SHIFT INVARIANT PROPERTIES OF THE DUAL-TREE COMPLEX WAVELET TRANSFORM

Nick Kingsbury

Signal Processing Group, Dept. of Engineering, University of Cambridge,
Cambridge CB2 1PZ, UK. E-mail: ngk@eng.cam.ac.uk

ABSTRACT

We discuss the shift invariant properties of a new implementation of the Discrete Wavelet Transform, which employs a *dual tree* of wavelet filters to obtain the real and imaginary parts of *complex* wavelet coefficients. This introduces limited redundancy (2^m:1 for m-dimensional signals) and allows the transform to provide approximate shift invariance and directionally selective filters (properties lacking in the traditional wavelet transform) while preserving the usual properties of perfect reconstruction and computational efficiency with good well-balanced frequency responses.

1. INTRODUCTION

The Discrete Wavelet Transform (DWT) in its maximally decimated form (Mallat's dyadic filter tree [1]) has established an impressive reputation as a tool for signal compression, but its use for other signal analysis and reconstruction tasks has been hampered by two main disadvantages:

- Lack of *shift invariance*, which means that small shifts in the input signal can cause major variations in the distribution of energy between DWT coefficients at different scales.
- Poor *directional selectivity* for diagonal features, because the wavelet filters are separable and real.

A well-known way of providing shift invariance is to use the undecimated form of the dyadic filter tree, but this suffers from increased computation requirements and high redundancy in the output information, making subsequent processing expensive too.

In [4, 5], we introduced the Dual-Tree Complex Wavelet Transform (DT CWT) with the following properties:

- Approximate **shift invariance**;
- Good **directional selectivity** in 2-dimensions (2-D) with Gabor-like filters (also true for higher dimensionality, m-D);
- **Perfect reconstruction** (PR) using short linear-phase filters;
- **Limited redundancy**, independent of the number of scales, $2:1$ for 1-D ($2^m:1$ for m-D);
- **Efficient order-N computation** – only twice the simple DWT for 1-D (2^m times for m-D).

2. THE DUAL FILTER TREE

Our work with complex wavelets for motion estimation [2] showed that complex wavelets could provide approximate shift invariance. Unfortunately we were unable to obtain PR and good frequency characteristics using short support complex FIR filters in a single tree (eg. fig. 1 Tree a).

However we can achieve approximate shift invariance with a *real* DWT by doubling the sampling rate at each level of the tree. For this to work, the samples must be evenly spaced. We can double all the sampling rates in Tree a of fig. 1 by eliminating the down-sampling by 2 after the level 1 filters, H_{0a} and H_{1a}. This is equivalent to two parallel fully-decimated trees, a and b, provided that the delays of H_{0b} and H_{1b} are one sample offset from H_{0a} and H_{1a}. We then find that, to get uniform intervals between samples from the two trees *below* level 1, the filters in one tree must provide delays that are half a sample different (at each filter input rate) from those in the other tree. For linear phase, this requires *odd-length* filters in one tree and *even-length* filters in the other. Greater symmetry between the two trees occurs if each tree uses odd and even filters alternately from level to level, but this is not essential. In fig. 2 we show the positions of the output samples when the filters are odd and even as in fig. 1. To invert the transform, we apply the PR filters G in the usual way to invert each tree separately and finally we average the two results.

In order to show the shift invariant properties of the dual tree, we shall consider what happens when we choose to retain the coefficients of just one type (wavelet or scaling function) from just one level of the dual tree. For example we might choose to retain only the level-3 wavelet coefficients x_{001a} and x_{001b} from

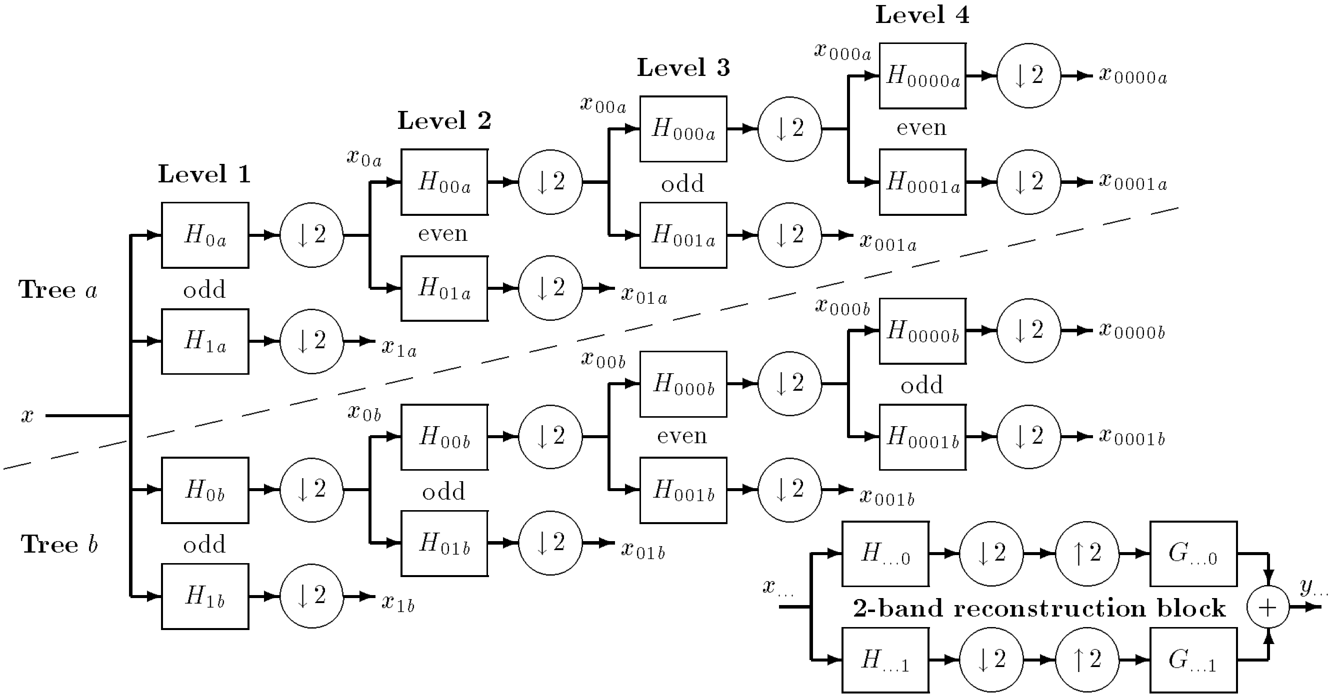

Figure 1: Dual tree of real filters for the CWT, giving real and imaginary parts of complex coefficients.

fig. 1, and set all others to zero. If the signal y, reconstructed from just these coefficients, is free of aliasing then the transform is shift invariant at that level.

Input samples	← Block of 16 input samples →															
x:	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Level 1 samples																
odd Lo x_{0a}:	a		a		a		a		a		a		a		a	
odd Lo x_{0b}:		b		b		b		b		b		b		b		b
odd Hi x_{1a}:		a		a		a		a		a		a		a		a
odd Hi x_{1b}:	b		b		b		b		b		b		b		b	
Level 2 samples																
even Lo x_{00a}:		a				a				a				a		
odd Lo x_{00b}:				b				b				b				b
Hi x_{01a}, x_{01b}:		*				*				*				*		
Level 3 samples																
odd Lo x_{000a}:		a								a						
even Lo x_{000b}:						b								b		
Hi x_{001a}, x_{001b}:						*								*		
Level 4 samples																
even Lo x_{0000a}:						a										
odd Lo x_{0000b}:														b		
Hi x_{0001a}, x_{0001b}:						*										

Figure 2: Effective sampling points of odd and even filters in fig. 1 assuming zero phase responses.

Fig. 3 shows the simplified analysis and reconstruction parts of the dual tree when coefficients of just one type and level are retained. All down(up)-sampling operations are moved to the output (input) of the analysis (reconstruction) filter banks and the cascaded filter transfer functions are combined. $M = 2^m$ is the total downsampling factor. For example if we retain only x_{001a} and x_{001b}, then $M = 8$, $A(z) = H_{0a}(z)\, H_{00a}(z^2)\, H_{001a}(z^4)$ and $B(z)$, $C(z)$, $D(z)$ are obtained similarly.

3. SHIFT INVARIANT FILTER DESIGN

Letting $W = e^{j2\pi/M}$, multi-rate analysis of fig. 3 gives:

$$Y(z) = \frac{1}{M} \sum_{k=0}^{M-1} X(W^k z)[A(W^k z)\, C(z) + B(W^k z)\, D(z)] \tag{1}$$

For shift invariance, the aliasing terms (for which $k \neq 0$) must be negligible. So we design $B(W^k z)\, D(z)$ to cancel out $A(W^k z)\, C(z)$ for all non-zero k which give overlap of the pass or transition bands of the filters $C(z)$ or $D(z)$ with those of the shifted filters $A(W^k z)$ or $B(W^k z)$. Separate strategies are needed depending on whether the filters are lowpass (for scaling functions) or bandpass (for wavelets).

For level m in the dual tree, the lowpass filters have passbands $-f_s/2M \to f_s/2M$ (f_s is the input sampling frequency). The W^k terms in (1) shift the passbands in multiples of f_s/M. If $A(z)$ and $C(z)$ have similar frequency responses (as required for near-orthogonal filter sets) and significant transition bands, it is not possible to make $A(Wz)\, C(z)$ small at all frequencies $z = e^{j\theta}$, whereas we can quite easily make $A(W^2 z)\, C(z)$ small since the frequency shift is twice as great. Hence for the lowpass case, we design $B(W^k z)\, D(z)$ to cancel

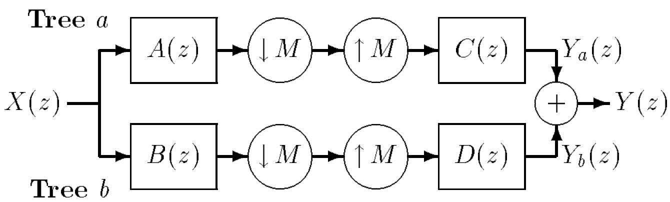

Figure 3: Basic configuration of the dual tree if either wavelet or scaling-function coefficients from just level m are retained ($M = 2^m$).

$A(W^k z)\,C(z)$ when k is odd by letting:

$$B(z) = z^{\pm M/2} A(z) \quad \text{and} \quad D(z) = z^{\mp M/2} C(z) \qquad (2)$$

so that $B(W^k z)\,D(z) = (-1)^k A(W^k z)\,C(z)$.

Now consider the bandpass case. Here we find that the edges of the positive frequency passband of C or D, $f_s/2M \rightarrow f_s/M$, will tend to overlap with the edges of the negative frequency passband of A or B, that gets shifted either to $0 \rightarrow f_s/2M$ or to $f_s/M \rightarrow 3f_s/2M$ when $k = 1$ or 2 respectively. Similarly for the opposite passbands when $k = -1$ or -2. Since the aliasing terms are always caused by the overlap of opposing frequency passbands, whereas the wanted terms ($k = 0$) are produced by overlap of same-frequency passbands, the solution here is to give B and D positive and negative passbands of opposite polarity while A and C have passbands of the same polarity (or vice versa). Suppose we have prototype *complex* filters $P(z)$ and $Q(z)$, each with just a single passband $f_s/2M \rightarrow f_s/M$ and zero gain at all negative frequencies, then we let:

$$\begin{aligned} A(z) &= \Re[2P(z)] = P(z) + P^*(z) \\ B(z) &= \Im[2P(z)] = -j[P(z) - P^*(z)] \\ C(z) &= \Re[2Q(z)] = Q(z) + Q^*(z) \\ D(z) &= \Im[-2Q(z)] = j[Q(z) - Q^*(z)] \qquad (3) \end{aligned}$$

where conjugation is given by $P^*(z) = \sum_r p_r^* z^{-r}$ and produces negative frequency passbands. The overlap terms are of the form $Q(z)\,P^*(W^k z)$ for $k = 1, 2$ and $Q^*(z)\,P(W^k z)$ for $k = -1, -2$ which all cancel when $B(W^k z)\,D(z)$ is added to $A(W^k z)\,C(z)$ in (1):

$$\begin{aligned} & A(W^k z)\,C(z) + B(W^k z)\,D(z) \\ &= 2P(W^k z)\,Q(z) + 2P^*(W^k z)\,Q^*(z) \qquad (4) \end{aligned}$$

Hence we now need only design the filters such that the positive frequency complex filter $Q(z)$ does not overlap with shifted versions of the similar filter $P(z)$, which is quite easy since the filter bandwidths are only $f_s/2M$ while the shifts are in multiples of f_s/M. For octave band filters in which the upper transition band is twice as wide as the lower transition band, this implies that the pass and transition bands should lie within the frequency range $f_s/3M \rightarrow 4f_s/3M$. The formulations in equations (3) show that the highpass filter outputs from trees a and b should be regarded as the *real and imaginary parts of complex processes.* We may also regard the pairs of lowpass outputs in this way.

In practice, the filters will not have zero gain in their stop bands and the above relationships will be approximate. So the transform will only be *approximately* shift invariant. However good performance is possible with quite low complexity filters.

For the lowpass filters, equation (2) implies that the tree b samples should interpolate midway between the tree a samples, effectively doubling the sampling rate, as shown in fig 2. This may achieved by two identical lowpass filters (either odd or even) at level 1, offset by 1 sample delay, and then by pairs of odd and even length filters at further levels to achieve the extra delay difference of $M/4$ samples, to make the total difference $M/2$ at each level.

The responses of A and B also need to match, which can only be achieved approximately beyond level 1. We do this by designing $H_{00a}(z^2)$ to give minimum mean squared error in the approximation $z^{\pm 2} H_{0a}(z)\,H_{00a}(z^2) \approx H_{0b}(z)\,H_{00b}(z^2)$. Then H_{01a} can be designed to form a perfect reconstruction set with H_{00a} such that the reconstruction filters G_{00a} and G_{00b} also match each other closely.

Finally the symmetry of the odd-length highpass filters and the anti-symmetry of the even-length highpass filters produce the required phase relationships between the positive and negative frequency passbands, and equations (3) are approximately satisfied too.

These filters can then be used for all subsequent levels of the transform. Good shift invariance (and wavelet smoothness) requires that frequency response sidelobes of the cascaded multirate filters should be small. This is achieved if each lowpass filter has a stopband covering $\frac{1}{3}$ to $\frac{2}{3}$ of its sample rate, so as to reject the image frequencies due to subsampling in the next lowpass stage. If the highpass filters then mirror this characteristic, the conditions for no overlap of the shifted bandpass responses in (4) are automatically satisfied.

As an example, we selected two linear-phase PR biorthogonal filter sets which meet the above conditions quite well and are also nearly orthogonal. For the odd-length set, we designed (13,19)-tap filters using the (1-D) transformation of variables method [3], and then a (12,16)-tap even-length set to match. Fig. 4 shows the frequency responses of the reconstruction filter bank; and the analysis filters are very similar. The analysis coefficients are listed in the following table (the reconstruction filters are obtained by negating alternate coefficients and swapping bands).

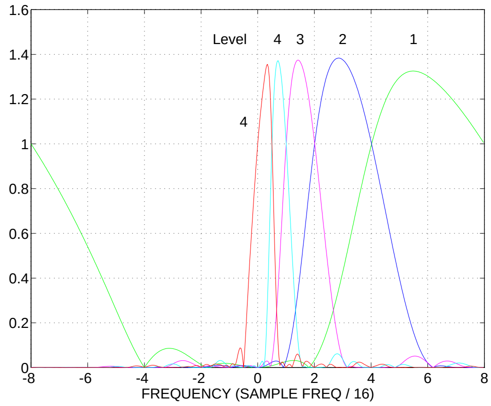

Figure 4: Frequency reponses of complex wavelets at levels 1 to 4 and of the level 4 scaling function.

odd $H_{\dots 0}$ 13-tap	odd $H_{\dots 1}$ 19-tap	even $H_{\dots 0}$ 12-tap	even $H_{\dots 1}$ 16-tap
	-0.0000706		
	0		-0.0004645
-0.0017581	0.0013419		0.0013349
0	-0.0018834	-0.0058109	0.0022006
0.0222656	-0.0071568	0.0166977	-0.0130127
-0.0468750	0.0238560	-0.0000641	0.0015360
-0.0482422	0.0556431	-0.0834914	0.0869008
0.2968750	-0.0516881	0.0919537	0.0833552
0.5554688	-0.2997576	0.4807151	-0.4885957
0.2968750	0.5594308	0.4807151	0.4885957
-0.0482422	-0.2997576	0.0919537	-0.0833552
⋮	⋮	⋮	⋮

Fig. 5 demonstrates the shift invariance of the DT CWT with these filters. The input is a unit step, shifted to 16 adjacent sampling instants in turn. Fig. 5a shows the input steps and the components of the DT CWT output, reconstructed from the wavelet coefficients at each of levels 1 to 4 in turn and from the scaling function coefficients at level 4. Summing these components reconstructs the input steps perfectly. For comparison fig. 5b shows the equivalent components if the real DWT is used. The CWT responses are clearly much more consistent with shift (shift invariant). The energies of the DT CWT coefficients at each level vary over the 16 shifts by no more than 1.025 : 1, whereas the DWT coefficient energies vary by up to 5.45 : 1 !

4. EXTENSION TO M-DIMENSIONS

Extension to 2-D is achieved by separable filtering along columns and then rows. However, if column and row filters both suppress negative frequencies, then only the first quadrant of the 2-D signal spectrum is retained. Two adjacent quadrants of the spectrum are required to represent fully a real 2-D signal, so we also filter with complex conjugates of the row filters. This gives 4 : 1 redundancy in the transformed 2-D signal.

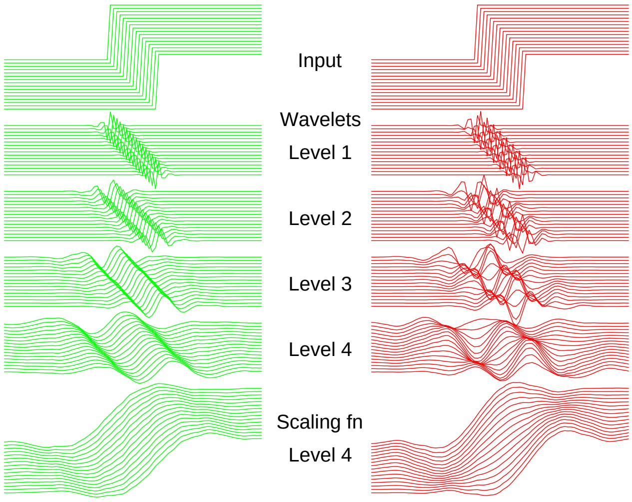

(a) Dual Tree CWT

(b) Real DWT

Figure 5: Wavelet and scaling function components at levels 1 to 4 of 16 shifted step responses of the DT CWT (a) and real DWT (b).

If the signal exists in more than 2-D, then further conjugate pairs of filters are needed for each dimension leading to redundancy of $2^m : 1$.

Complex filters in multiple dimensions provide true directional selectivity, despite being implemented separably, because they are still able to separate all parts of the m-D frequency space. For example a 2-D CWT produces six bandpass subimages of complex coefficients at each level, which are strongly oriented at angles of $\pm 15°, \pm 45°, \pm 75°$. We believe this is an important feature for many applications, including motion estimation and compensation, texture synthesis, image denoising, edge enhancement, segmentation, and image classification. Some of these are discussed in [4, 5].

5. REFERENCES

[1] S G Mallat, "A theory for multiresolution signal decomposition: The wavelet representation", *IEEE Trans. PAMI*, 11(7), pp 674-693, 1989.

[2] J F A Magarey and N G Kingsbury, "Motion estimation using a complex-valued wavelet transform", *IEEE Trans. Signal Proc.*, 46(4), April 1998.

[3] D B H Tay and N G Kingsbury, "Flexible design of multidimensional perfect reconstruction FIR 2-band filters using transformations of variables", *IEEE Trans. Image Proc.*, 2(4), pp 466-480, 1993.

[4] N G Kingsbury, "The dual-tree complex wavelet transform: a new technique for shift invariance and directional filters", *Proc. 8th IEEE DSP Workshop*, Bryce Canyon, Aug 1998.

[5] N G Kingsbury: "The dual-tree complex wavelet transform: a new efficient tool for image restoration and enhancement", *Proc. EUSIPCO 98*, Rhodes, Sept 1998.