TRANSFORM/SUBBAND REPRESENTATIONS FOR SIGNALS WITH ARBITRARILY SHAPED REGIONS OF SUPPORT

John G. Apostolopoulos and Jae S. Lim

Research Laboratory of Electronics
Massachusetts Institute of Technology
Cambridge, MA USA
japos@image.mit.edu

ABSTRACT

Transform/subband representations form a basic building block for many signal processing algorithms and applications. Most of the effort 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 signals may have arbitrary length or arbitrarily shaped (AS) regions of support (ROS). We present a novel framework for creating critically sampled perfect reconstruction transform/subband representations for AS signals. Our method selects an appropriate subset of vectors from an (easily obtained) basis for a larger (superset) signal space, in order to form a basis for the AS signal. In particular, we have developed a number of promising wavelet representations for arbitrary-length 1-D signals and AS 2-D/M-D signals that provide high performance with low complexity.

1. INTRODUCTION

Many applications today and in the near future will entail efficient processing of multi-dimensional (M-D) signals with arbitrarily shaped (non-rectangular) regions of support. 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 problems defined over arbitrarily shaped domains.

In this paper we propose a novel approach for creating efficient transform/subband representations for discrete 1-D, 2-D, and general M-D signals defined over arbitrarily shaped regions of support. For brevity, we will refer to these signals as AS signals. Specifically, we assume that the ROS of the signal is given, and the goal is to represent the signal's amplitude over the ROS. We desire a representation that can achieve the same high performance with low complexity for AS signals as is currently achieved for signals with convenient supports. Note that complexity is of prime importance, since even a small M-D signal can make most algorithms impractical, e.g. a 3-D AS signal of only 30-sample "diameter" may require the processing of up to 27,000 samples.

This research was supported by the MIT Advanced Television Research Program and by an AT&T Bell Laboratories fellowship.

This paper begins with a brief overview of the previous research in representing AS signals, and then presents our general approach. We then focus on creating a wavelet-type representation, which appears to be a natural approach for many AS problems. We present two different wavelet representations and discuss the relative merits of each.

2. PREVIOUS RESEARCH

Considerable research has focused on representing finite-length 1-D signals and a number of highly successful and practical 1-D methods exist. However these methods do not appear applicable for representing M-D AS signals. For example, the symmetric extension method provides a very simple and successful approach for producing a critically sampled (CS) perfect reconstruction (PR) representation for 1-D signals using linear phase filters. However, determining an extrapolation for an M-D AS signal that enables CS and PR appears to be a much more complex problem.

Boundary filter construction is an elegant approach for creating a CS PR representation with nice boundary filters for processing 1-D signals [1, 2]. An important practical feature for 1-D is that since there are a very small number of boundary cases, the complex boundary filter construction can be performed once for a given filter set and stored for repeated application. However, 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 cases. Overall, boundary filter construction is complex and cumbersome (indexing, etc.), thereby severely limiting its application to M-D problems. Similar difficulties also exist for [3].

A CS PR representation can easily be constructed for an M-D AS signal by applying 1-D schemes along each of the dimensions, e.g. a 2-D signal can be processed by applying an invertible transform along all the rows and subsequently along all the columns [4, 5]. These approaches require very low complexity, however they possess certain undesirable properties. For example, while they process the rows and columns separately, the transform is not separable, i.e. the result depends on whether the rows or columns are processed first.

A natural approach for representing M-D AS signals is to use a transform/subband scheme defined over a circumscribing square [6, 7]. This guarantees PR and facilitates the use of the separable/fast properties of the underlying superset basis. This is a promising approach, however its overcomplete nature makes processing more complex.

3. OVERVIEW OF GENERAL APPROACH

The goal of this work is to construct a CS PR representation for signals with arbitrarily shaped regions of support. This is equivalent to determining a basis for the signal space defined by the ROS of the signal. Our approach is 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 employ our knowledge and ability in representing signals of convenient support, in order to produce a practical and successful method for representing signals of arbitrarily shaped supports. To illustrate this approach, consider a 2-D AS signal and an initial transform/subband representation defined over a support that contains the given signal. 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 that they are defined. In addition to spanning their conveniently defined space, they also span the signal space given by the AS signal. Therefore, we would like to *select an appropriate subset of these vectors which will provide a basis for the AS signal.* The general approach can therefore be summarized as follows:

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

Beginning with a basis over a circumscribing space guarantees that there exists (at least one) basis for the desired signal space — typically there are many subsets of vectors that lead to a basis. In addition, the chosen basis for the AS signal can inherit some of the important approximation and computational properties of the basis over the larger space.

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, as well as placement of the signal within the ROS; subset of vectors (which basis) to choose for the AS signal. Furthermore, the selected vectors may be used for "analyzing" or "synthesizing" the signal. In effect, we are creating a biorthogonal representation and we are explicitly choosing *either* the analysis or the synthesis.

The desired representation (basis) for the AS signal should be chosen so that it provides high performance and low complexity. The specifics of high performance will vary with the individual application. For example, 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, and the coefficient computation being relatively fast.

4. WAVELET-TYPE REPRESENTATIONS

A wavelet-type representation appears to be a natural choice for AS signals. Wavelet representations achieve high performance in representing images, and therefore appear promising for representing the arbitrarily shaped objects or regions in an image. Wavelets with local support 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. In addition, the ROS of the initial wavelet representation does not matter — 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 (unlike a global DCT). These properties greatly simplify the analysis and computational issues.

The remainder of this paper will focus on 1-D and 2-D AS signals, Daubechies orthogonal FIR (closest to linear phase) filters and their use as *synthesis vectors.* The general case can be found in [8]. Consider the problem of determining a single-level wavelet decomposition of an AS 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 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 $V_{\text{AS-sig}}$ be the signal space and $V_{\text{bnd-trunc}}$ be the subspace spanned by the boundary vectors when truncated to the AS ROS:

$$\begin{aligned} V_{\text{2D-plane}} &= V_{\text{int}} \oplus V_{\text{bnd}} \oplus V_{\text{ext}} \\ V_{\text{AS-sig}} &= V_{\text{int}} \oplus V_{\text{bnd-trunc}} \,. \end{aligned}$$

Determining a basis for $V_{\text{AS-sig}}$ is equivalent to determining a basis for each V_{int} and $V_{\text{bnd-trunc}}$ — thus 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, $V_{\text{int}} \perp V_{\text{bnd-trunc}}$, however the 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 [1, 2] where a basis for $V_{\text{bnd-trunc}}$ is explicitly constructed, we simply select or discard each boundary vector.

4.1. Selecting the Vectors/Choosing a Basis

Given a 2-D AS signal and an initial wavelet representation defined over the entire 2-D plane, we select a subset of the vectors from the initial representation that provide a basis for the AS signal. There are typically many possible subsets of vectors that lead to a basis, where some choices are better in a certain sense than others. We will examine two possible selection schemes. The first is primarily of theoretical interest, while the second may be more successful in practical applications. Because of the limited space, we will focus on the basic concepts and limitations of each approach, and present some brief examples of their application in Section 5. Further details and analysis may be found in [8].

Polynomial Accuracy Property A natural criterion for the selection is to preserve the important wavelet property of representing polynomials. An L-tap Daubechies HP filter has $p = \frac{L}{2}$ vanishing moments. Any polynomial of order less than p is exactly represented by the lowpass subband, i.e. all the energy (information) is contained in the lowpass subband, and all the bandpass and highpass subbands are zero. In addition, the lowpass subband will also be a polynomial of similar order. The LP vectors exactly reproduce polynomial signals. To retain this desirable property we retain all the LP vectors that overlap the signal's ROS. The resulting selection algorithm therefore selects (1) all the interior vectors (LP and HP), (2) all the boundary LP vectors, and (3) an appropriate number of boundary HP vectors to complete the basis. In the particular case of a contiguous arbitrary-length 1-D signal, selecting all the interior vectors and only the boundary LP vectors is sufficient to guarantee a basis. For M-D AS signals we also select additional boundary vectors.

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 the preconditioned representation in [1] 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. 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 [8]. Therefore, a smooth signal will have a smooth representation, and a smooth representation will correspond to a smooth signal. These properties are beneficial for interpretation and compression.

This proposed selection 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 large change in the coefficient amplitudes, i.e. similar signals can have vastly different coefficient amplitudes. This may prohibit its usefulness in all but very 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 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 selection may have very limited practical usefulness.

High Performance and General ROS's Polynomial accuracy is theoretically appealing, however its usefulness for general signal/image processing is less evident. The ability to accurately approximate a constant signal (no DC-leakage into HP subbands) is very important, but higher order accuracy may be of limited value. More important properties include good conditioning, good frequency responses for the boundary filters, and flexibility in the possible ROS's that may be processed. These properties can be readily achieved by *selecting the vector centered at each sample in the signal's ROS*. For even-length vectors, the center can be taken to be the largest of the two center taps, i.e. the LP vector will be centered at one sample, and the HP at the other.

For each sample in the signal's ROS, we select 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 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 AS ROS. This approach works for a large class of filters: Daubechies 4-16 tap, Smith and Barnwell's orthogonal filters, and some biorthogonal filters. However, it does not work for all filters, e.g. Daubechies 2-tap (Haar) and some biorthogonal filters. In addition, since this approach selects centered vectors which have large overlap with the ROS, the resulting representation is much better conditioned and the "analysis" vectors have much better frequency responses than the previous scheme, where the selected vectors often are very small tails. Furthermore, there is considerably more freedom in the possible ROS's that may be represented: the ROS can have any shape (no restrictions on its curvature as in the previous scheme), it can contain holes, it can be disconnected, and it can include isolated samples. While we have not shown that we have complete freedom in the ROS, empirical evidence suggests that this is plausible.

This approach provides a number of advantages: a very simple selection based on the ROS, a large class of filters to choose from, good conditioning and frequency responses, and potentially complete freedom in the possible ROS's. Therefore it appears to be very promising as a general purpose wavelet-type representation for AS signals.

4.2. Computational Issues

Some applications, such as compression, require separate analysis and synthesis stages. The complexity of the analysis depends on: (1) determining the appropriate vectors to select based on the ROS, and (2) computing the coefficient amplitudes of the selected vectors. The selection process can be extremely simple, as discussed in the previous section. The coefficients can be computed in a variety of ways: (1) set up and solve the linear system, (2) construct the dual vectors and compute their inner products with the signal, (3) use iterative methods which exploit the structured initial basis, and (4) what may be referred to as the "peel-away" method. (1) and (2) are typically too cumbersome and complex for any large signal, while (3) requires a number of iterations and also is somewhat unsatisfying since we never get the exact answer. The "peel-away" method appears to be promising, especially for large ROS signals.

The "peel-away" method is easiest to visualize for an arbitrary-length 1-D signal, and when only selecting the LP boundary vectors. 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 for the innermost boundary vector (longest overlap with the signal) can be computed independent 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. In the general case when both LP and HP boundary vectors are selected, the supports are no longer staggered, however the problem can still be decomposed into a number of very low complexity problems. For example, in 2-D, every 2×2 set of samples in the boundary residual are represented by one to four boundary vectors — resulting in a low-dimensional linear system of equations to solve. In addition, since only a small number of different linear systems arise, all the inverses can be precomputed and stored. In essence, the "peel-away" method solves for the interior and then goes around the boundary, spiraling or peeling-away outward.

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, and the inverse transform is a simple inverse *block* transform possibly followed by cropping to the original ROS — thus decoding an AS signal is almost as easy as decoding a signal with rectangular support.

5. EXAMPLES AND SUMMARY

A two-level CS PR wavelet decomposition based on the second selection criterion is illustrated in Figure 1 for two AS signals. The top signal contains 7632 pixels, while the bottom one (with a very complex ROS) contains 3143 pixels. 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 AS 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.

We have developed a general approach for creating critically sampled perfect reconstruction representations; this approach is applicable to 1-D, 2-D, and general M-D AS problems, and it provides considerable freedom in the design. The wavelet-based representation appears very promising, and may potentially provide complete freedom in the ROS. This last feature is very important since it may often be difficult or impractical to impose constraints on the size, shape, or the number of samples in the ROS of AS M-D signals.

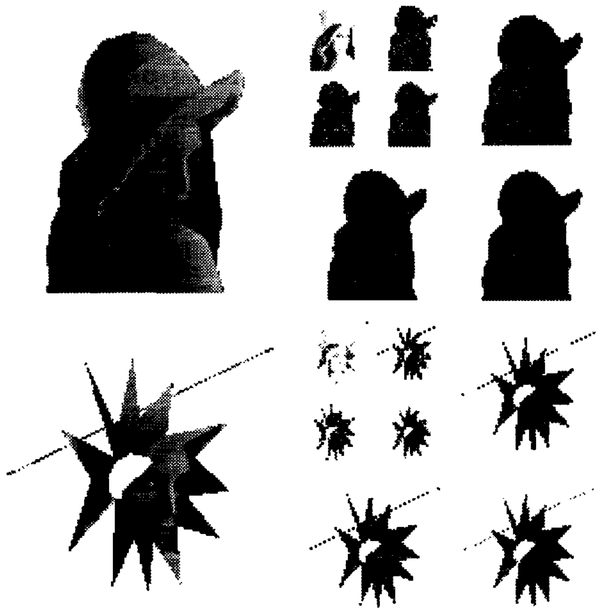

Figure 1: The wavelet representations of two 2-D AS signals using Daubechies 8-tap filters. To aid in interpretation, the lowpass subband has been scaled down by a factor of 2 and the other subbands have been offset to gray.

6. REFERENCES

[1] 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.

[2] 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.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] H. Chen, M. Civanlar, and B. Haskell, "A block transform coder for arbitrarily shaped image segments," *Inter. Conf. on Image Proc.*, vol. 1, pp. 85–89, Nov. 1994.

[8] J. Apostolopoulos. PhD thesis, MIT, To be submitted Spring 1997.