CRITICALLY SAMPLED WAVELET REPRESENTATIONS FOR MULTIDIMENSIONAL SIGNALS WITH ARBITRARY 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 an important element of many signal processing algorithms and applications. Until recently, representations have typically been designed for signals with convenient supports, e.g. 2-D signals with rectangular supports. However, a number of applications require representations for signals with arbitrary (non-rectangular) regions of support. We present a novel algorithm for creating critically sampled perfect reconstruction wavelet representations for signals defined over arbitrary supports. The proposed algorithm selects a subset of vectors from a convenient superset basis which under appropriate conditions provides a basis over the given arbitrary support. The algorithm can be interpreted as solving a corresponding sampling problem.

1. INTRODUCTION

Many applications involve multidimensional (M-D) signals defined over arbitrary regions of support. Examples include object/region-based representations for image and video, medical imaging, and the numerical solution of partial differential equations over arbitrary domains. In these applications, it is often useful to have transform/subband representations for signals with arbitrary supports.

In this paper we present a novel approach for creating critically sampled perfect reconstruction wavelet representations for discrete 1-D, 2-D, and general M-D signals defined over arbitrary regions of support (A-ROS). Specifically, we assume the A-ROS is given, and the goal is to represent the signal's amplitude over the A-ROS. This paper begins with a brief overview of previous research in this general area. We then review our general proposed approach [1] and present some interesting interpretations. We then present our wavelet-based approach and discuss some of its properties.

2. PREVIOUS RESEARCH

The previous research was primarily motivated by the problem of representing (transforming) signals defined over a single, contiguous, arbitrarily shaped support, such as an object in an image. The previous research can be roughly grouped into three classes. The first class of approaches embeds the signal in a superset space over which a transform can be conveniently defined/computed, e.g. a 2-D A-ROS signal can be embedded in a circumscribing square over which a DCT can be computed. A variety of different approaches based on this theme have been investigated [2, 3]. A common feature of these approaches is that they are overcomplete (oversampled) since there are more coefficients in the transform over the square than there are samples in the signal's ROS.

The second class of approaches constructs a basis over the A-ROS, for example via Gram-Schmidt [4, 5] or the KLT. This class of approaches may achieve the highest performance, since they attempt to construct a basis that is in some sense optimal. However, the complexity of the construction appears to limit their practical applications. The lifting scheme [6] does not explicitly construct a basis, but instead performs (invertible) local operations that ensure critical sampling and perfect reconstruction. Lifting as well as the sophisticated approach presented in [7] may potentially be very useful for applications that demand very high quality processing.

The third class of approaches apply 1-D transforms along each of the dimensions, e.g. a 2-D A-ROS signal can be processed by applying invertible 1-D transforms first along each of the rows and subsequently along each of the columns [8, 9]. Although these approaches are appealing in that they require only 1-D transforms, the fact that they are nonseparable may complicate their interpretation and subsequent processing.

3. PROPOSED GENERAL APPROACH

In this section we present a general approach that provides a number of desirable properties. The representation is based on known, separable transforms as defined over convenient supports, and also provides critical sampling and perfect reconstruction as follows from a basis defined over the A-ROS.

The approach can be achieved in the following manner [1, 10]. For purposes of illustration, and without loss of generality, consider a 2-D signal with an arbitrary (bounded) support. Embed the A-ROS signal in a superset space (support) over which a transform may be conveniently defined, e.g. the entire 2-D plane or a circumscribing square. For simplicity, consider a transform over a circumscribing square as illustrated in Figure 1. The transform over the square provides a basis over the square. The set of vectors in the basis over the square spans not only the signal space defined by the square, but also the signal space defined by the arbitrary support. We select an appropriate subset of the vectors defined over the square, such that they are linearly independent over the arbitrary support and thereby provide a basis over the arbitrary support.

Examining the Proposed General Approach: The proposed approach provides a significant amount of flexibility. For instance, it can in principle be applied using any superset transform/subband

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

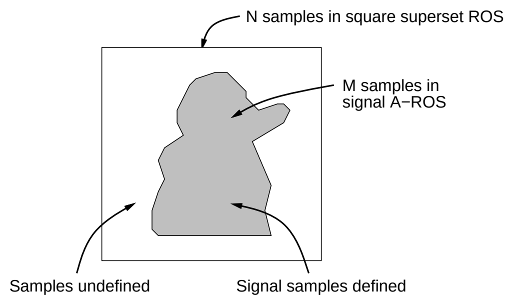

Figure 1: Example of general proposed approach: The given 2-D A-ROS signal whose support is shown in gray and contains M samples is embedded in an N-sample circumscribing square over which a transform (basis) is defined. We select M out of the N vectors from the superset basis such that they are linearly independent over the A-ROS and thereby provide a basis over the A-ROS.

representation, e.g. DFT, DCT, wavelet. In addition, given any superset basis, typically there are many subsets of M out of N vectors that provide a basis over the given A-ROS, where some of these bases may be in some sense better than the others. Furthermore, the selected basis may be used for either "analyzing" or "synthesizing" the signal—we create a biorthogonal representation and explicitly choose *either* the analysis or the synthesis basis. We next briefly consider these two possibilities as they provide different insights and may be useful for different applications.

Selecting a basis to analyze the signal (analysis basis) is equivalent to the following *sampling problem*. Begin by extrapolating the A-ROS signal to fill the circumscribing square, for instance by zero padding, and then compute a transform over the square, as shown in Figure 2. Since the square contains N samples, the transform over the square produces N coefficients. Selecting a basis to analyze the signal is equivalent to selecting (sampling) M out of the N coefficients such that with only those M coefficients the original signal can be perfectly recovered.

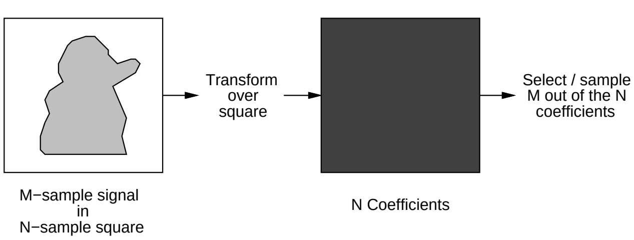

Figure 2: Selecting a basis to analyze the A-ROS signal.

Selecting a basis to synthesize the signal (synthesis basis) is equivalent to determining which M out of N coefficients in the square one needs (the remaining $(N-M)$ coefficients are zero) such that computing an inverse transform over the square recovers the original signal, as shown in Figure 3. The inverse transform produces the A-ROS signal with some extrapolation, but the signal can simply be picked out since its support is known.

Determining for any arbitrary support a subset of M coefficients (vectors) that enables perfect reconstruction is in general a very difficult problem. While the linear independence of various subsets of M vectors can be explicitly examined (e.g. by Gram-

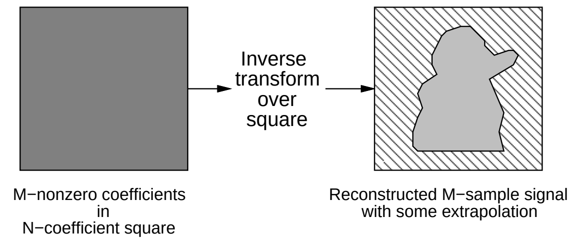

Figure 3: Selecting a basis to synthesize the A-ROS signal.

Schmidt), this is very computationally complex. In Section 4.1 we present a simple algorithm that selects an appropriate subset of vectors to provide a basis (solves the sampling problem) for any arbitrary support.

4. WAVELET-BASED REPRESENTATIONS

In this section we present a critically sampled wavelet representation which is an extension of our previous work [1]. We focus on a wavelet-based representation because wavelets are useful in many applications, and also because they lead to a particularly simple and appealing solution. We consider the case of a 2-channel filterbank in 1-D or a 2×2 channel filterbank in 2-D, since once we can solve the problem of determining a basis over any arbitrary support for these filterbanks, the solution can be recursively applied to any of the subbands to create wavelet- or wavelet-packet-based representations.

In [1] it was shown that the important polynomial accuracy property (provided by conventional wavelet transforms of signals with convenient supports) can be preserved within the context of our simple approach for determining a basis over any A-ROS by selecting/discarding vectors from a superset basis. However, two disadvantages arise which limit the practical applications: as the filter length increases (1) there are more stringent constraints on the possible supports that may be represented and (2) the resulting representation becomes increasingly ill-conditioned. In [1] we also briefly described how relaxing the polynomial accuracy property may enable an algorithm that overcomes these disadvantages.

4.1. Algorithm to Select a Basis/Solve the Sampling Problem

In this section, we present a simple algorithm for selecting a subset of vectors from a superset basis which under certain conditions that we develop provides a basis (solves the sampling problem) over any arbitrary support [10]. We first describe the algorithm and then analyze its properties and computational issues.

This algorithm establishes a 1-to-1 association between samples in the superset ROS and vectors in the superset basis; given any arbitrary ROS, the associated vectors are selected as a (potential) basis. This approach provides the correct number of vectors needed for a basis, and the association must guarantee that the selected vectors are linearly independent over the A-ROS. A remarkable feature of this approach is that for a very broad class of possible lowpass/highpass (LP/HP) filter pairs, there exists a simple association that guarantees a basis for all possible arbitrary supports. We next present an example of one possible association.

For simplicity, consider the 1-D problem and a 2-channel filterbank with an appropriately chosen LP/HP filter pair. The basis vectors for the superset basis over $l^2(\mathbb{Z})$ are formed by all even translates of the LP and HP filter impulse responses. For purpose of illustration, consider typical odd-length, symmetric filters as are commonly used in the image processing community, e.g. Daubechies 9/7-tap biorthogonal filters. These filters have natural center taps, and the center taps of the LP and HP filters are located at adjacent locations (not at the same location). Hence, the center of each of the superset basis vectors is located at a different location. An appealing and simple 1-to-1 mapping between samples in the support and vectors to select can be obtained by associating to each sample the vector whose center tap is located at that sample. This association strategy generalizes to a much larger class of filters, and also extends trivially to 2-D or general M-D. Note that for a given filter pair there may exist a number of possible associations that guarantee a basis. One appealing and generally applicable association is based on the largest filter taps since this (1) is natural, (2) leads to a well-conditioned representation, and (3) simplifies the proof that the association provides a basis.

4.2. Analysis of Algorithm

We next develop the required property for the simple association method to provide a basis over any possible arbitrary support [10]. In particular, we consider the problem of selecting a basis to synthesize the signal (the problem of selecting an analysis basis is analogous). Transform synthesis can be expressed as

$$\underset{N\times N}{\begin{bmatrix} & & \\ & T^{-1} & \\ & & \end{bmatrix}} \underset{N\times 1}{\begin{bmatrix} \\ c \\ \\ \end{bmatrix}} = \underset{N\times 1}{\begin{bmatrix} \\ f \\ \\ \end{bmatrix}}$$

where f represents the signal, T^{-1} is the transform synthesis matrix whose columns contain the basis vectors, and c contains the coefficients describing the linear combination of the basis vectors that would recover the signal. For the purpose of illustration, and without loss of generality, consider the case when the signal has an arbitrary support given by $\alpha, \beta,$ and γ, and the samples outside the support are denoted by don't cares "$\times$".

$$\underset{N\times N}{\begin{bmatrix} \times & \times & \times & \times & \times \\ & & & & \\ & & & & \\ \times & \times & \times & \times & \times \\ & & & & \end{bmatrix}} \underset{N\times 1}{\begin{bmatrix} \\ \\ \\ \\ \\ \end{bmatrix}} = \underset{N\times 1}{\begin{bmatrix} \times \\ \alpha \\ \beta \\ \times \\ \gamma \end{bmatrix}}$$

Due to the $\times$'s in f some rows of the matrix are useful while others are not (signified by $\times$'s) and can be discarded—only the portion of each superset vector that overlaps the A-ROS is useful for representing the A-ROS signal. Collapsing the matrix and vector leads to an underdetermined linear system of equations or an overcomplete representation.

$$\underset{M\times N}{\begin{bmatrix} & & \\ & & \\ & & \end{bmatrix}} \underset{N\times 1}{\begin{bmatrix} \\ \\ \\ \\ \\ \end{bmatrix}} = \underset{M\times 1}{\begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix}}$$

Since T^{-1} is nonsingular there exists at least one subset of M columns that are linearly independent, and although one can determine a subset by explicitly examining the vectors that would be excessively complex. Our proposed association corresponds to a simple algorithm for selecting M out of the N columns such that the resulting $M \times M$ submatrix is nonsingular; this holds true for any arbitrary support.

Our proposed association for selecting vectors (coupled with appropriate indexing/shifting of the filter impulse responses) corresponds to selecting columns with the same indices as the rows that were retained. The resulting $M \times M$ submatrix is then a principal submatrix of the original matrix.

$$\underset{N\times N}{\begin{bmatrix} \times & \times & \times & \times & \times \\ & & & & \\ & & & & \\ \times & \times & \times & \times & \times \\ & & & & \end{bmatrix}} \Longrightarrow \underset{N\times N}{\begin{bmatrix} \times & \times & \times & \times & \times \\ \times & & & \times & \\ \times & & & \times & \\ \times & \times & \times & \times & \times \\ \times & & & \times & \end{bmatrix}} \Longrightarrow \underset{M\times M}{\begin{bmatrix} & & \\ & & \\ & & \end{bmatrix}}$$

Therefore, determining whether the proposed association provides a basis over the given A-ROS is equivalent to determining whether the resulting $M \times M$ principal submatrix is nonsingular. Furthermore, determining whether *the proposed representation provides a basis for all possible arbitrary supports* is equivalent to determining whether *all possible principal submatrices are nonsingular*. This reformulation provides significant structure which greatly facilitates further analysis. For instance, given an $N \times N$ matrix A, if $x^T Ax \neq 0$ for all nonzero $N \times 1$ vectors x, then all possible principal submatrices of A are nonsingular.

We refer to a LP/HP filter pair as admissible (for a particular dimension, e.g. 1-D or 2-D) if the proposed representation guarantees a basis for any arbitrary ROS (in that dimension). For 1-D, we have developed a very simple and successful sufficient condition on the filter taps such that if satisfied, the proposed representation is guaranteed to provide a basis over any possible arbitrary support [10]. We examined a number of popular filters including orthogonal filters (Daubechies 2-16 tap, Smith-Barnwell 8-tap), pseudo-semiorthogonal filters (Simoncelli 9-tap), and linear-phase biorthogonal filters (Daubechies 9/7-tap, Sweldens and Strang "Binary" 9/7-tap, and Le Gall 3/5-tap). Every filter pair tested is admissible in 1-D—the representation provides a basis for any arbitrary 1-D support using any of the tested filters. Inadmissible filters can be constructed, however these are contrived cases and do not correspond to conventional filters in the sense that they do not have meaningful LP/HP frequency responses.

The 1-D analysis and conditions can be extended to 2-D and general M-D, but a straightforward extension leads to relatively narrow sufficient conditions, i.e. conditions that cover only a small fraction of the filter sets that appear to be admissible. Admissibility can be examined by individually analyzing each filter pair and exploiting its specific features [10]. To briefly summarize, of all the filters tested, the Haar filters are the only filters that have been shown to be inadmissible in 2-D, i.e. there exist 2-D A-ROS's for which the association with Haar filters does not provide a basis. Theoretical considerations and empirical evidence strongly suggest that a very broad class of filters are admissible, while a very small and special class of filters (such as Haar) are not [10]. Details will appear in a forthcoming paper.

4.3. Computational Issues

We consider the problem of analyzing (decomposing) and synthesizing (reconstructing) an A-ROS signal with the proposed approach. Specifically, we examine the case where a synthesis basis is selected, as discussed in Section 3. The case where an analysis basis is selected is similar, with the analysis and synthesis operations exchanged. The complexity of this approach depends on three operations: (1) selecting a basis given the signal's support, (2) computing the coefficient amplitudes to represent the signal with respect to the basis, and (3) using the basis and coefficients to reconstruct the signal.

The proposed association provides an extremely simple method for determining a basis over any A-ROS. For illustration, consider 1-D supports and a 2-channel filterbank. Shift the LP and HP filter impulse responses so that their "chosen taps" occur at $n = 0$ and $n = 1$, respectively, and express the output of the two-channel filterbank in its transform representation (alternating LP and HP coefficients as opposed to collecting all the LP and all the HP together as in the subband representation). With this organization the binary (select/discard) mask for selecting vectors is exactly given by the binary (within-A-ROS/outside-A-ROS) mask given by the signal's support. This simple relationship between selection and support extends directly to 2-D and general M-D A-ROS signals.

Practically, it may be beneficial to use an iterative algorithm for accurately estimating the coefficient amplitudes while requiring a relatively small amount of complexity. This approach was motivated by [3] where the coefficient amplitudes are estimated by computing separable/fast transforms over a square. In the case of an orthogonal transform over the square, we have orthogonal projections onto two convex sets within the superset space: (1) a subspace given by the span of the selected vectors, and (2) an affine (translated) subspace corresponding to all possible extrapolations of the A-ROS signal. Alternating (orthogonal) projections onto these two convex sets is guaranteed to converge to an element in their union, and since the sets intersect at a unique point, the iteration is guaranteed to converge to the solution.

The theory of POCS guarantees convergence for all admissible orthogonal filters, but it does not apply for biorthogonal filters because they involve oblique (non-orthogonal, i.e. non-closest-point) projections. However, empirical evidence shows that the iteration does converge for typical admissible biorthogonal filters. Specifically, the structure produced by the proposed association (selection/sampling) produces a converging sequence. Furthermore, the convergence for typical biorthogonal filters is faster than for orthogonal filters. The fast convergence (depending on association, filters, A-ROS, signal, and relaxation; approximately 10 dB/iteration) suggests that sufficient accuracy for typical image processing applications may be achieved in a very small number of iterations.

The signal reconstruction is very simple; it requires determining the basis given the A-ROS, computing an inverse transform over the superset basis, and extracting the signal as shown in Figure 3.

5. CONCLUDING REMARKS

The proposed approach provides a simple algorithm to determine a basis (solve the sampling problem) for signals with completely arbitrary supports. Specifically, for admissible filters, the approach is applicable to any possible arbitrary support. Theoretical considerations and empirical evidence strongly suggest that there exists a very large class of admissible filters.

The approach provides good quality representations for objects in an image, as shown in Figure 4. The representation also enables a natural reconstruction of the A-ROS signal at different scales, and at each scale the representation is critically sampled.

Polynomial accuracy is in general not preserved by the proposed association. While fully preserving polynomial accuracy may often not be important, it is important to minimize the DC-leakage into the highpass subbands. The proposed approach may exhibit undesirable DC-leakage into the boundary highpass coefficients depending on the filters, specifics of the association, and A-ROS. The filters and association can be designed to minimize this effect.

The representation appears to be stable to perturbations in the signal. Specifically, a change in the A-ROS produces a change of the same "size" in the basis. Also, the representation appears to be well-conditioned to changes in the signal's amplitude.

Computationally, one operation (analysis or synthesis) is simple and fast while the other is simple and relatively fast, requiring a small number of iterations.

Overall, this wavelet-based representation appears promising for representing signals with arbitrary supports, and in particular the objects/regions within an image or video.

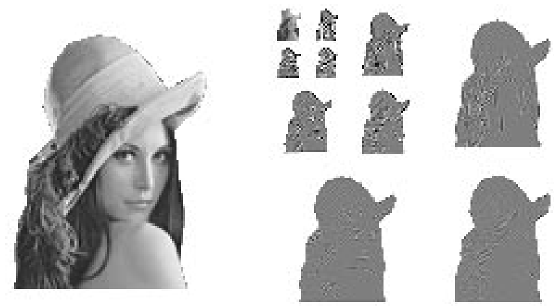

Figure 4: The wavelet representation of a 2-D A-ROS signal where a synthesis basis was selected and using Daubechies 9/7-tap biorthogonal filters. To aid in interpretation, the lowpass subband has been attenuated and the other subbands have been offset to gray.

6. REFERENCES

[1] J. Apostolopoulos and J. Lim, "Transform/subband representations for signals with arbitrarily shaped regions of support," in *Inter. Conf. Acoustics, Speech and Signal Processing*, April 1997.

[2] 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.

[3] 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.

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

[5] C. Herley and J. Kovačević, "Spatially varying two-dimensional filter banks," *Proc. Asilomar Conf.*, Nov. 1993.

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

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

[8] 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, The Netherlands: Elsevier, 1994.

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

[10] J. Apostolopoulos, *Transform/Subband Representations for Multidimensional Signals with Arbitrary Regions of Support*. PhD thesis, MIT, May 1997.