cohen, I. Daubechies, and P. Vial, "Wavelets on the interval and fast wavelet transforms," *J. Applied Computational and Harmonic Analysis*, vol. 1, pp. 54–81, Dec. 1993.

[30] C. Herley and M. Vetterli, "Orthogonal time-varying filter banks and wavelet packets," *IEEE Transactions on Signal Processing*, vol. 42, pp. 2650–2663, October 1994.

[31] C. Herley, "Boundary filters for finite-length signals and time-varying filter banks," *IEEE Transactions on Circuits and Systems-II: Analog and Digital Signal Processing*, vol. 42, pp. 102–114, February 1995.

[32] T. Sikora and B. Makai, "Shape-adaptive DCT for generic coding of video," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 5, pp. 59–62, February 1995.

[33] T. Sikora, S. Bauer, and B. Makai, "Efficiency of shape-adaptive 2-D transforms for coding of arbitrarily shaped image segments," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 5, pp. 254–258, June 1995.

[34] H. Barnard, J. Weber, and J. Biemond, "A region-based discrete wavelet transform for image coding," in *Wavelets In Image Communication* (M. Barlaud, ed.), ch. 4, Amsterdam, The Netherlands: Elsevier, 1994.

[35] C. Stiller and J. Konrad, "Eigentransforms for region-based image processing," in *Proc. IEEE International Conf. on Consumer Electronics*, pp. 286–287, June 1995.

[36] C. Stiller and J. Konrad, "Region-adaptive transform based on a stochastic model," in *Proceedings of the IEEE International Conference on Image Processing*, vol. 2, pp. 264–267, October 1995.

[37] M. Gilge, T. Engelhardt, and R. Mehlan, "Coding of arbitrarily shaped image segments based on a generalized orthogonal transform," *Signal Processing: Image Communication*, vol. 1, pp. 153–180, 1989.

[38] A. Cohen, W. Dahmen, and R. DeVore, "Multiscale decompositions on bounded domains," *Preprint*, Feburary 12 1996.

[39] W. Sweldens, "The lifting scheme: A construction of second generation wavelets," *SIAM J. Math. Anal.*, To appear.

[40] W. Sweldens, "Wavelets and the lifting scheme: A 5 minute tour," *Z. Angew. Math. Mech.*, vol. 76 (Suppl. 2), pp. 41–44, 1996.

[41] W. Sweldens and P. Schröder, "Building your own wavelets at home," in *Wavelets in Computer Graphics*, ACM SIGGRAPH Course Notes, 1996.

[42] I. Daubechies and W. Sweldens, "Factoring wavelet transforms into lifting steps," tech. rep., Bell Laboratories, Lucent Technologies, 1996.

[43] P. Schröder and W. Sweldens, "Spherical wavelets: Texture processing," in *Rendering Techniques '95* (P. Hanrahan and W. Purgathofer, eds.), Wien, New York: Springer Verlag, August 1995.

[44] P. Schröder and W. Sweldens, "Spherical wavelets: Efficiently representing functions on the sphere," *Computer Graphics Proceedings (SIGGRAPH 95)*, pp. 161–172, 1995.

[45] J. Apostolopoulos and J. Lim, "Representing arbitrarily-shaped regions: A case study of overcomplete representations," in *Proceedings of the IEEE International Conference on Image Processing*, vol. 1, pp. 426–429, October 1995.

[46] G. Davis, S. Mallat, and M. Avellaneda, "Adaptive nonlinear approximations," tech. rep., New York University, Courant Institute, 1994.

[47] J. Apostolopoulos and J. Lim, "Coding of arbitrarily-shaped regions," *Visual Communications and Image Processing*, pp. 1713–1726, May 1995.

[48] H. Chen, M. Civanlar, and B. Haskell, "A block transform coder for arbitrarily shaped image segments," *Proceedings of the IEEE International Conference on Image Processing*, vol. 1, pp. 85–89, Nov. 1994.

[49] D. Youla and H. Webb, "Image restoration by the method of convex projections: Part 1-theory," *IEEE Transactions on Medical Imaging*, vol. MI-1, pp. 81–94, Oct. 1982.

[50] M. Sezan, "An overview of convex projections theory and its application to image recovery problems," *Ultramicroscopy*, vol. 40, pp. 55–67, 1992.

[51] A. Levi and H. Stark, "Image restoration by the method of generalized projections with application to restoration from magnitude," *J. Opt. Soc. Am.*, vol. 1, pp. 932–, Sept. 1984.

[52] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, pp. 3397–3415, Dec. 1993.

[53] G. Davis, S. Mallat, and Z. Zhang, "Adaptive time-frequency decompositions," *Optical Engineering*, vol. 33, pp. 2183–2191, July 1994.

[54] J. Högbom, "Aperture synthesis with a non-regular distribution of interferometer baselines," *Astron. Astrophys. Suppl.*, vol. 15, pp. 417–426, 1974.

[55] R. Bracewell, *Two-Dimensional Imaging*. Englewood Cliffs, NJ: Prentice Hall, 1995.

[56] P. Jansson, *Deconvolution: With Applications to Spectroscopy*. Orlando, FL: Academic Press, Inc., 1984.

[57] J. Friedman and W. Stuetzle, "Projection pursuit regression," *Journal American Statistical Association*, vol. 76, pp. 817–823, 1981.

[58] P. Huber, "Projection pursuit," *The Annals of Statistics*, vol. 13, no. 2, pp. 435–475, 1985.

[59] L. Jones, "Constructive approximations for neural networks by sigmoidal functions," *Proceedings of the IEEE*, vol. 78, pp. 1586–1589, October 1990.

[60] A. Kaup and T. Aach, "A new approach towards description of arbitrarily shaped image segments," *IEEE Inter. Workshop on Intell. Sig. Proc. Com. Sys.*, pp. 543–553, Mar. 1992.

[61] U. Desai, "Coding of segmented image sequences," Master's thesis, MIT, 1994.

[62] S.-F. Chang and D. Messerschmitt, "Comparison of transform coding techniques for two-dimensional arbitrarily shaped images," *ACM/Springer-Verlag Multimedia Systems*, pp. 231–239, 1994.

[63] V. Goyal, M. Vetterli, and N. Thao, "Quantization of overcomplete expansions," *Proceedings of Data Compression Conference*, pp. 13–22 1995.

[64] S. Mallat, U. Desai, and J. Apostolopoulos, November 1994. Personal communications.

[65] U. Desai, "DCT and wavelet based representations of arbitrarily shaped image segments," in *Proceedings of the IEEE International Conference on Image Processing*, vol. 1, pp. 558–561, October 1995.

[66] S. Chen and D. Donoho, "Basis pursuit," *Asilomar Conference on Signals, Systems, and Computers*, Nov. 1994.

[67] S. Chen and D. Donoho, "Atomic decomposition by basis pursuit," tech. rep., Statistics Dept., Stanford University, May 1995.

[68] G. Strang, *Linear Algebra and Its Applications*. Fort Worth, TX: Saunders College Publishing, 3rd ed., 1988.

[69] B. Jacob, *Linear Algebra*. New York, NY: W.H. Freeman and Company, 1990.

[70] M. Smith and T. Barnwell, "Exact reconstruction techniques for tree-structured subband coders," *IEEE Transactions on Signal Processing*, vol. 34, pp. 434–441, June 1986.

[71] E. Adelson, E. Simoncelli, and R. Hingorani, "Orthogonal pyramid transforms for image coding," *SPIE Vol. 845 Visual Communications and Image Processing II*, 1987.

[72] D. Le Gall and A. Tabatabai, "Sub-band coding of digital images using symmetric short kernel filters and arithmetic coding techniques," *International Conference on Acoustics, Speech, and Signal Processing*, vol. M2.3, 1988.

[73] J. Apostolopoulos and J. Lim, "Video compression for digital advanced television systems," in *Motion Analysis and Image Sequence Processing* (M. Sezan and R. Lagendijk, eds.), ch. 15, Kluwer Academic Publishers, 1993.

[74] S. Cheung, *Biorthogonality in Lapped Transforms: A Study in High-Quality Audio Compression*. PhD thesis, Massachusetts Institute of Technology, 1996.

[75] G. Golub and C. Van Loan, *Matrix Computations*. Baltimore, Maryland: Johns Hopkins Press, 3rd ed., 1997.

[76] Y. Saad, *Iterative Methods for Sparse Linear Systems*. Boston, MA: PWS Publishing Company, 1996.

[77] R. Barrett, M. Berry, T. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. Van der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*. 1996.

[78] N. Jayant, J. Johnston, and B. Safranek, "Signal compression based on models of human perception," *Proceedings of the IEEE*, vol. 81, pp. 1385–1422, October 1993.

[79] D. J. LeGall, "MPEG: A video compression standard for multimedia applications," *Communications of the ACM*, vol. 34, pp. 47–58, April 1991.

[80] *IEEE Transactions on Circuits and Systems for Video Technology: Special Issue on MPEG-4*, vol. 7, February 1997.

[81] K. Ramchandran and M. Vetterli, "Best wavelet packet bases in a rate-distortion sense," *IEEE Transactions on Image Processing*, vol. 2, pp. 160–175, April 1993.

[82] P. Moulin, "A multiscale relaxation algorithm for SNR maximization in nonorthogonal subband coding," *IEEE Transactions on Image Processing*, vol. 4, pp. 1269–1281, September 1995.

[83] R. Neff and A. Zakhor, "Very low bit rate video coding based on matching pursuits," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 7, pp. 158–171, February 1997.

[84] M. Banham and J. Brailean, "A selective update approach to matching pursuits video coding," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 7, pp. 119–129, February 1997.

[85] J.-H. Moon, G. Park, S.-M. Chun, and S.-R. Choi, "Shape-adaptive region partitioning method for shape-assisted block-based texture coding," *IEEE Transactions on Circuits and Systems for Video Technology*, 1997.

[86] A. Zakhor, *Reconstruction of Multidimensional Signals from Multiple Level Threshold Crossings*. PhD thesis, Massachusetts Institute of Technology, 1987.