Directional Complex-Wavelet Processing

Felix C. A. Fernandes[r], Rutger van Spaendonck[d], Mark J. Coates[r] and C. Sidney Burrus[r]

[r]Department of Electrical and Computer Engineering, Rice University, Houston, TX 77251-1892, USA

[d]Applied Earth Sciences, Delft University of Technology, Delft, Netherlands

ABSTRACT

Poor directional selectivity, a major disadvantage of the separable 2D discrete wavelet transform (DWT), has previously been circumvented either by using highly redundant, nonseparable wavelet transforms or by using restrictive designs to obtain a pair of wavelet trees. In this paper, we demonstrate that superior directional selectivity may be obtained with *no* redundancy in *any* separable wavelet transform. We achieve this by projecting the wavelet coefficients to separate approximately the positive and negative frequencies. Subsequent decimation maintains non-redundancy. A novel reconstruction step guarantees perfect reconstruction within this critically-sampled framework. Although our transform generates complex-valued coefficients, it may be implemented with a fast algorithm that uses only real arithmetic. We also explain how redundancy may be judiciously introduced into our transform to benefit certain applications. To demonstrate the efficacy of our projection technique, we show that it achieves state-of-the-art performance in a seismic image-processing application.

Keywords: directional, non-redundant, complex, wavelet, projection, separable, redundant, orientation, directionality

1. INTRODUCTION

The separable 2D discrete wavelet transform (DWT) is a powerful image-processing tool, but in some applications its poor directional selectivity is a serious disadvantage. The transform can only distinguish between three different orientations of spatial features. Nonseparable 2D DWTs can provide directional selectivity,[1,2] but these involve complicated design problems and are more computationally expensive. Kingsbury's dual-tree complex wavelet transform[3] is separable and has impressive directional selectivity due to an approximate quadrature relationship between its trees. However, Kingsbury's transform involves a redundancy of $2N$ in N dimensions, and neither tree in the dual-tree structure genuinely corresponds to a wavelet transform. Enforcing the quadrature relationship between the two trees is complicated, impeding the incorporation of other wavelet criteria in the design.

This paper extends our earlier work on critically sampled, directional, complex wavelet transforms.[4] In Section 3, we propose the non-redundant post-projection transform, a separable wavelet transform that provides comparable directional selectivity to Kingsbury's transform, but introduces *no* redundancy. The method involves the application of projection filters to discriminate between positive and negative frequency components of 2D separable DWT wavelet coefficients. This results in complex-valued coefficients. Decimation of these coefficients preserves directional selectivity but makes the transform non-redundant. In Section 5 we show that we can achieve perfect reconstruction using a novel synthesis filter bank structure, if the projection filters satisfy certain criteria. The ability to obtain a directional, non-redundant, complex-valued decomposition of images allows for interesting variations on traditional processing applications such as compression and denoising. Since some applications benefit from a small amount of transform redundancy, in Section 4 we introduce the redundant pre-projection transform, which has the same redundancy as the dual-tree wavelet transform.

2. POOR DIRECTIONAL SELECTIVITY IN THE DWT

In this section, we explain the poor directional selectivity of the 2D separable DWT based on the Fourier-plane partitioning associated with it. We then show how non-redundant post-projection uses separable complex-coefficient filters to provide a different Fourier-plane partitioning with improved directional selectivity.

This work was supported by DARPA, Texas Instruments, the Dutch Science Foundation's NEESDI program and the Consortium for Computational Seismic Interpretation, Rice University.

Email: felixf@rice.edu, csb@rice.edu

Web: http://www.dsp.rice.edu

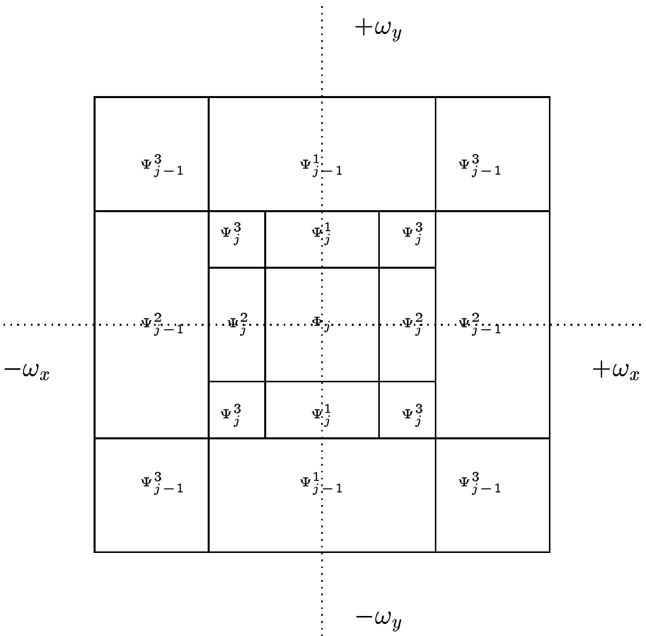

Figure 1. Frequency-domain energy localization of tensor wavelets in a two-level 2D DWT.

Figure 1 shows the Fourier-plane partitioning obtained using a separable, real-valued, 2-level 2D DWT. The tensor wavelet $\psi^3_{j-1}(x,y)$ is associated with diagonally-oriented spatial features at scale $j-1$ and concentrates its energy in the four blocks labeled ψ^3_{j-1}. High energy at the output of the filter in the subband associated with $\psi^3_{j-1}(x,y)$ indicates the presence of this class of features. The ψ^3_{j-1} blocks in the upper-left and lower-right corners indicate features at -45 degrees, while those in the upper-right and lower-left corners indicate features at $+45$ degrees. Since all four blocks are associated with the output of one filter in the DWT, we cannot differentiate between these two orientations. Similarly, features oriented at $+15$ degrees and -15 degrees are indistinguishable from their wavelet coefficients since they are both associated with energy in the ψ^2_{j-1} subband. Features oriented at $+75$ degrees and -75 degrees suffer a similar fate. It is evident that the poor directional selectivity of the real DWT is due to positively- and negatively-oriented blocks both being associated with the same subband.

3. NON-REDUNDANT POST-PROJECTION

Decoupling the positively- and negatively-oriented blocks associated with each subband of the DWT improves directional selectivity. Figure 2 shows how this may be done using our non-redundant post-projection framework. The first two columns of filters obtain the first level of the DWT and output the subbands with the $j-1$ subscripts in Figure 1. The filters in the third and fourth columns have $+(-)$ superscripts to indicate that they retain positive (negative) frequencies and suppress negative (positive) frequencies. These $+(-)$ filters are one-dimensional projection filters that project real signals onto subspaces consisting of only positive (negative) frequencies on the Fourier line. These projection filters enhance the directional selectivity of the real DWT transform that precedes them by decoupling the subband blocks in Figure 1 to obtain the subbands in Figure 3.

Consider the subband corresponding to the two blocks labeled ψ^1_{j-1} in Figure 1. If we filter this subband in the y direction (along columns) with a filter h^+_y as shown in Figure 2, we obtain a new subband containing the two blocks labeled ψ^{1-}_{j-1} and ψ^{1+}_{j-1} in Figure 3. We obtain a decoupled subband ψ^{1+}_{j-1} by filtering this new subband in the x direction (along rows) with h^+_x. The other decoupled subband ψ^{1-}_{j-1} is similarly obtained by row filtering with h^-_x. The decoupled subband ψ^{1-}_{j-1} indicates features at -75 degrees, while the other decoupled subband ψ^{1+}_{j-1} indicates features at $+75$ degrees.

By applying the projection filters to the other subbands in Figure 1, we can now distinguish between features oriented in 6 directions at any particular scale. At scale $j-1$, these 6 directions correspond to the 6 directional subbands labeled ψ^{2+}_{j-1}, ψ^{3+}_{j-1}, ψ^{1+}_{j-1}, ψ^{1-}_{j-1}, ψ^{3-}_{j-1} and ψ^{2-}_{j-1} in Figure 3. These directional subbands correspond to

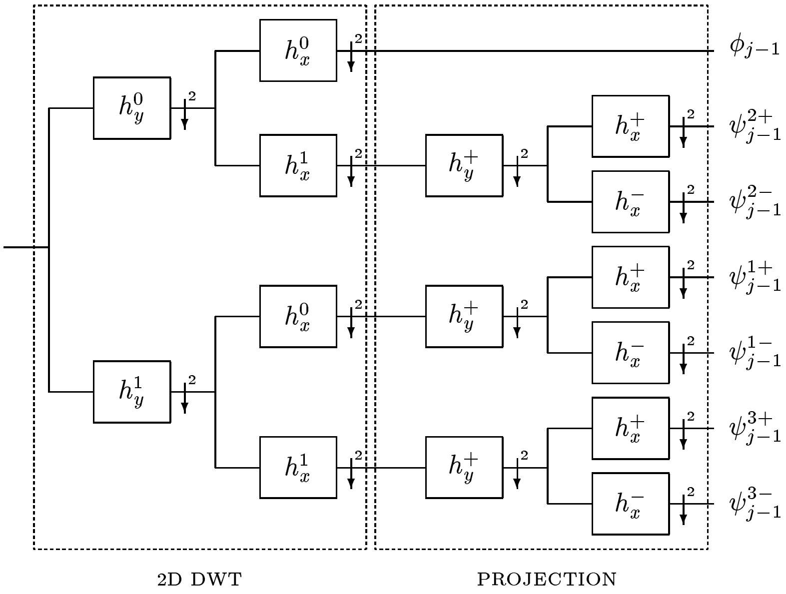

Figure 2. Non-redundant post-projection analysis filter-bank structure for separable, directional, complex 2D DWT . The y (x) subscripts indicate filtering along the columns (rows) of an image. The $+$ ($-$) superscripts indicate projection filters that attenuate negative (positive) frequencies.

features oriented at +15, +45, +75, -75, -45 and -15 degrees respectively. To get the directional subbands at scale j in Figure 3, we iterate the processing in Figure 2 on the ϕ_{j-1} subband. We can then use the projection filters to decouple the lowpass subband ϕ_j into ϕ_{j-1}^- and ϕ_{j-1}^+ blocks that discriminate between low-scale features oriented at -45 and +45 degrees.

In practice, it is impossible to retain positive (negative) frequencies and suppress negative (positive) frequencies perfectly, with a finite length filter. Instead, we use projection filters that attenuate negative (or positive) frequencies. To compare the directional selectivity of various different transforms, we used the zone-plate image displayed in Figure 4 (a). It contains features oriented at all possible directions with low-scale features in the center and high scale features at the periphery.

The separable, single-level Daubechies 2D DWT of the image is displayed in Figure 4 (c). The wavelet coefficients for both the +45 and -45 degree orientations appear in the diagonal-orientation subband and cannot be distinguished. The transform also suffers from severe aliasing artifacts. Figure 4 (b) shows the 2D, single-level CDWT obtained using Kingsbury's dual-tree (upper right). The transform shows enhanced directional selectivity but also some aliasing. Furthermore the lowpass band has not been transformed, but is cut into halves. This representation is four times redundant for images. Figure 4 (d) presents the 2D, single-level post-projection CDWT obtained using the filter-bank structure of Figure 2. Both the dual-tree transform and the post-projection transform provide six highpass directional subbands. In addition, the post-projection transform discriminates between features oriented at +45 and -45 degrees in the lowpass subband.

Although the post-projection CDWT uses length-4 filters and the dual-tree CDWT uses length-10 filters, the features are better localized in each subband of the former. We emphasize that the dual-tree transform has a redundancy of four whereas our post-projection transform is critically sampled. However, there are some applications for which transforms with a small amount of redundancy yield significantly improved results. In the next section, we introduce an alternative, redundant projection-based transform that proves beneficial for such applications.

4. REDUNDANT PRE-PROJECTION

The non-redundant post-projection transform introduced in the previous section employs projection filters on the subbands of a DWT to obtain the Fourier-plane tiling shown in Figure 3. Another way to arrive at this tiling is to perform first the projection operation on the image to suppress its positive (or negative) frequencies, and

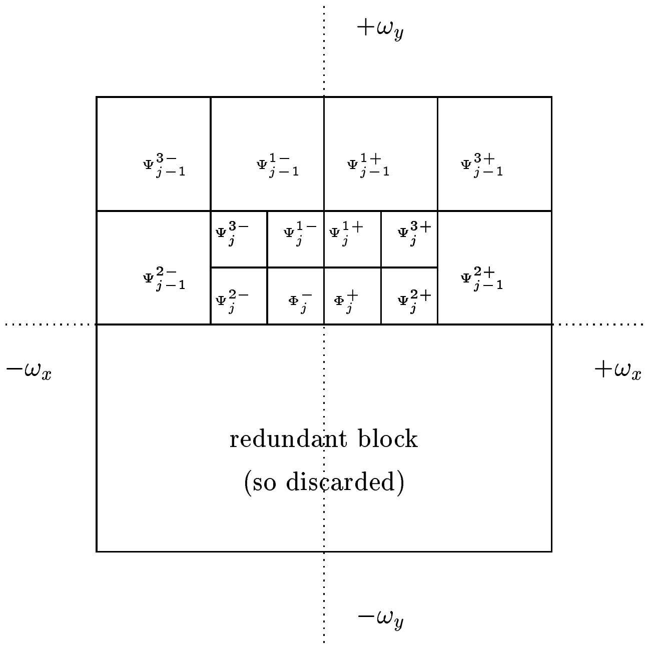

Figure 3. Fourier plane partitioning after post-projection.

then compute the DWT of the image projection. We call this the non-redundant pre-projection transform since the projection precedes the DWT. The filter bank implementation of the non-redundant pre-projection transform is shown in Figure 5(a). Eliminating the downsamplers in Figure 5(a) results in the redundant pre-projection transform shown in Figure 5(b).

Pre-projection provides a convenient mechanism for restricting the redundancy required to both suppress aliasing and provide directional selectivity. The imaginary part of the image projection is approximately the Hilbert transform of the real part of the image projection. Therefore the real and imaginary parts of the projected image are related through an approximate quadrature relationship. This is useful for certain applications.[5] Moreover, it means that less redundancy (a factor of two times the number of dimensions) is required to restrict aliasing and provide a genuine indication of where image energy is located in the frequency domain.

Having established that the redundant pre-projection transform performs better than the redundant post-projection transform, we now compare the non-redundant versions of these transforms. To do so, we emphasize that in either transform, the DWT performs a multiscale decomposition of the input while the projection operator extracts directional information from its input. In the non-redundant post-projection transform the multiscale decomposition is performed on an unaliased image while the directional information is extracted from an aliased image. Hence aliasing artifacts impair the directional selectivity of the transform rather than its multiscale nature. On the other hand, in the non-redundant pre-projection transform aliasing artifacts impair the multiscale nature of the transform rather than its directional selectivity. Since the multiscale property is usually more important than directional selectivity in most applications, post-projection fares better than pre-projection in a non-redundant framework.

We tested the directional selectivity of the redundant pre-projection transform on the zone-plate image in Figure 4(a). The resulting transform envelope is shown in Figure 4(e). We used the length-4 Daubechies filters to generate the projection filter and compute the DWT. The transform shows the same angles as the transform with post-projection but is better localized in space because there is no downsampling after projection.

5. PERFECT RECONSTRUCTION

In Sections 3 and 4, we introduced projection filters that provide transform coefficients with directional information. This information may be exploited for image analysis or image processing. In an image processing application, a

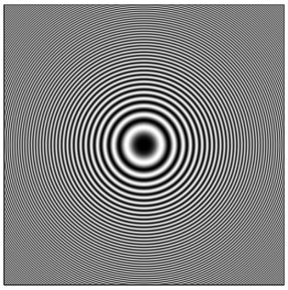

(a) zone-plate

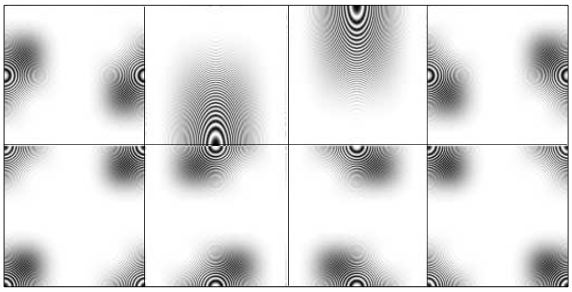

(b) dual-tree

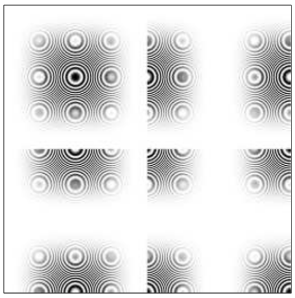

(c) DWT

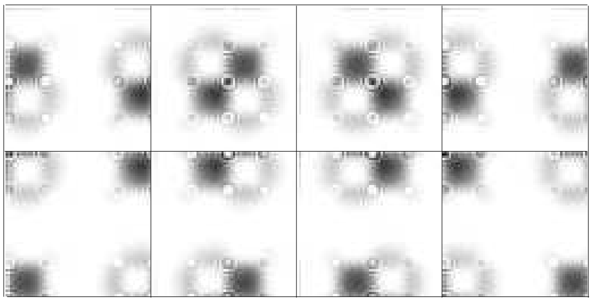

(d) post-projection

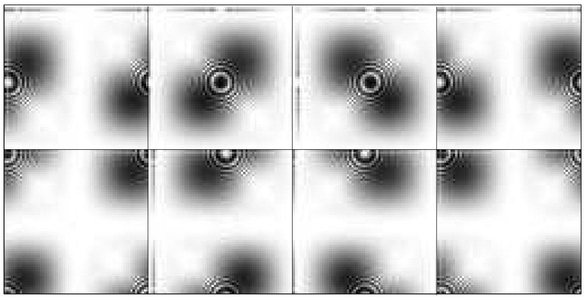

(e) pre-projection

Figure 4. **(a)** Zone-plate Image. **(b)** Dual-tree CDWT. Clockwise from upper-rightmost block: -75, +45, -15, +15, -45, +75 degree highpass subband, upper half of lowpass subband, lower half of lowpass subband. **(c)** 2D Daubechies DWT. Clockwise from upper-right block: vertical-, diagonal-, horizontal-orientation subbands and lowpass subband. **(d)** Directionally post-projected DWT. Clockwise from upper-rightmost block: -75, +45, -15, +15, -45, +75 degree highpass subbands, -45, +45 degree lowpass subbands. Length-4 Daubechies filters are used for the wavelet decomposition and to generate the projection filter. **(e)** Directionally pre-projected DWT.

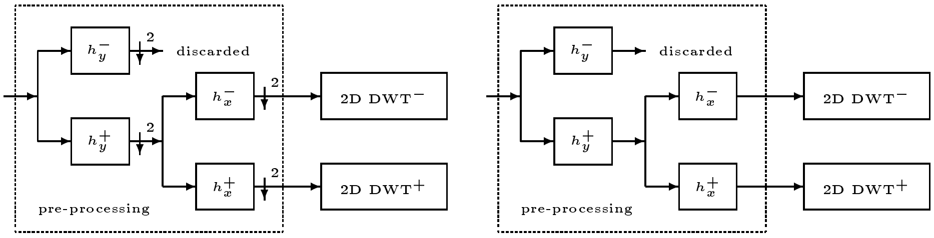

(a) with downsampling (b) without downsampling

Figure 5. Analysis filter bank structure for the separable complex wavelet transform with pre-projection (**a**) pre-projection with down-sampling: no redundancy (**b**) pre-projection without down-sampling: redundancy = 4.

processed version of the image is reconstructed from its processed transform coefficients. In the absence of any processing, this is called the perfect-reconstruction problem.

Obtaining perfect reconstruction for the non-redundant post-projection transform requires the design of a synthesis filter bank that inverts the effect of the analysis filter bank in Figure 2. The DWT is easily inverted by a synthesis filter bank that performs the inverse DWT. The filter pairs (h_x^+, h_x^-) operate on a complex input signal that is the downsampled output of a column projection filter h_y^+. As explained in implication (3) of Theorem 1 below, (h_x^+, h_x^-) are selected to be the analysis filters of a 2-band filter bank. Their input is then reconstructed by the corresponding synthesis filters of the 2-band filter bank. Implication (1) of Theorem 1 proves that the synthesis block in Figure 6 inverts the projection block in that figure. This projection block corresponds to the column projection filters h_y^+ that have real-valued input in Figure 2. Therefore the synthesis block in Figure 6 may be used to reconstruct the input to the column projection filters h_y^+. The synthesis structures discussed above may also be used to reconstruct the inputs to the projection filters in the non-redundant pre-projection transform in Figure 5(a).

We now consider perfect reconstruction for the redundant pre-projection transform in Figure 5(b). In this case, the filter pairs (h_x^+, h_x^-) operate on a complex input signal that is the output of a column projection filter h_y^+. As explained in implication (4) of Theorem 1 below, (h_x^+, h_x^-) are selected to be the analysis filters of a 2-band non-decimated filter bank. Their input is then reconstructed by the corresponding synthesis filters of the 2-band non-decimated filter bank. Implication (2) of Theorem 1 proves that the synthesis block in Figure 7 reconstructs the input to the column projection filters h_y^+ that have real-valued input in Figure 2.

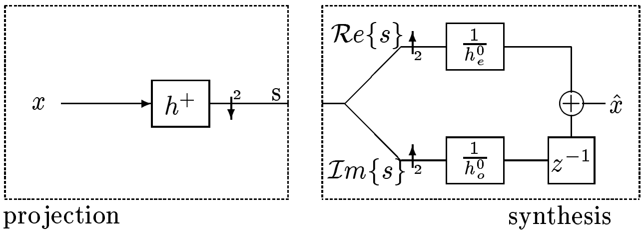

Figure 6. Reconstruction of real-valued input after projection and downsampling. The symbols $\frac{1}{h_e^0}, \frac{1}{h_e^0}$ represent the IIR filters $\frac{1}{H_e^0(-z^2)}, \frac{1}{H_e^0(-z^2)}$ respectively.

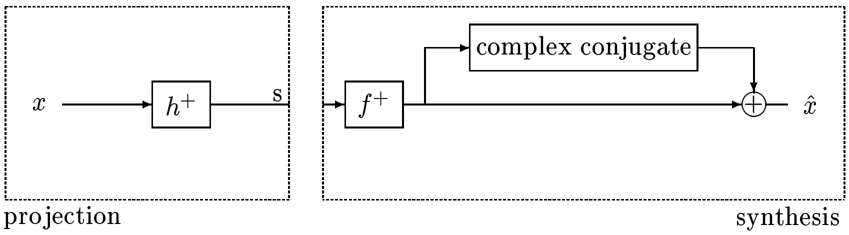

Figure 7. Reconstruction of real-valued input after projection without downsampling.

Theorem 1. Let h^0 be the lowpass analysis filter of a two-band, real-coefficient filter bank (h^0, h^1, f^0, f^1). Denote the nonzero polyphase components of h^0 by h_e^0, h_o^0. Create a projection filter h^+ by frequency shifting the low-pass transform H^0 by $\pi/2$, so that

$$H^+(z) = H^0(-jz). \tag{1}$$

Then the following statements are true.
1. If $x \in \mathcal{L}^2(\mathcal{R} \to \mathcal{R})$, then $\hat{x} = x$ in Figure 6.
2. If $x \in \mathcal{L}^2(\mathcal{R} \to \mathcal{R})$ and $F^+(z) = F^0(-jz)$, then $\hat{x} = x$ in Figure 7.
3. If $x \in \mathcal{L}^2(\mathcal{R} \to \mathcal{C})$, then the 2-band filterbank (h^+, h^-, f^+, f^-) generated from h^+ perfectly reconstructs x.
4. If $x \in \mathcal{L}^2(\mathcal{R} \to \mathcal{C})$, then the 2-band non-decimated filterbank (h^+, h^-, f^+, f^-) generated from h^+ perfectly reconstructs x.

Proof:

1. We have $H_e^+(z) = H_e^0(-z)$ and $H_o^+(z) = jH_o^0(-z)$. Let $\downarrow h^+$ denote the linear transform that maps the input x to its projection s. Then $S(z) = H_e^+(z)X_e(z) + z^{-1}H_o^+(z)X_o(z)$ and the kernel of the transform is $\ker(\downarrow h^+) = \{X(z) : H_e^0(-z)X_e(z) = -jz^{-1}H_o^0(-z)X_o(z)\} = \{0\}$. Therefore, the mapping $\downarrow h^+$ is invertible. The inverse mapping to $\hat{x}$ is given by $\hat{x} = \mathcal{R}e\{S(z^2)\}/H_e^0(-z^2) + z^{-1}\mathcal{I}m\{S(z^2)\}/H_o^0(-z^2)$, where $\frac{1}{H_e^0(-z^2)}$ and $\frac{1}{H_o^0(-z^2)}$ are infinite-impulse response (IIR) synthesis filters. This proves (1).

2. We have

$$\begin{aligned}
\hat{X}(z) &= (F^+(z)H^+(z) + [F^+(z)H^+(z)]^*)X(z), \\
&= (F^0(-jz)H^0(-jz) + F^0(jz)H^0(jz))X(z), \\
&= (F^0(\nu)H^0(\nu) + F^0(-\nu)H^0(-\nu))X(z), \quad \text{where } \nu = -jz, \\
&= X(z), \quad \text{by the halfband condition.}
\end{aligned}$$

3. Implication (3) follows from the definition of a 2-band filter bank.

4. Implication (4) follows from the definition of a 2-band non-decimated filter bank. □

6. CRITERIA FOR FILTER DESIGN

In Sections 3 and 4, we explained how the use of projection filters that attenuate negative (or positive) frequencies may be used to enhance the directional selectivity of the DWT. The directional selectivity is directly related to the frequency-domain performance of the projection filter; longer projection filters have better frequency-domain characteristics, and hence can impart greater directional selectivity to the DWT. However, in Theorem 1, we proved that to obtain perfect reconstruction in a non-redundant framework, it is necessary to use IIR synthesis filters obtained from the polyphase components of the projection filter. For practical reasons, the design of the projection filter $h+$ must ensure that the IIR synthesis filters $1/H_e^+(z^2)$ $1/H_o^+(z^2)$ may be implemented in a stable manner. The experiments in this paper have used the orthonormal, length-4, Daubechies scaling filter for the h^0 filter that is used to generate the projection filter h^+ according to Equation 1. The associated IIR synthesis filters have no poles on or outside the unit circle and hence enjoy a fast causal implementation. In general, however, we may need to use longer projection filters to obtain improved directional selectivity. This section explains how the IIR synthesis filters that are associated with arbitrary projection filters may be implemented in a stable manner.

Using IIR filters requires the imposition of boundary conditions on the IIR filter inputs in order to determine the initial conditions under which they operate. When processing finite-length images, the symmetric extension method in which an image is extended by reflection at its boundaries, performs well compared to other extension methods.[6] Since the DWT is a recursive transform, the input to each level must be symmetrically extended. This is achieved by symmetrically extending the input to the first level and using symmetric scaling and wavelet filters at all levels. Under these assumptions, the input to the projection filter may also be symmetrically extended. Figure 8 is a polyphase implementation of the projection and reconstruction blocks in Figure 6. This polyphase implementation allows for fast generation of the complex wavelet coefficients using only real arithmetic. The E operator performs whole-point symmetric extension[6] on the even and odd polyphase components x_e, x_o of the input x. The real and imaginary parts of the projected coefficients s are generated separately, by filtering with the polyphase components $h_e^+, \mathcal{I}m\{h_o^+\}$. If the polyphase components h_e^0, h_o^0 of the lowpass filter h^0 that generates h^+ are odd-length symmetric filters, then $h_e^+, \mathcal{I}m\{h_o^+\}$ are also odd-length symmetric filters and so the inputs to the IIR synthesis filters $\frac{1}{h_e^+}, \mathcal{I}m\{\frac{1}{h_o^+}\}$ will also have whole-point symmetry. This allows us to invoke the symmetric-extension boundary conditions for initialization of the IIR filters prior to recovery of the symetrically extended polyphase components Ex_e and Ex_o. Restriction of these signals[6] then enables the reconstruction of the input $\hat{x} = x_e(z^2) + z^{-1}x_o(z^2)$. Note that $h_e^+, \mathcal{I}m\{h_o^+\}$ cannot be even-length symmetric filters or anti-symmetric filters, because these filters would then necessarily have zeros on the unit circle.[7] The IIR synthesis filters $\frac{1}{h_e^+}, \frac{1}{\mathcal{I}m\{h_o^+\}}$ would then be unstable. Hence we are restricted to using a projection filter generated from a lowpass filter with symmetric, odd-length polyphase components.

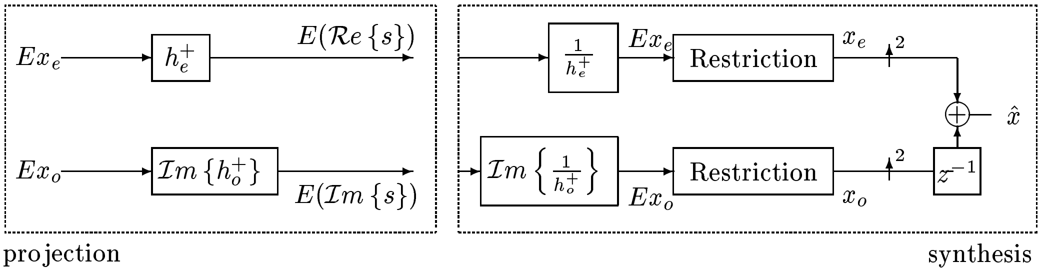

Figure 8. Perfect reconstruction with symmetric extension. The symbols h_e^+ and h_o^+ represent the polyphase components of the projection filter h^+.

Next, we focus on the implementation of the symmetric IIR synthesis filters. This problem has been studied in an image-processing context by Smith et al.[8] and more recently by Unser et al..[9,10] Our approach here is inspired by the fast algorithms proposed in these references. Let $1/H_s(z)$ denote the z transform of a symmetric IIR filter with real-valued coefficients and no poles on the unit circle. These properties imply that the zeros of $H_s(z)$ occur in sets that may have either two or four distinct elements. The two-element sets consist of a real-valued zero and its reciprocal while the four-element sets contain a complex-valued zero, its complex-conjugate and their reciprocals. This enables us to factor $H_s(z)$ as

$$H_s(z) = \left[\prod_{i=1}^{N} H_c(z; z_i) H_a(z; z_i)\right] \left[\prod_{j=1}^{M} G_c(z; z_j) G_a(z; z_j)\right], \tag{2}$$

where

$$\begin{aligned} H_c(z; z_i) &= 1 - z_i z^{-1}, \\ H_a(z; z_i) &= 1 - z_i z, \\ G_c(z; z_i) &= (1 - z_i z^{-1})(1 - z_i^* z^{-1}), \\ G_a(z; z_i) &= (1 - z_i z)(1 - z_i^* z), \end{aligned}$$

and the 'c' ('a') subscripts refer to real-coefficient filters that are implemented causally (anti-causally). Therefore, the symmetric IIR filter $1/H_s(z)$ may be implemented as the cascade of N sections as in Figure 9 and M sections as in Figure 10. Unser et al.[9,10] provide more details on the fast implementation of of these sections as well as the derivation of initial conditions for the IIR filters in them.

We summarize the findings of this section by pointing out that in addition to reducing artifacts at the image boundaries, the symmetric-extension boundary conditions also allow for fast implementation of the IIR synthesis filters. However, to realize these particular boundary conditions at the input to the IIR synthesis filter, the projection filter must have symmetric odd-length polyphase components. This entails the design of an even-length lowpass filter with symmetric odd-length polyphase components. The design process should also guarantee that the symmetric polyphase components $h_e^0(-z), h_o^0(-z)$ have no zeros on the unit circle. This ensures that the IIR synthesis filters may be implemented in a stable manner.

7. SEISMIC ATTRIBUTE ANALYSIS

This section shows how the directional transforms introduced in Sections 3 and 4 provide state-of-the-art performance in a seismic image-processing application. Seismic imagery of the earth's subsurface plays a critical role in all

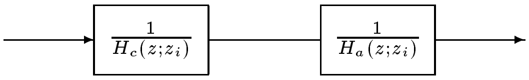

Figure 9. Cascade implementation of two-element set of zeros

aspects of oil and gas exploration and production — from the location of reserves to their appraisal and subsequent monitoring. In oil and gas exploration, seismic cross-sections are scrutinized by interpreters who search for features that indicate possible hydrocarbon reservoirs. Previously, interpreters dealt with large plots of 2D cross sections; they now work on computers with 3D volumes comprising gigabytes of data. Local signal attributes aid the interpretation of seismic data, elucidating its salient characteristics.

A particularly useful attribute is the local angle (dip) of the reflecting surface. Dip representations enable 3D interpretation of *structures* using seismic depth-slices. Channels and faults appear as dip variations, but are often barely visible in amplitude slices. Previously, we used the complex steerable pyramid (CSP[1]) to develop attribute representations that provide very accurate angle indications for 2D cross-sections.[11] However, the overwhelming redundancy of the pyramid representation prohibits its application to 3D seismic volumes because of computational expense.

We now outline a two-stage method for developing local angle representations for 2D seismic images using the novel complex transforms presented in this paper. In the first stage, we apply a projection-based transform to the seismic image. For the non-redundant post-projection transform we use different length filters in the vertical and horizontal directions. We generate h_y^+ from the short, length-4 Daubechies scaling filter since the IIR filters in the corresponding synthesis block may be implemented causally. In the horizontal direction the h_x^+ and h_x^- projection filters are generated from the length-20 Daubechies filters. These longer filters suppress aliasing and improve differentiation between angles by enhancing the distinction between positive and negative lateral frequencies. For the redundant pre-projection transform we have used long filters in both directions as the reconstruction constraints allow this.

In the second stage of the method, we perform six separate reconstructions by considering each angle in isolation. Each reconstructed image reflects the extent to which the seismic cross-section is oriented in the corresponding direction. We estimate the local angle as the weighted average of the six angles.

Figure 11 compares the angle representations generated for a seismic cross-section. The geometry in the section (Figure 11(a)) shows a structure with substantial angle variations and several smaller scale events such as faults. The CSP angle representation in Figure 11(b) provides an accurate and smooth representation of the local angles in the seismic section. Steep uphill dipping reflections are indicated with high intensity, steeply downward dipping reflections are represented in dark gray-value. The non-redundant post-projection transform generates a noisier and less smooth representation with a somewhat finer resolution (Figure 11(c)). This transform however requires *14 times* less redundancy. The result of the angle detection with the redundant pre-projection transform in Figure 11 (d) looks almost equally informative as the result with CSP (b), but has a higher spatial resolution and is only 4 times redundant.

8. CONCLUSION

This paper demonstrates persuasively that the projection technique improves directional selectivity. Both our methods partition highpass coefficients into six directional subbands at each scale, and lowpass coefficients into two directional subbands. Although our transform uses shorter filters than Kingsbury's dual-tree transform, it achieves

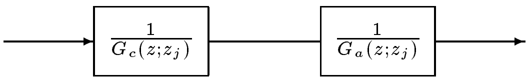

Figure 10. Cascade implementation of four-element set of zeros

tighter localization of features in each of the subbands. The novel synthesis filter bank structure is an important contribution, since it guarantees perfect reconstruction in a non-redundant framework. In conclusion, we stress the flexibility and elegance of our technique, since it enhances *any* DWT with superior directional selectivity.

REFERENCES

1. E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable multi-scale transforms," *IEEE Trans. Inform. Theory* **38**, pp. 587–607, Mar. 1992.
2. R. H. Bamberger and M. J. T. Smith, "A filter bank for the directional decomposition of images: Theory and design," *IEEE Transactions on Signal Processing* **40**, pp. 882–893, Apr. 1992.
3. N. Kingsbury, "Image processing with complex wavelets," *Phil. Trans. R. Soc. Lond.* , 1997 (Submitted).
4. R. van Spaendonck, F. C. A. Fernandes, M. Coates, and C. S. Burrus, "Non-redundant, directionally selective, complex wavelets." To appear at IEEE Proc. Int. Conf. Image Processing, 2000.
5. R. L. C. van Spaendonck, F. M. Hindriks, F. C. A. Fernandes, and G. G. Drijkoningen, "Three-dimensional attributes for seismic interpretation." To appear at the 70th Annual meeting SEG (Society of Exploration Geophysicists) Calgary, Canada, August 2000.
6. G. Strang and T. Nguyen, *Wavelets and Filter Banks*, Wellesley-Cambridge Press, Wellesley, MA, 1996.
7. T. W. Parks and C. S. Burrus, *Digital Filter Design*, John Wiley & Sons. Inc., New York, NY, 1st ed., 1987.
8. M. J. T. Smith and S. L. Eddins, "Analysis/synthesis techniques for subband image coding," *IEEE Transactions on Acoustics, Speech, and Signal Processing* **38**, pp. 1446–1456, Aug. 1990.
9. M. Unser, A. Aldroubi, and M. Eden, "Fast b-spline transforms for continuous image representation and interpolation," *IEEE Trans. Patt. Anal. and Machine Intell.* **13**, pp. 277–285, Mar. 1991.
10. M. Unser, A. Aldroubi, and M. Eden, "B-spline signal processing: Part II - efficient design and applications," *IEEE Transactions on Signal Processing* **41**, pp. 834–848, Feb. 1993.
11. R. L. C. van Spaendonck and R. G. Baraniuk, "Directional scale analysis for seismic interpretation," in *69th Annual meeting SEG (Society of Exploration Geophysicists)*, 1999. Houston, Texas.

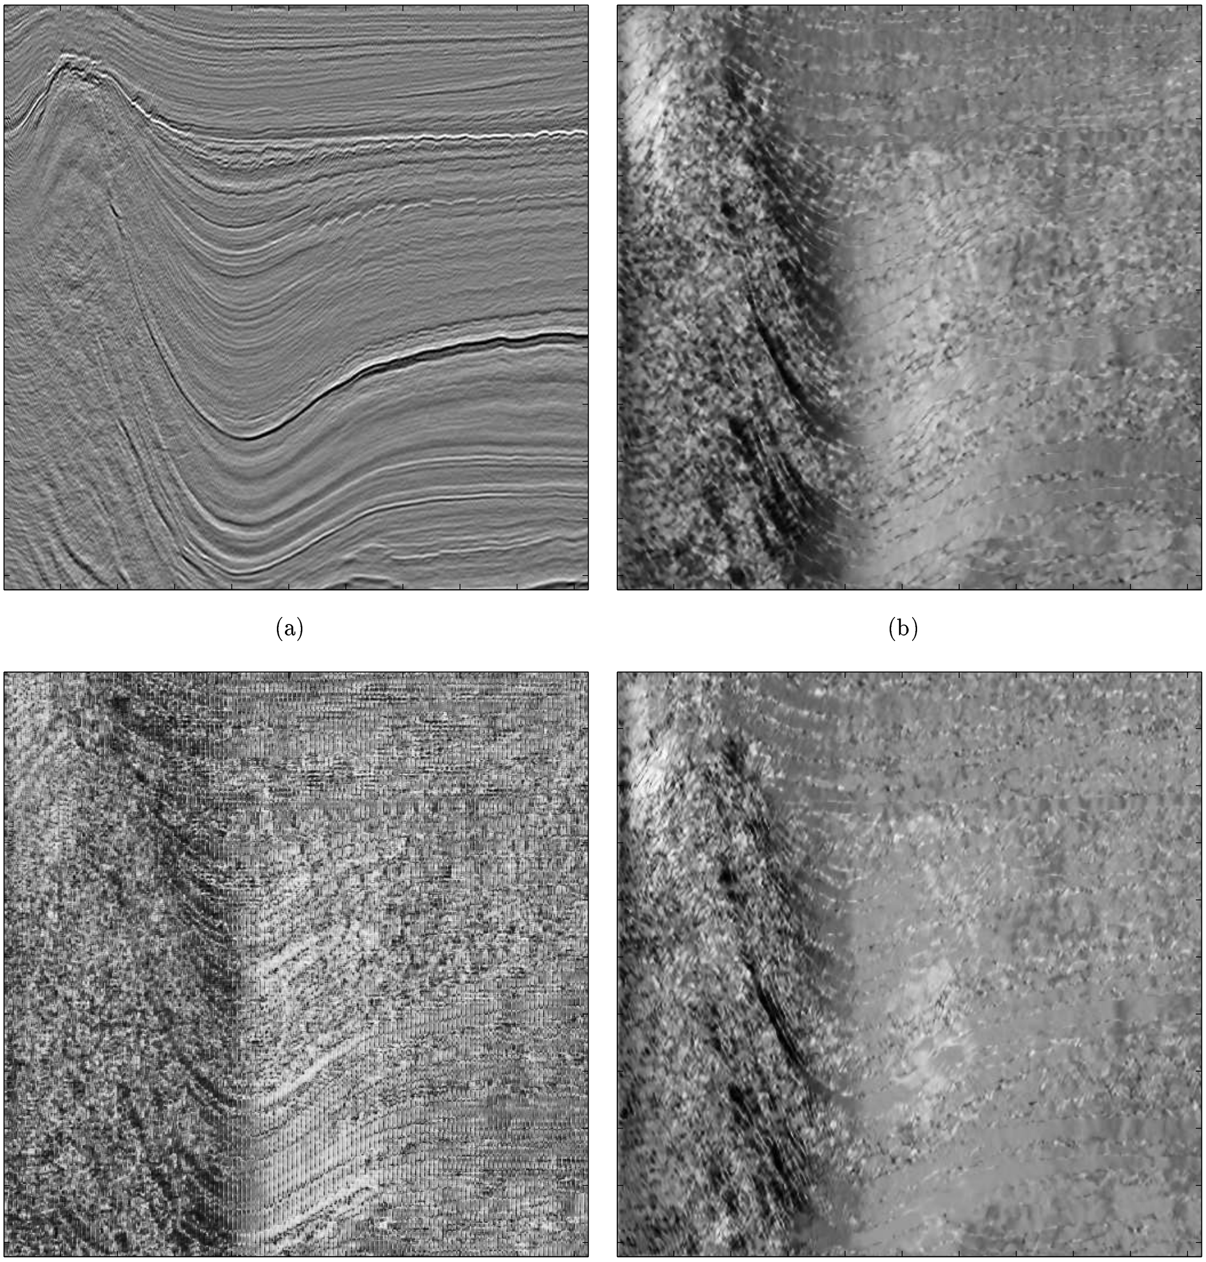

(a) (b)

(c) (d)

Figure 11. (a) Seismic section showing substantial angle variations (b) Angle analysis results using complex steerable pyramid. (c) Angle analysis results using the non-redundant post-projection transform. (d) Angle analysis using the redundant pre-projection transform. In (b),(c) and (d) intensity represents angle, black (-90 degrees) through white (+90 degrees).