Complex-valued Wavelets, the Dual Tree, and Hilbert Pairs: why these lead to Shift Invariance and Directional M-D Wavelets?

Nick Kingsbury

Signal Processing & Communications Group, Dept. of Engineering
University of Cambridge, Cambridge CB2 1PZ, UK.
ngk@eng.cam.ac.uk www.eng.cam.ac.uk/~ngk

Meeting on Mathematics & Information in Cambridge – February 2012

UNIVERSITY OF CAMBRIDGE

Important Continuous-time Complex-valued Wavelets

- The **Gabor** function:

$$\psi(t) = k\; e^{-t^2/2\sigma^2}\; e^{i\omega_0 t} \qquad\qquad \hat{\psi}(\omega) = K\; e^{-\sigma^2(\omega-\omega_0)^2/2}$$

- The **Morlet** wavelet:

$$\psi(t) = k\; e^{-t^2/2\sigma^2}(e^{i\omega_0 t} - \kappa_0) \qquad\qquad \hat{\psi}(\omega) = K\left(e^{-\sigma^2(\omega-\omega_0)^2/2} - \kappa_0\; e^{-\sigma^2\omega^2/2}\right)$$

- The **Cauchy** wavelet:

$$\psi(t) = k\,(1 - i\beta t)^{-\alpha} \qquad\qquad \hat{\psi}(\omega) = \begin{cases} K\,\omega^{\alpha-1}e^{-\omega/\beta}, & \omega \geq 0 \\ 0\,, & \omega < 0 \end{cases}$$

Max gain is at $\omega_0 = \beta(\alpha - 1)$, and typically for one octave half-power bandwidth, $\alpha \approx 8$.

Plots of $\psi(t)$ and $|\hat{\psi}(\omega)|$ for Morlet, Cauchy (continuous) and dual-tree (discrete) complex wavelets

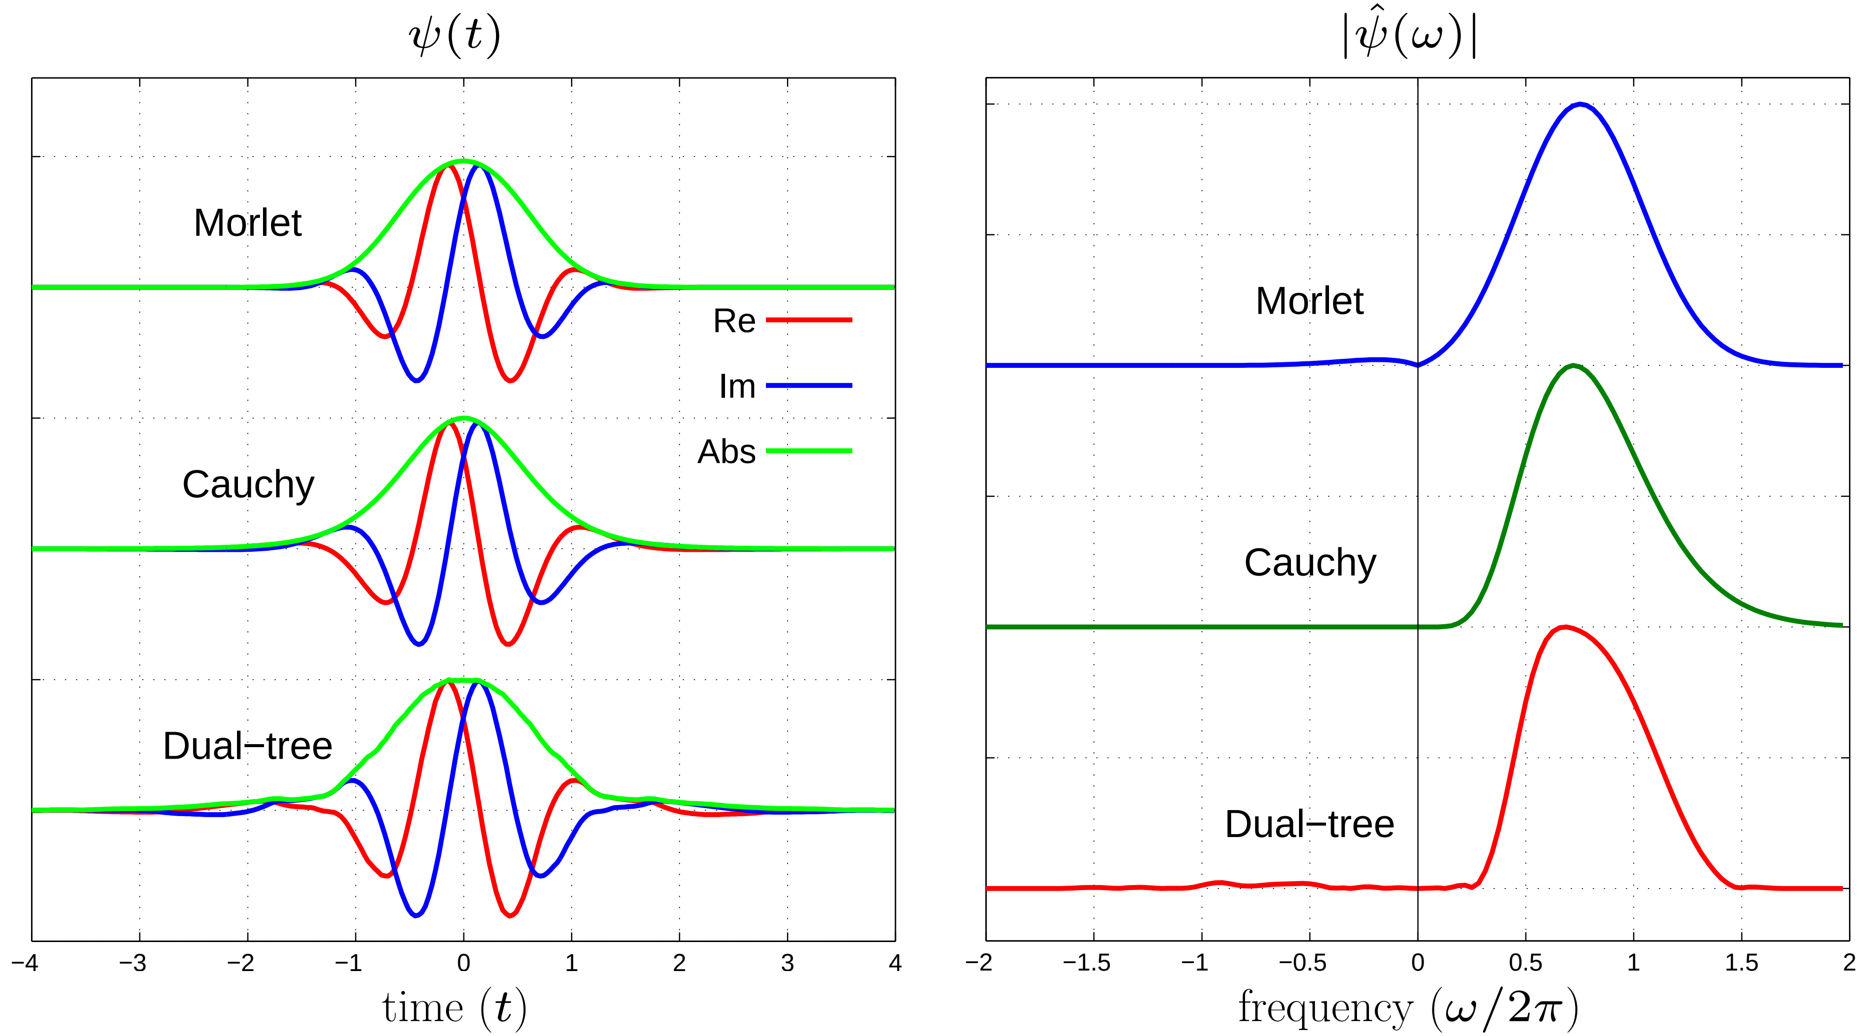

Note: $k = (1+i)/\sqrt{2}$; and the dual-tree filters used here are 18-tap Q-shift filters.

COMPLEX-VALUED WAVELETS, THE DUAL TREE, AND HILBERT PAIRS

How can we produce *good* discrete wavelet transforms?

- What are the problems with real-valued wavelet bases?
- Why do we need the Dual Tree?
- What is the Hilbert Pair delay condition?
- Why does this give shift invariance?
- Why do we use Q-shift filters?
- How do we extend the dual-tree to multi-dimensions?
- Why do we get good directional filters in m-D?
- What are some applications of the DT $\mathbb{C}$WT?

Features of the (Real) Discrete Wavelet Transform (DWT)

- **Good compression** of signal energy.
- **Perfect reconstruction** with short support filters.
- **No redundancy**.
- **Very low computation** – order-N only.

But what are the problems of the DWT?

- **Severe shift dependence** (due to aliasing in down-samplers).
- **Poor directional selectivity** in 2-D, 3-D etc. (due to separable real filters).

The DWT is normally implemented with a tree of highpass and lowpass filters, separated by $2:1$ decimators.

Real Discrete Wavelet Transform (DWT) in 1-D

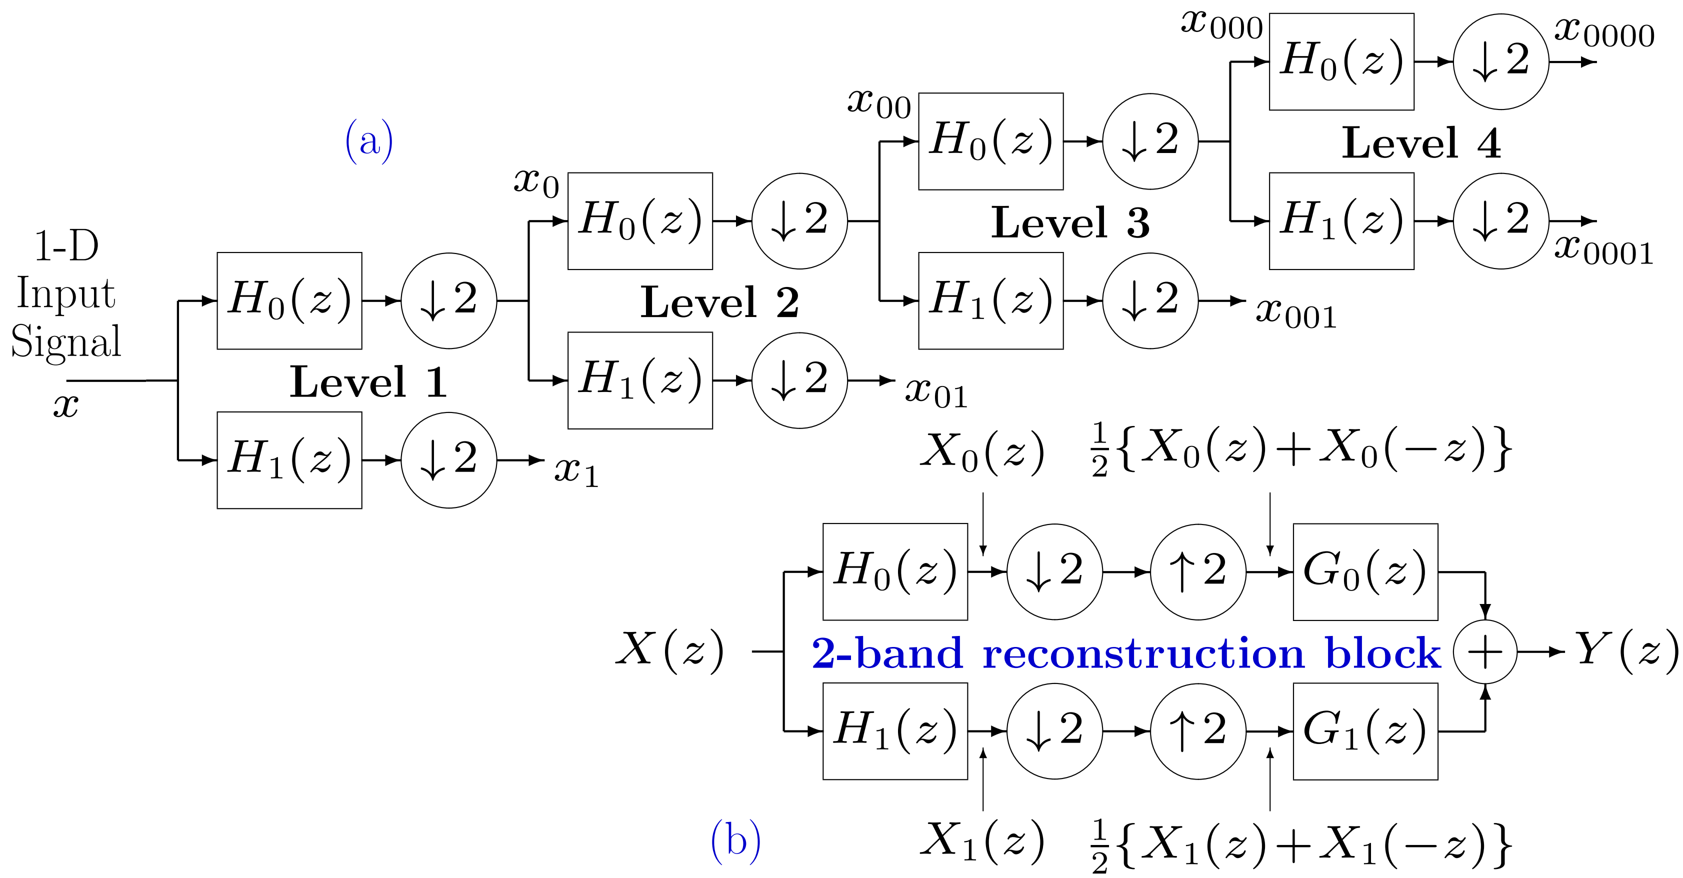

Figure 1: (a) Tree of real filters for the DWT. (b) Reconstruction filter block for 2 bands at a time, used in the inverse transform.

Shift Invariance of Complex DT ℂWT vs Real DWT

(a)

DT ℂWT

(18-tap filter)

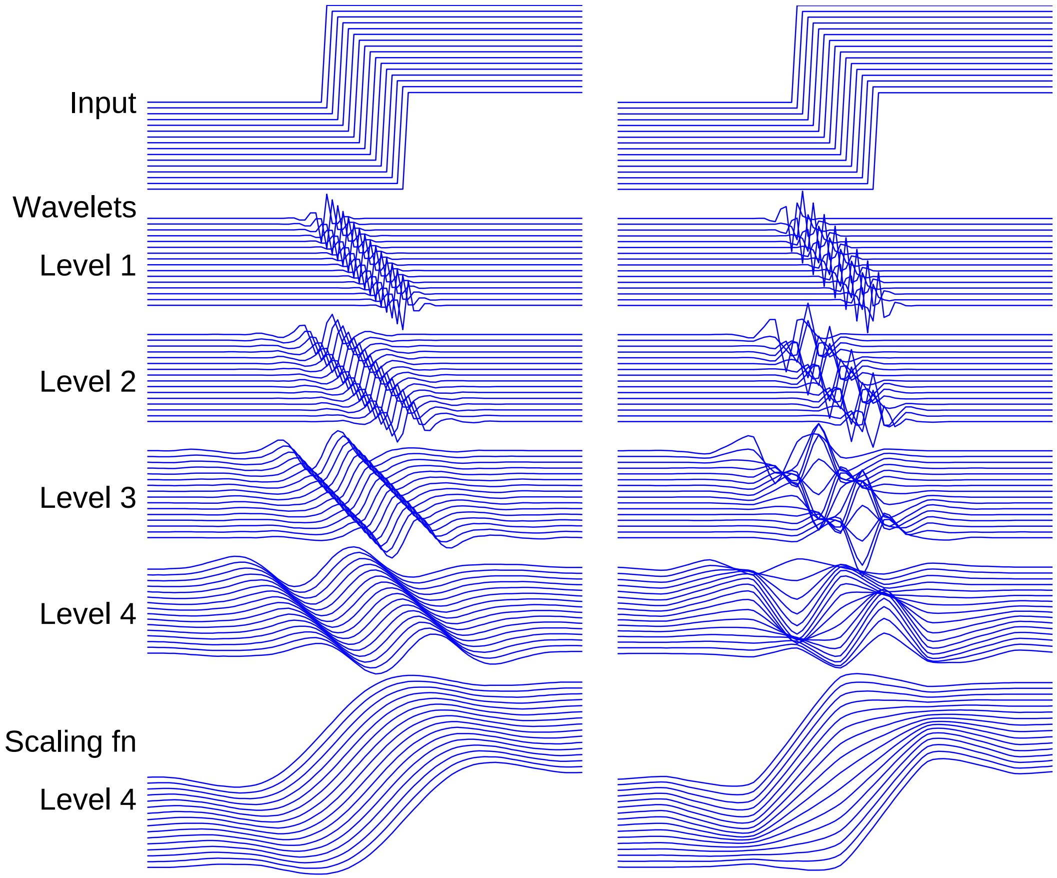

(b)

Real DWT

(13,19-tap filter)

Figure 2: 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 there is good shift invariance, all components at a given level should be similar in shape, as in (a).

2-D Shift Invariance of Complex DT ℂWT vs Real DWT

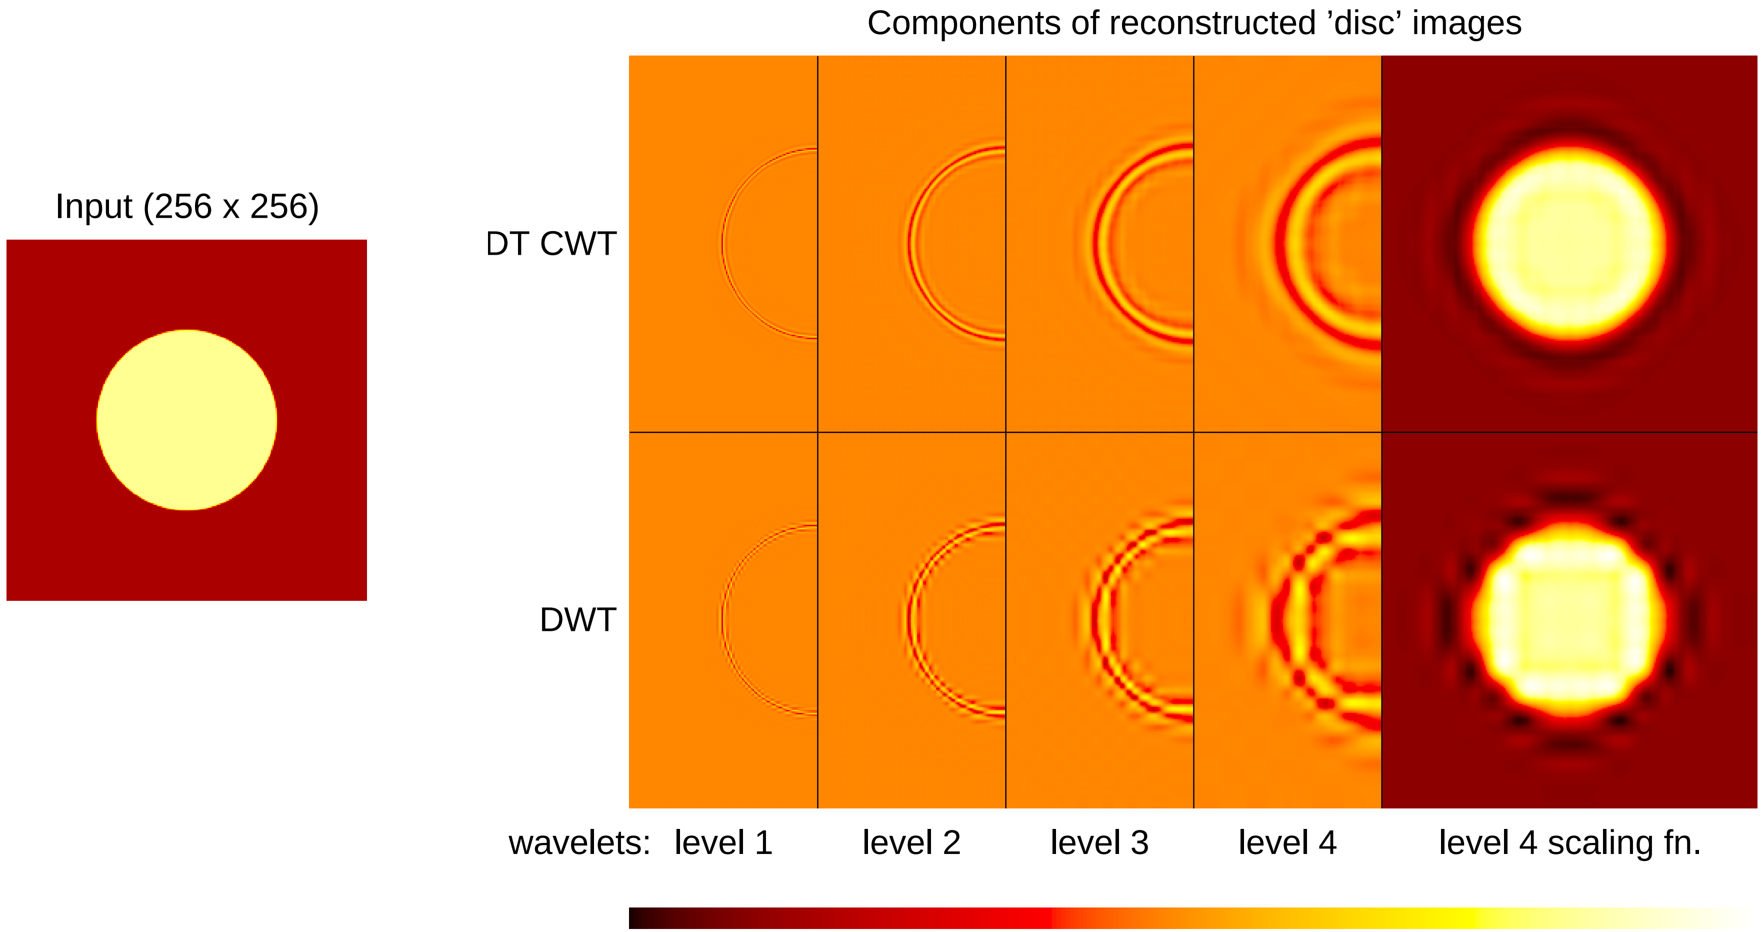

Figure 3: Wavelet and scaling function components at levels 1 to 4 of an image of a light circular disc on a dark background, using the 2-D DT ℂWT (upper row) and 2-D DWT (lower row). Only half of each wavelet image is shown in order to save space.

Why do we need the Dual Tree?

Making the wavelet responses **analytic** is a good way to halve their bandwidth and hence minimise aliasing.

BUT we cannot use complex filters in Fig 1 to obtain analyticity and perfect reconstruction together, because of conflicting requirements in Fig 1b – analytic filters must suppress negative frequencies, while perfect reconstruction requires a flat overall frequency response.

So we use the **Dual Tree**:

- to create the **real** and **imaginary** parts of the analytic wavelets separately, using 2 trees of **purely real** filters;
- to efficiently synthesise a multiscale **shift-invariant** filterbank, with perfect reconstruction and **only 2:1 redundancy** (and computation);
- to produce complex coefficients whose **amplitude varies slowly** and whose **phase shift** depends approximately **linearly** on displacement;

Q-SHIFT DUAL TREE COMPLEX WAVELET TRANSFORM IN 1-D

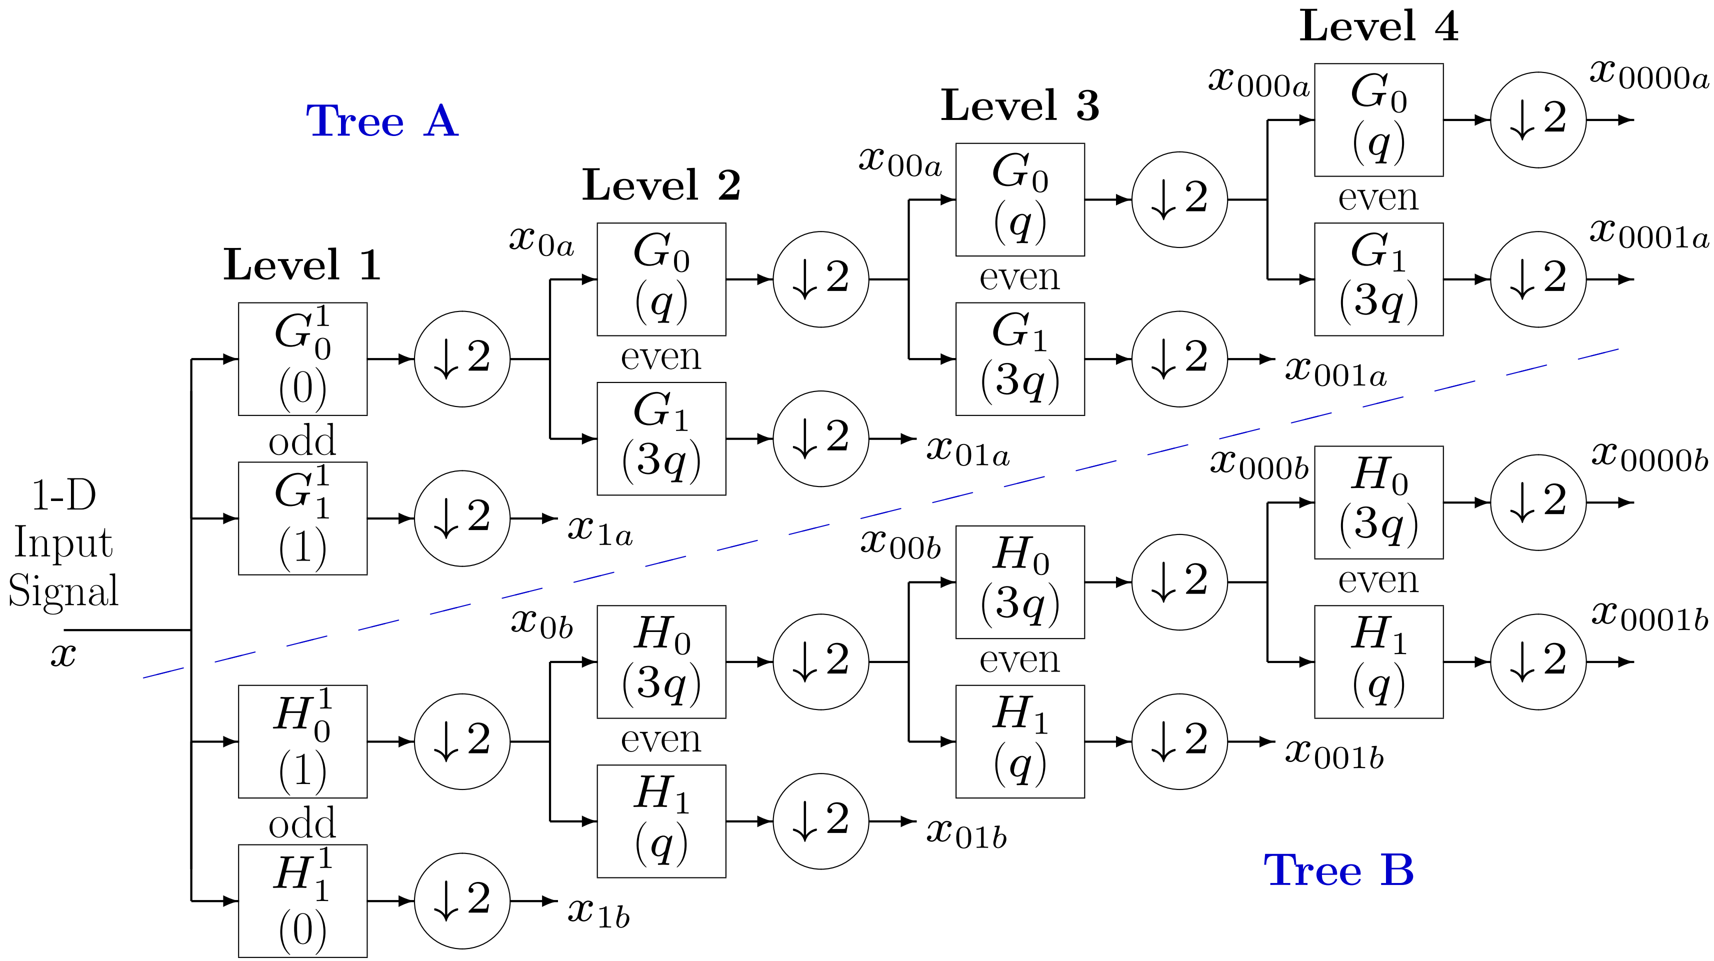

Figure 4: Dual tree of real filters for the Q-shift $\mathbb{C}$WT, giving real and imaginary parts of complex coefficients from tree A and tree B respectively. Figures in brackets indicate the approximate delay for each filter, where $q = \frac{1}{4}$ sample period. Special level 1 filters, G^1 and H^1, allow for the finite number of levels.

Q-shift DT ℂWT Basis Functions – Levels 1 to 3

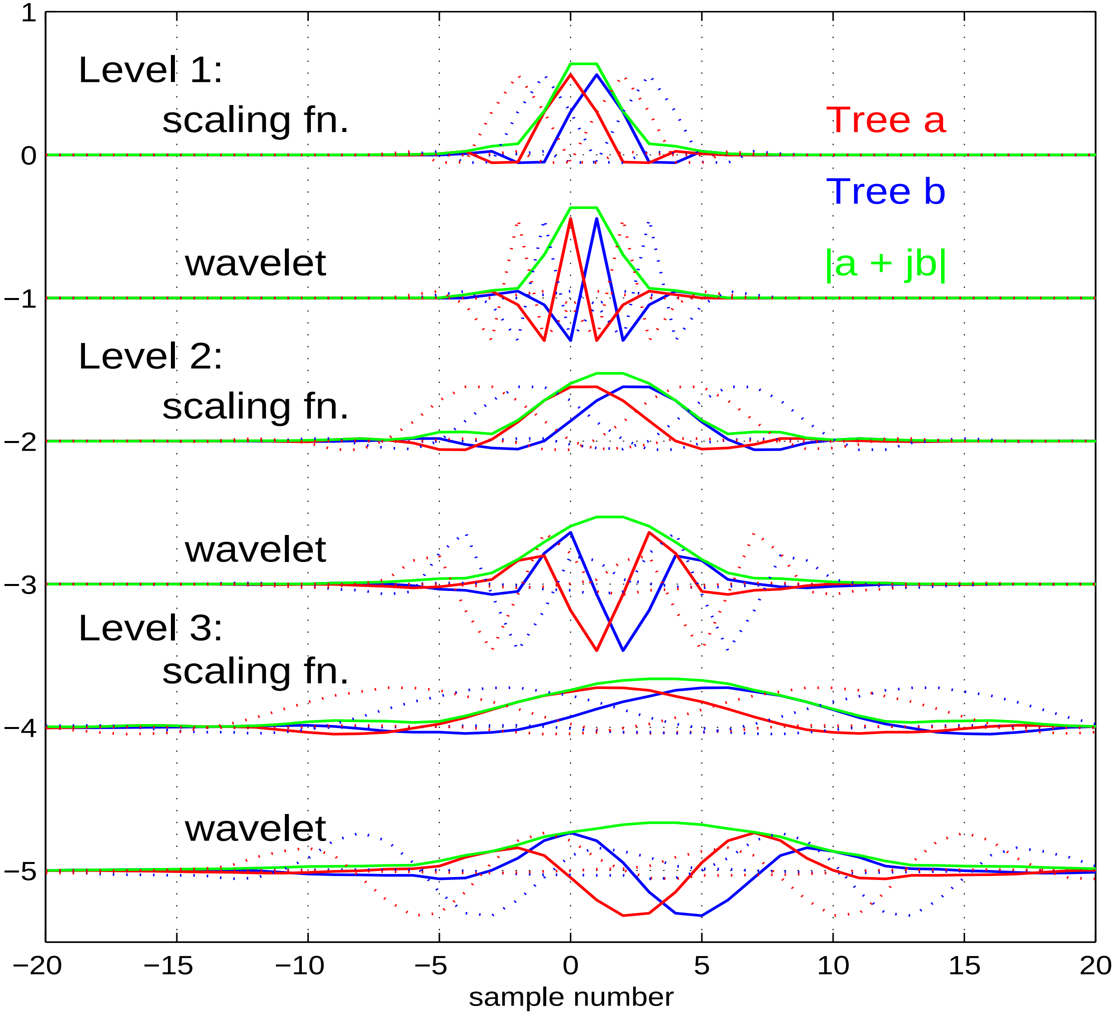

Figure 5: Basis functions for adjacent sampling points are shown dotted.

WHAT IS THE HILBERT PAIR DELAY CONDITION ?

- **Given two parallel orthonormal discrete wavelet transforms, what is the constraint on the lowpass filters in each transform, such that the resulting continuous wavelets from each transform form a Hilbert Pair?** (This question and its answer are due to Ivan Selesnick in Signal Proc. Letters, June 2001.)

- A pair of wavelets, $\psi_g(t)$ and $\psi_h(t)$, are a **Hilbert Pair** if the complex function $\psi_g(t) + i\,\psi_h(t)$ is **analytic** (i.e. its Fourier transform is zero for $\omega < 0$).

- We shall show that this requires the lowpass filters, $g_0(n)$ and $h_0(n)$, of the two transforms to be related by the **half-sample delay condition**, expressed in the frequency domain as

$$H_0(\omega) = e^{-i\omega/2} G_0(\omega)$$

2-scale condition on the Tree A filters of a dyadic wavelet transform

$$\textbf{Scaling function:} \quad \phi_g(t) = 2\sum_n g_0(n)\, \phi_g(2t-n) \tag{1}$$

$$\textbf{Mother wavelet:} \quad \psi_g(t) = 2\sum_n g_1(n)\, \phi_g(2t-n) \tag{2}$$

Taking the Fourier transform of (1) gives the frequency domain relationship

$$\begin{aligned}
\hat{\phi}_g(\omega) &= \int_{-\infty}^{\infty} 2\sum_n g_0(n)\, \phi_g(2t-n)\, e^{-i\omega t}\, dt \\
&= \int_{-\infty}^{\infty} \sum_n g_0(n)\, \phi_g(u)\, e^{-i\omega u/2}\, e^{-i\omega n/2}\, du \quad \text{where } u = 2t-n \\
&= \sum_n g_0(n)\, e^{-in\omega/2} \;.\; \int_{-\infty}^{\infty} \phi_g(u)\, e^{-iu\omega/2}\, du \\
&= G_0(\tfrac{\omega}{2}) \,.\, \hat{\phi}_g(\tfrac{\omega}{2}) \qquad (3)
\end{aligned}$$

Infinite Product formulae

Iterating on (3):

$$\hat{\phi}_g(\omega) = G_0(\tfrac{\omega}{2})\ G_0(\tfrac{\omega}{4})\ \hat{\phi}_g(\tfrac{\omega}{4}) = \cdots = \left[\prod_{k=1}^{\infty} G_0(2^{-k}\omega)\right] \hat{\phi}_g(0) \tag{4}$$

Similarly, from (2) and (4):

$$\hat{\psi}_g(\omega) = G_1(\tfrac{\omega}{2})\ \hat{\phi}_g(\tfrac{\omega}{2}) = G_1(\tfrac{\omega}{2}) \left[\prod_{k=2}^{\infty} G_0(2^{-k}\omega)\right] \hat{\phi}_g(0) \tag{5}$$

And similarly, for the Tree B filters:

$$\hat{\phi}_h(\omega) \quad = \quad \left[\prod_{k=1}^{\infty} H_0(2^{-k}\omega)\right] \hat{\phi}_h(0) \tag{6}$$

$$\hat{\psi}_h(\omega) \quad = \quad H_1(\tfrac{\omega}{2}) \left[\prod_{k=2}^{\infty} H_0(2^{-k}\omega)\right] \hat{\phi}_h(0) \tag{7}$$

The amplitude scaling of $\phi_g(t)$ and $\phi_h(t)$ is arbitrary, so we choose $\hat{\phi}_g(0) = \hat{\phi}_h(0) = 1$ to give them both unit area.

Conjugate Quadrature Filterbank (CQF)

In an **orthonormal** wavelet transform, G_1 **and** G_0 **form a CQF** (also known as a Quadrature Mirror Filterbank, QMF), such that

$$G_1(\omega) = e^{-im\omega} \; G_0^*(\omega \pm \pi) \tag{8}$$

where we use $\pm\pi$ to emphasise the 2π-periodic nature of G_0 and G_1, and G_0^* means complex conjugate of G_0. The delay shift of m samples must be an odd integer and is usually chosen to approximately equalise the group delay or the support of G_0 and G_1.

Similarly

$$H_1(\omega) = e^{-im\omega} \; H_0^*(\omega \pm \pi) \tag{9}$$

Hence we can now express the wavelet frequency responses, $\hat{\psi}_g(\omega)$ and $\hat{\psi}_h(\omega)$, purely in terms of the two lowpass filters G_0 and H_0.

THE HILBERT PAIR CONDITION

This condition is

$$\frac{\hat{\psi}_h(\omega)}{\hat{\psi}_g(\omega)} = \begin{cases} i & \text{if } \omega < 0 \\ -i & \text{if } \omega > 0 \end{cases} \tag{10}$$

Note that the behaviour of the RHS at (and near) $\omega = 0$ is immaterial, because, for the wavelets to be admissible bandpass functions, $\hat{\psi}_g(0) = \hat{\psi}_h(0) = 0$.

Substituting (8) into (5) and (9) into (7), we get the following expression for this ratio

$$\begin{aligned} \frac{\hat{\psi}_h(\omega)}{\hat{\psi}_g(\omega)} &= \frac{e^{-im\omega/2}\ H_0^*(\frac{\omega}{2} \pm \pi)\left[\prod_{k=2}^{\infty} H_0(2^{-k}\omega)\right] \hat{\phi}_h(0)}{e^{-im\omega/2}\ G_0^*(\frac{\omega}{2} \pm \pi)\left[\prod_{k=2}^{\infty} G_0(2^{-k}\omega)\right] \hat{\phi}_g(0)} \\ &= R_0^*(\tfrac{\omega}{2} \pm \pi)\left[\prod_{k=2}^{\infty} R_0(2^{-k}\omega)\right] \end{aligned} \tag{11}$$

where $R_0(\omega) = H_0(\omega)/G_0(\omega)$ and is 2π-periodic. R_0 will give the desired relation between H_0 and G_0 if (10) and (11) are both satisfied.

Solving for $R_0(\omega)$

Since the modulus of the RHS of (10) is unity everywhere, and (11) contains an infinite product of terms in R_0, which will tend to grow or shrink if R_0 does not have unit magnitude, **we deduce that** $|R_0(\omega)| = 1$.

Now consider the phase $\theta(\omega)$ of R_0, by letting

$$R_0(\omega) = e^{i\theta(\omega)} \tag{12}$$

Equating the phases of (10) and (11), we require that

$$-\theta(\tfrac{\omega}{2} \pm \pi) + \sum_{k=2}^{\infty} \theta(2^{-k}\omega) = \begin{cases} \frac{\pi}{2} & \text{if } \omega < 0 \\ -\frac{\pi}{2} & \text{if } \omega > 0 \end{cases} \tag{13}$$

DEDUCING THE FORM OF $\theta(\omega)$

Because of the infinite sum in (13), we require $\theta(\omega) \to 0$ as $\omega \to 0$. Hence $\theta(0) = 0$.

Since $g_0(n)$ and $h_0(n)$ are purely real and lowpass, $R_0(\omega)$ must be conjugate symmetric and so $\theta(\omega)$ must be a continuous odd function about $\omega = 0$.

It can be shown (by Fourier analysis on $\theta'(\omega)$) that **any non-linear terms in $\theta(\omega)$ would prevent (13) from being satisfied**, because in (13) the gradient of the first term must cancel out the gradient of the summation terms at all $\omega \neq 0$.

Therefore we let

$$\theta(\omega) = \alpha\omega \quad \text{for } -\pi < \omega < \pi, \text{ where } \alpha \text{ is a constant.} \tag{14}$$

Hence

$$\theta(\tfrac{\omega}{2} \pm \pi) = \begin{cases} \alpha(\frac{\omega}{2} + \pi) & \text{if } -4\pi < \omega < 0 \\ \alpha(\frac{\omega}{2} - \pi) & \text{if } \quad 0 < \omega < 4\pi \end{cases} \tag{15}$$

Note that, since $\theta(\omega)$ must be 2π-periodic for $|\omega| \geq \pi$, it will have discontinuities at $\omega = \pm\pi$ if α is not an integer. These become discontinuities at $\omega = 0$ in $\theta(\frac{\omega}{2} \pm \pi)$.

TYPICAL PLOTS OF $\theta(\omega)$ AND TERMS IN EQU.(13)

$\theta(\omega) \approx \alpha\omega, \ -\pi < \omega < \pi$
and $\theta(\omega)$ is 2π-periodic

$-\theta(\frac{\omega}{2} \pm \pi)$

$\theta(\frac{\omega}{4})$

$\theta(\frac{\omega}{8})$

$-\theta(\frac{\omega}{2} \pm \pi) + \sum_{k=2}^{\infty} \theta(2^{-k}\omega)$

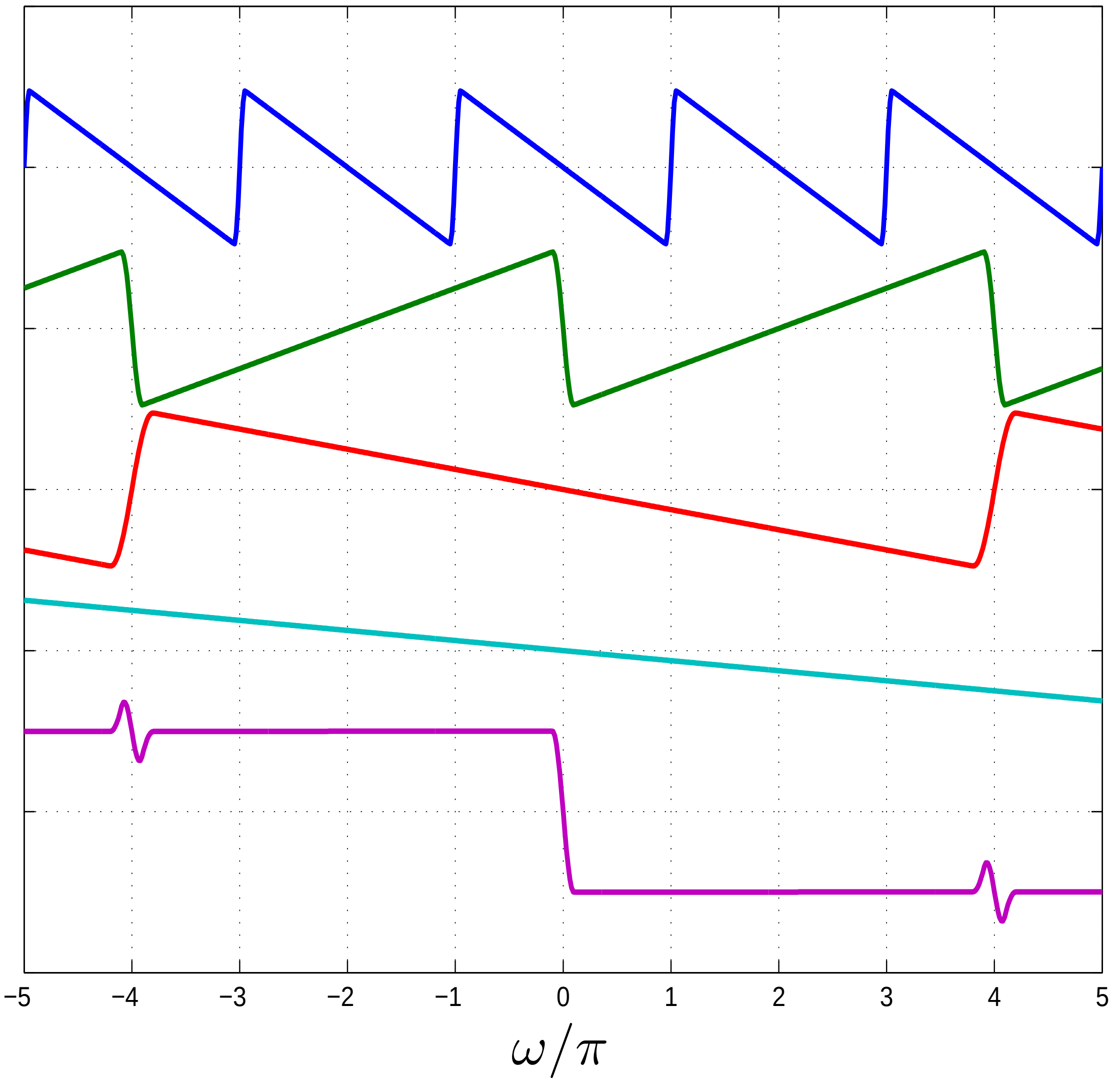

CALCULATING α

Noting that $\displaystyle\sum_{k=2}^{\infty} \theta(2^{-k}\omega) = \alpha\omega[\tfrac{1}{4} + \tfrac{1}{8} + \tfrac{1}{16} + \ldots] = \tfrac{\alpha\omega}{2}$ if $-4\pi < \omega < 4\pi$,

and substituting (15) into (13) gives

$$-\alpha(\tfrac{\omega}{2} + \pi) + \tfrac{\alpha\omega}{2} = \tfrac{\pi}{2} \quad \text{if } -4\pi < \omega < 0 \qquad (16)$$

$$\text{and} \quad -\alpha(\tfrac{\omega}{2} - \pi) + \tfrac{\alpha\omega}{2} = -\tfrac{\pi}{2} \quad \text{if } 0 < \omega < 4\pi \qquad (17)$$

(16) and (17) are both satisfied if $\alpha = -\frac{1}{2}$, and therefore

$$\frac{H_0(\omega)}{G_0(\omega)} = R_0(\omega) = e^{i\theta(\omega)} = e^{i\alpha\omega} = e^{-i\omega/2} \quad \text{for } -\pi < \omega < \pi \qquad (18)$$

This is the **half-sample delay** solution that makes $\psi_h(t)$ the Hilbert transform of $\psi_g(t)$.

Ozkaramanli and Yu (Dec 2005 and June 2006) have shown this solution to be **unique** and applicable to **biorthogonal** as well as **orthonormal** wavelet transforms.

Q-shift Dual Tree Complex Wavelet Transform in 1-D

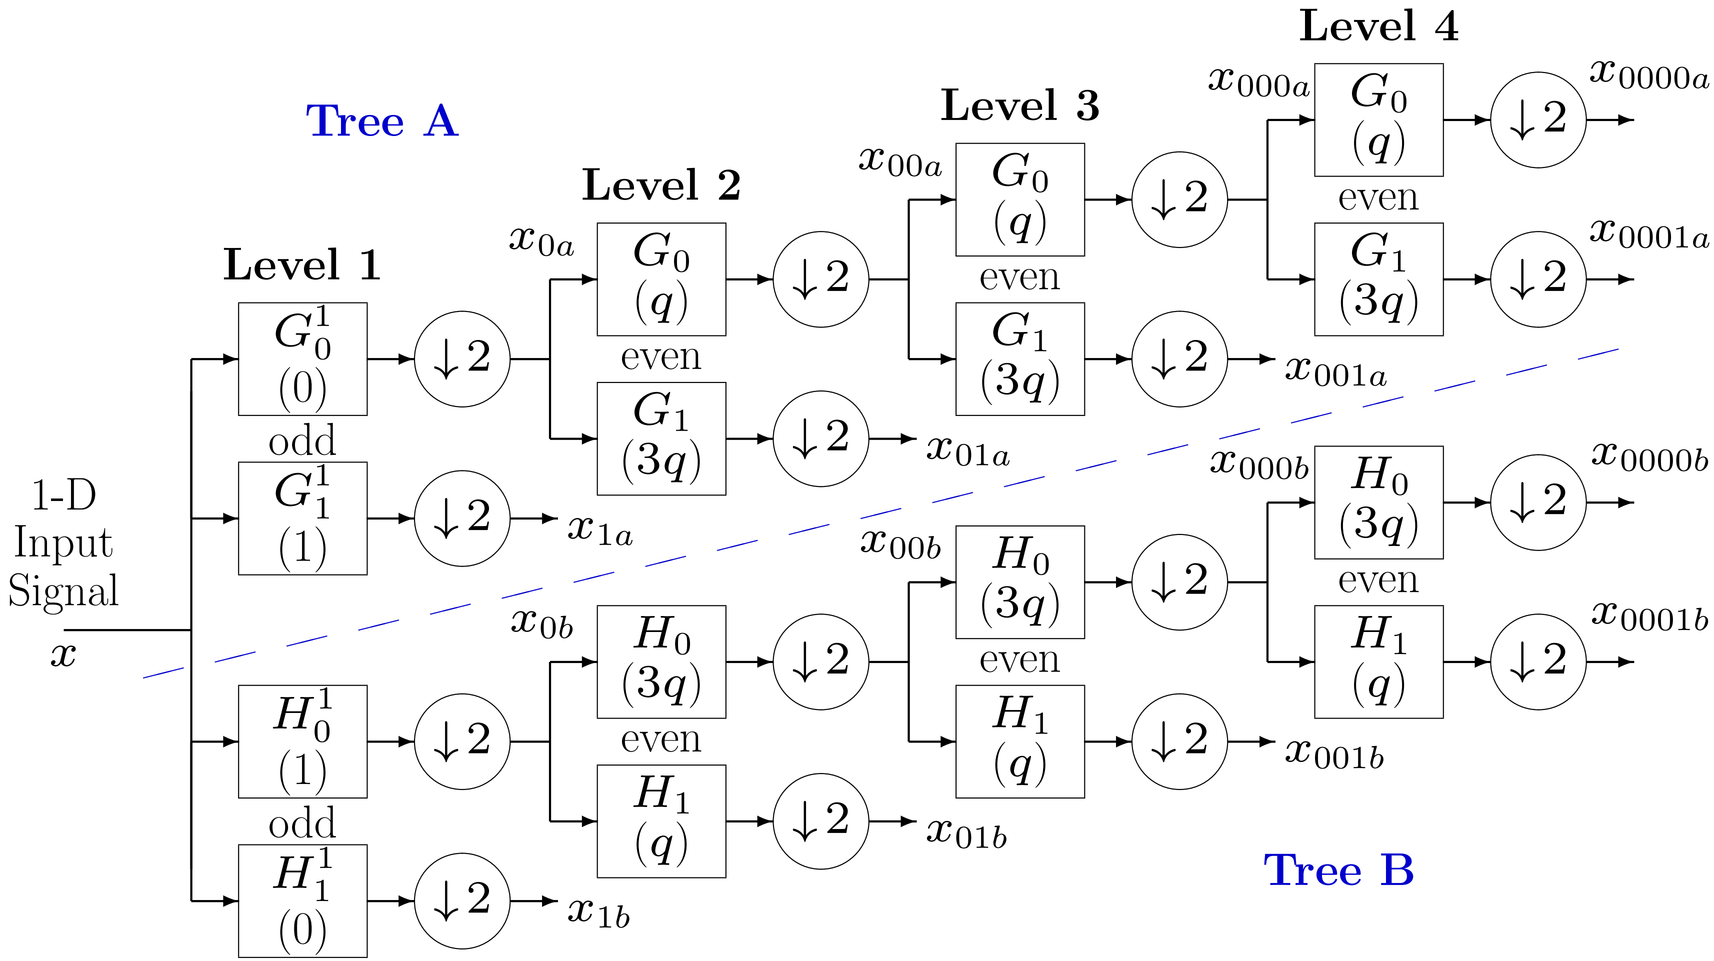

Figure 6: Dual tree of real filters for the Q-shift $\mathbb{C}$WT, giving real and imaginary parts of complex coefficients from tree A and tree B respectively. Figures in brackets indicate the approximate delay for each filter, where $q = \frac{1}{4}$ sample period. Special level 1 filters, G^1 and H^1, allow for the finite number of levels.

Why does the delay condition give shift invariance?

- The **half-sample delay** between the G_0 and H_0 **lowpass** filters means that their output samples are **uniformly interleaved** at all scales, and hence the sample rate is effectively **doubled** everywhere.

- The doubled sampling rate is sufficient to virtually **eliminate aliasing** if filters of 12 or more taps are used.

- If aliasing is eliminated in the **lowpass** branch of each 2-band reconstruction block, then it must also be eliminated in the **highpass** branch (to obtain perfect reconstruction).

- Elimination of aliasing means that each subband can be represented by a **unique z-transfer function**, and hence the filtering is **LTI**, linear time-invariant (i.e. shift-invariant).

At level 1 of a finite dual tree, the delay difference must **increase to 1 sample** to compensate for the absence of delay differences at finer levels.

WHY DO WE USE Q-SHIFT FILTERS (BELOW LEVEL 1)?

- Half-sample delay difference is obtained with filter delays of $\frac{1}{4}$ and $\frac{3}{4}$ of a sample period (instead of 0 and $\frac{1}{2}$ a sample for our original DT $\mathbb{C}$WT).
- This is achieved with an **asymmetric even-length** filter $G_0(z)$ and its time reverse $H_0(z) = z^{-1}\, G_0(z^{-1})$. $G_1(z)$ and $H_1(z)$ are the CQFs of these.
- Due to the asymmetry (like Daubechies filters), these may be designed to give an **orthonormal perfect reconstruction** wavelet transform in each tree.
- Tree **B** filters are the **reverse** of tree **A** filters, and reconstruction filters are the reverse of analysis filters, so **all filters** are from the **same orthonormal set**.
- Both trees have the **same frequency responses** (in magnitude).
- The combined **complex** impulse responses are **conjugate symmetric** about their mid points, even though the separate responses are asymmetric. Hence **symmetric extension** still works at image edges.

At level 1, **any DWT filters** can be used.

Q-shift DT ℂWT Basis Functions – Levels 1 to 3

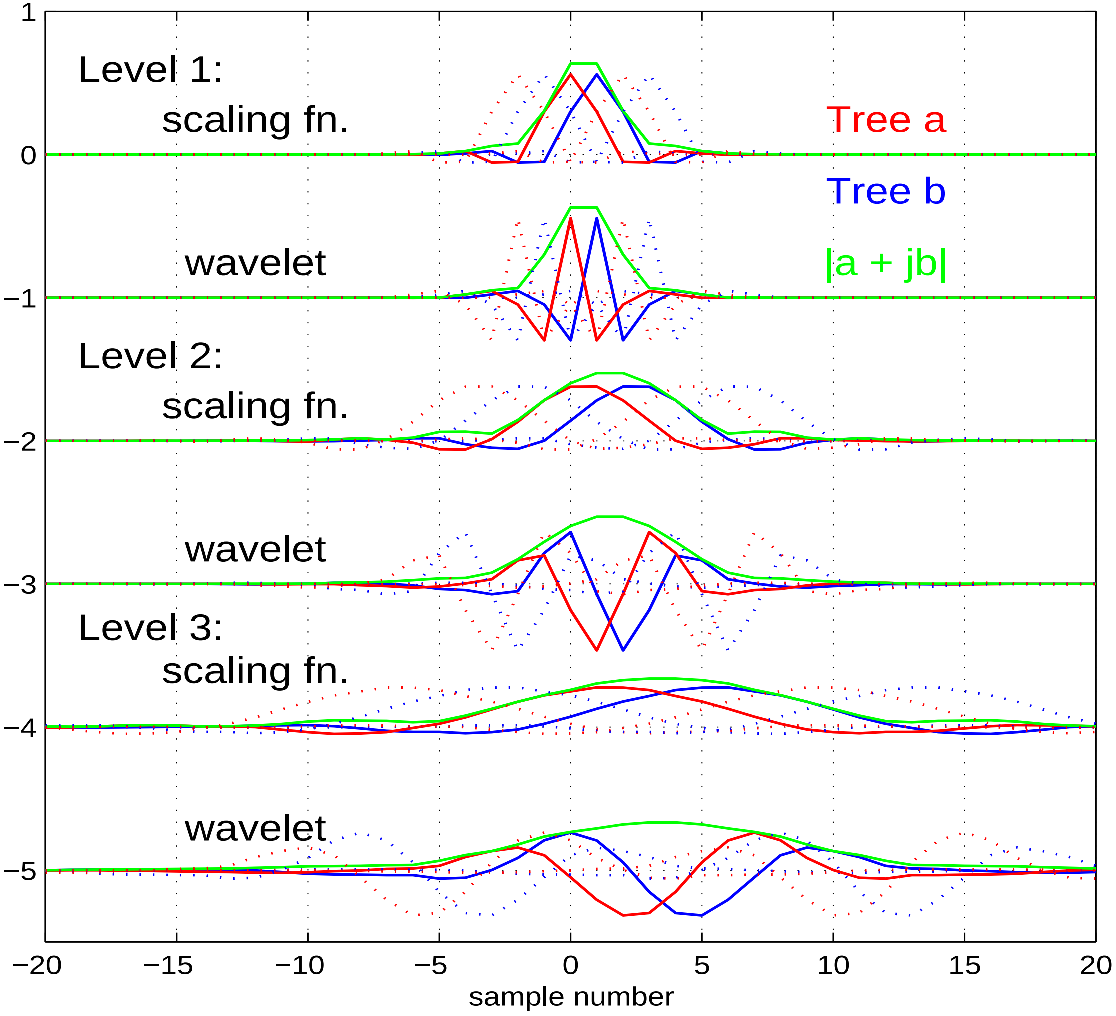

Figure 7: Basis functions for adjacent sampling points are shown dotted.

Frequency Responses of 18-tap Q-shift filters

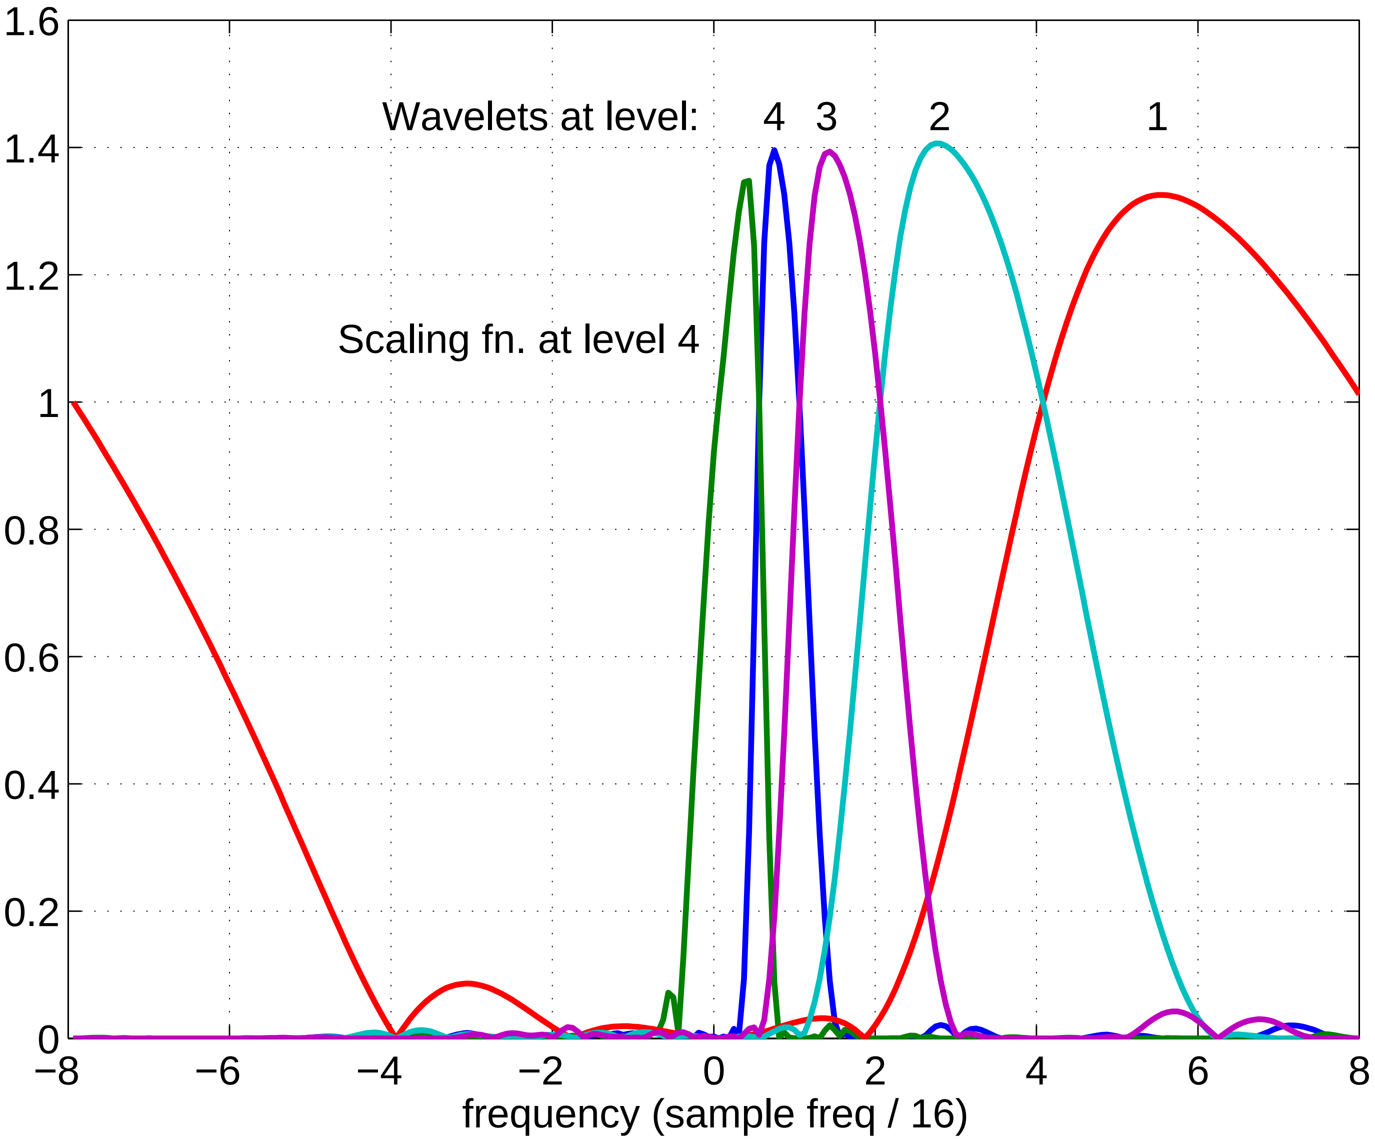

Visualising Shift Invariance

- Apply a standard input (e.g. unit step) to the transform for a **range of shift positions**.
- Select the transform coefficients from **just one wavelet level** at a time.
- Inverse transform each set of selected coefficients.
- Plot the component of the reconstructed output for each shift position at each wavelet level.
- Check for **shift invariance** (similarity of waveforms).

Fig 8 shows that the DT $\mathbb{C}$WT has near-perfect shift invariance, whereas the maximally-decimated real discrete wavelet transform (DWT) has substantial shift dependence.

SHIFT INVARIANCE OF COMPLEX DT ℂWT VS REAL DWT

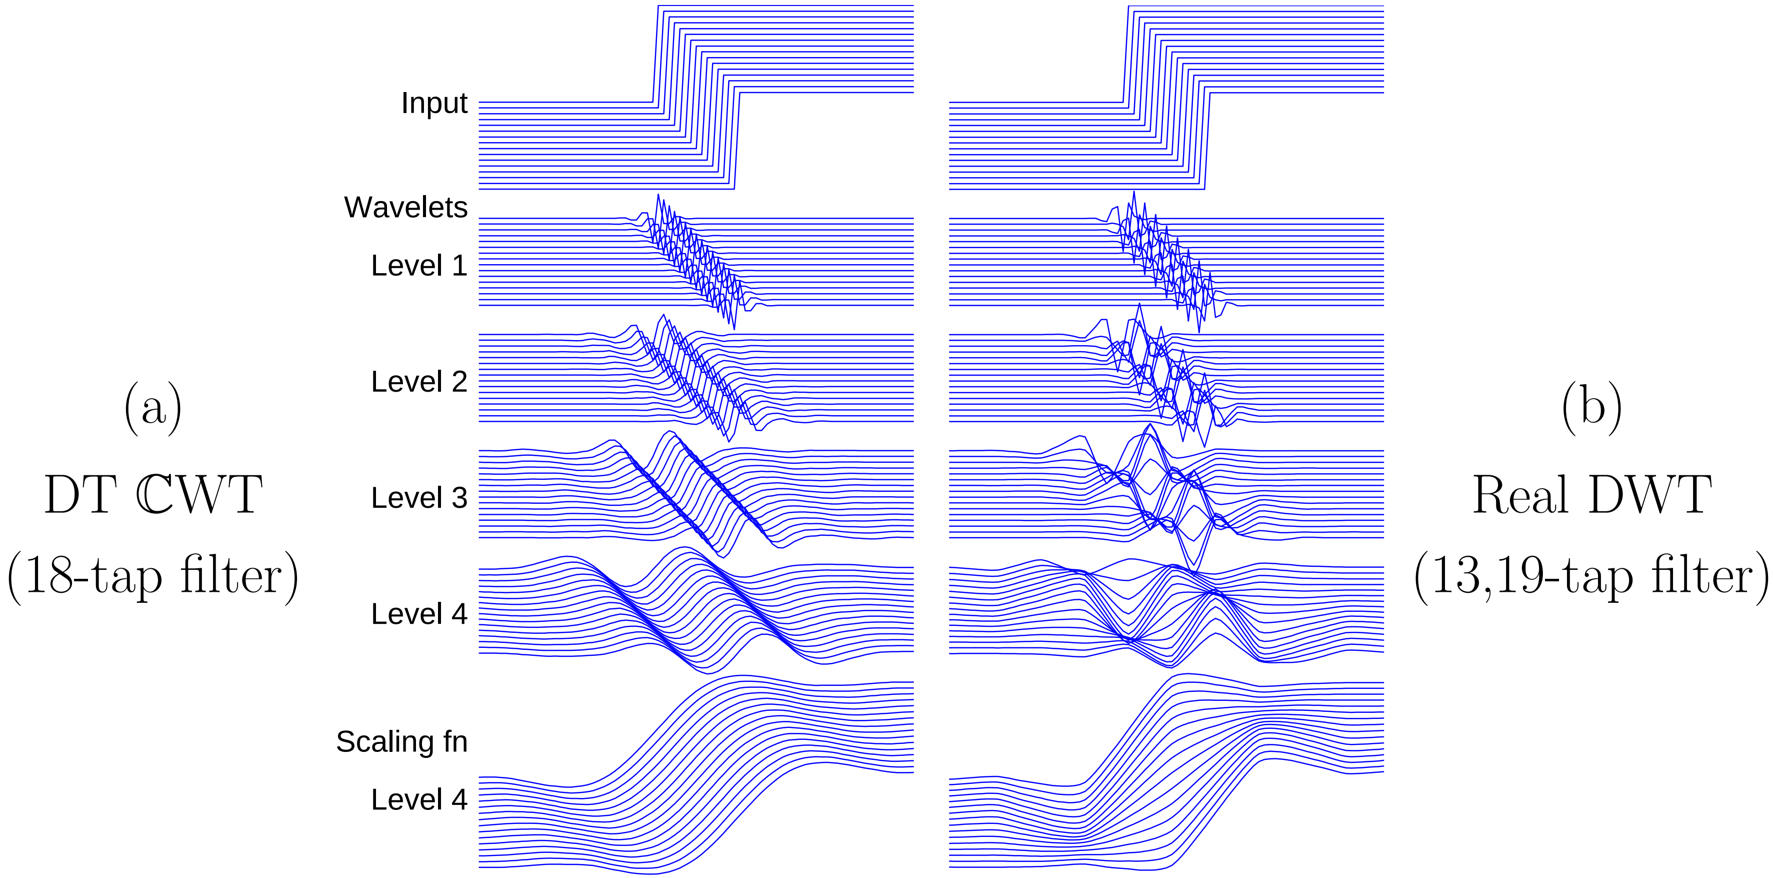

Figure 8: Wavelet and scaling function components at levels 1 to 4 of 16 shifted step responses of the DT ℂWT (a) and real DWT (b). If there is good shift invariance, all components at a given level should be similar in shape, as in (a).

Shift Invariance of simpler DT ℂWTs

(a)

DT ℂWT

(14-tap filter)

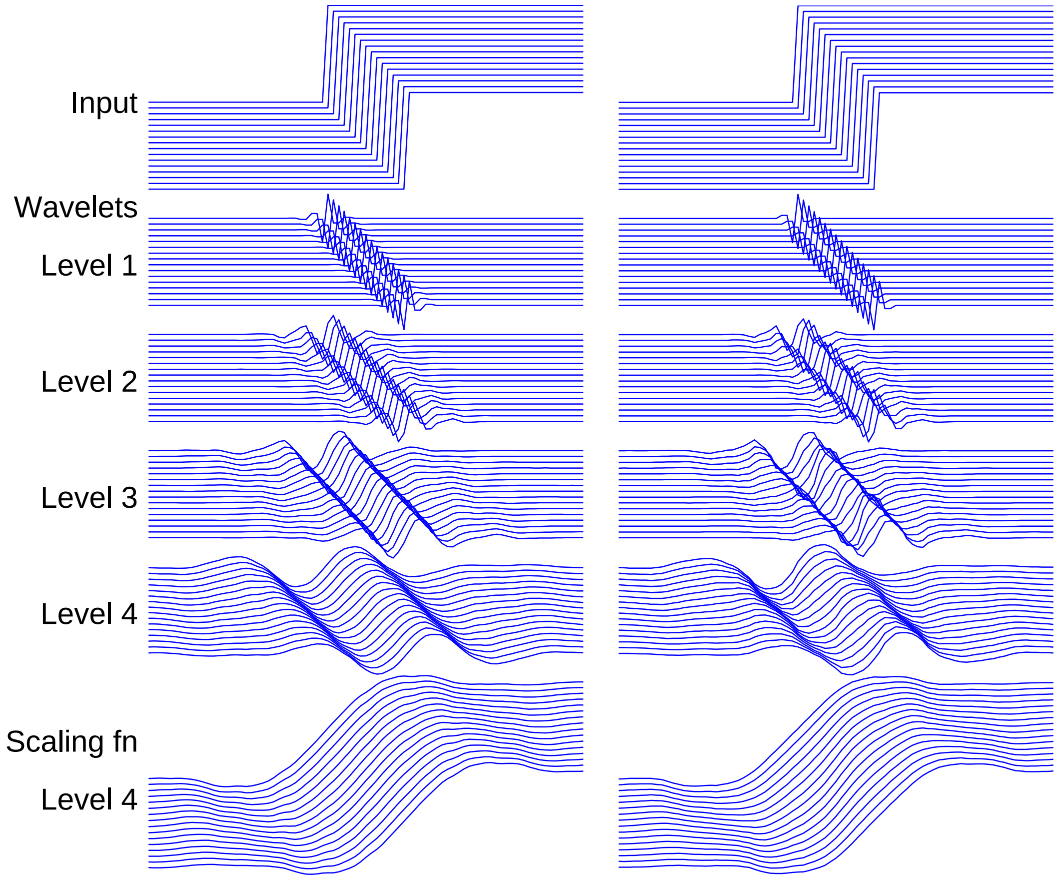

(b)

DT ℂWT

(6-tap filter)

Figure 9: Wavelet and scaling function components at levels 1 to 4 of 16 shifted step responses of simpler forms of the DT ℂWT, using (a) 14-tap and (b) 6-tap Q-shift filters.

Shift Invariance – Quantitative measurement

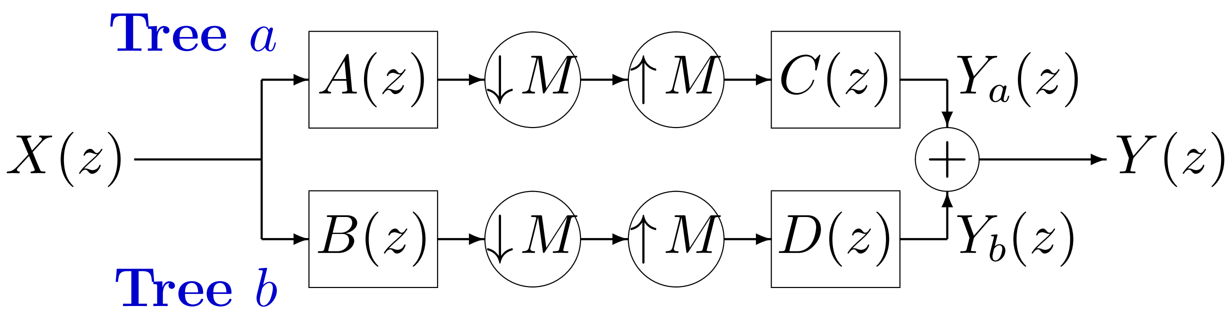

Basic configuration of the dual tree if either wavelet or scaling-function coefficients from just level m are retained ($M = 2^m$).

Letting $W = e^{i2\pi/M}$, **multi-rate analysis** gives:

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

For shift invariance, **aliasing terms ($k \neq 0$) must be negligible.** So we design $B(W^k z)\, D(z)$ to cancel $A(W^k z)\, C(z)$ for all non-zero k that give overlap of the passbands of filters $C(z)$ or $D(z)$ with those of shifted filters $A(W^k z)$ or $B(W^k z)$.

A Measure of Shift Invariance

Since

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

we quantify the shift dependence of a transform by calculating the ratio of the total energy of the **unwanted aliasing transfer functions** (the terms with $k \neq 0$) to the energy of the **wanted transfer function** (when $k = 0$):

$$R_a = \frac{\sum_{k=1}^{M-1} \mathcal{E}\{A(W^k z)\, C(z) + B(W^k z)\, D(z)\}}{\mathcal{E}\{A(z)\, C(z) + B(z)\, D(z)\}}$$

where $\mathcal{E}\{U(z)\}$ calculates the energy, $\sum_r |u_r|^2$, of the impulse response of a z-transfer function, $U(z) = \sum_r u_r z^{-r}$.

$\mathcal{E}\{U(z)\}$ may also be interpreted in the **frequency domain** as the integral of the squared magnitude of the frequency response, $\frac{1}{2\pi} \int_{-\pi}^{\pi} |U(e^{i\theta})|^2 \, d\theta$ from Parseval's theorem.

TYPES OF DT $\mathbb{C}$WT FILTERS

We show results for the following combinations of filters:

A **(13,19)-tap** and **(12,16)-tap** near-orthogonal odd/even filter sets.

B **(13,19)-tap** near-orthogonal filters at level 1, **18-tap** Q-shift filters at levels ≥ 2.

C **(13,19)-tap** near-orthogonal filters at level 1, **14-tap** Q-shift filters at levels ≥ 2.

D **(9,7)-tap** bi-orthogonal filters at level 1, **18-tap** Q-shift filters at levels ≥ 2.

E **(9,7)-tap** bi-orthogonal filters at level 1, **14-tap** Q-shift filters at levels ≥ 2.

F **(9,7)-tap** bi-orthogonal filters at level 1, **6-tap** Q-shift filters at levels ≥ 2.

G **(5,3)-tap** bi-orthogonal filters at level 1, **6-tap** Q-shift filters at levels ≥ 2.

Aliasing Energy Ratios,

Values of R_a **in dB**, for filter types A to G over levels 1 to 5.

Filters:	A	B	C	D	E	F	G	DWT
Complexity:	2.0	2.3	2.0	1.9	1.6	1.0	0.7	1.0
Wavelet								
Level 1	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	-9.40
Level 2	-28.25	-31.40	-29.06	-22.96	-21.81	-18.49	-14.11	-3.54
Level 3	-23.62	-27.93	-25.10	-20.32	-18.96	-14.60	-11.00	-3.53
Level 4	-22.96	-31.13	-24.67	-32.08	-24.85	-16.78	-15.80	-3.52
Level 5	-22.81	-31.70	-24.15	-31.88	-24.15	-18.94	-18.77	-3.52
Scaling fn.								
Level 1	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	-9.40
Level 2	-29.37	-32.50	-30.17	-24.32	-23.19	-19.88	-15.93	-9.38
Level 3	-28.17	-35.88	-29.21	-36.94	-29.33	-21.75	-20.63	-9.37
Level 4	-27.88	-37.14	-28.57	-37.37	-28.56	-24.37	-24.15	-9.37
Level 5	-27.75	-36.00	-28.57	-36.01	-28.57	-24.67	-24.65	-9.37

How do we extend the DT $\mathbb{C}$WT to multi-dimensions?

When the DT $\mathbb{C}$WT is applied to 2-D signals (images), it has the following features:

- It is performed separably, with 2 trees used for the rows of the image and 2 trees for the columns – yielding a **Quad-Tree** structure (4:1 redundancy).

- The 4 quad-tree components of each coefficient are combined by simple sum and difference operations to yield a **pair of complex coefficients**. These are part of two separate subbands in adjacent quadrants of the 2-D spectrum.

- This produces **6 directionally selective subbands** at each level of the 2-D DT $\mathbb{C}$WT. Fig 10 shows the basis functions of these subbands at level 4, and compares them with the 3 subbands of a 2-D DWT.

- The DT $\mathbb{C}$WT is directionally selective because the complex filters can **separate positive and negative frequency components** in 1-D, and hence **separate adjacent quadrants** of the 2-D spectrum. Real separable filters cannot do this!

WHY DO WE GET GOOD DIRECTIONAL FILTERS IN 2-D?

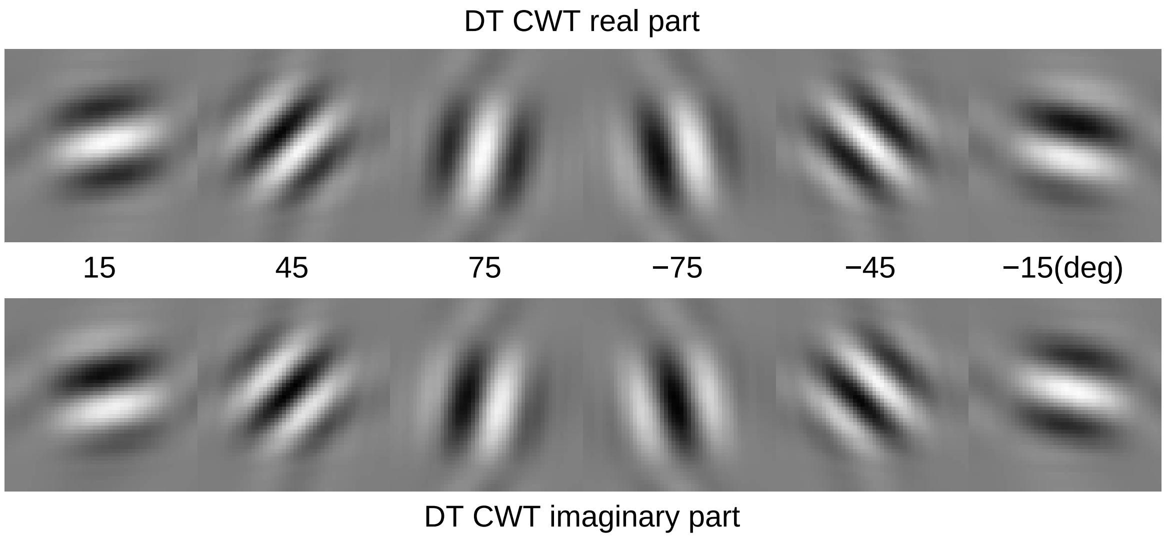

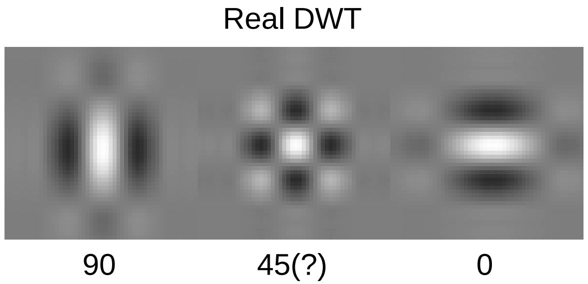

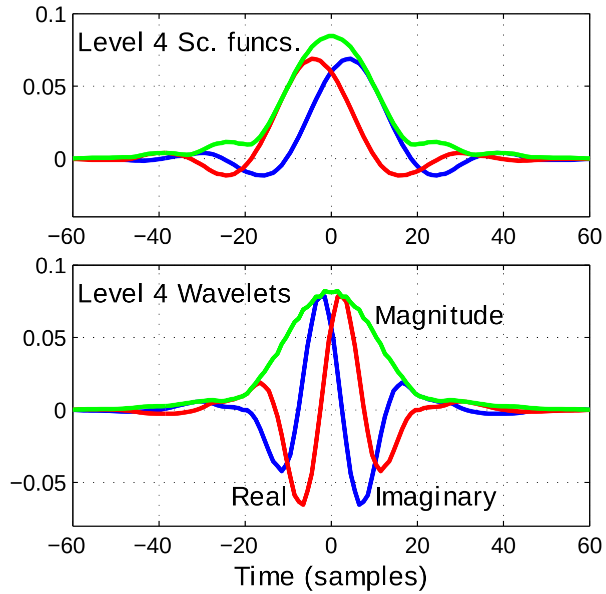

$$g(x)\, e^{i\omega_1 x} \;.\; g(y)\, e^{i\omega_2 y} =$$
$$g(x)\, g(y)\;\; e^{i(\omega_1 x + \omega_2 y)}$$

Figure 10: Basis functions of 2-D Q-shift complex wavelets (top), and of 2-D real wavelet filters (bottom), all illustrated at level 4 of the transforms. The complex wavelets provide 6 directionally selective filters, while real wavelets provide 3 filters, only two of which have a dominant direction. The 1-D bases, from which the 2-D complex bases are derived, are shown to the right.

Test Image and Colour Palette for Complex Coefficients

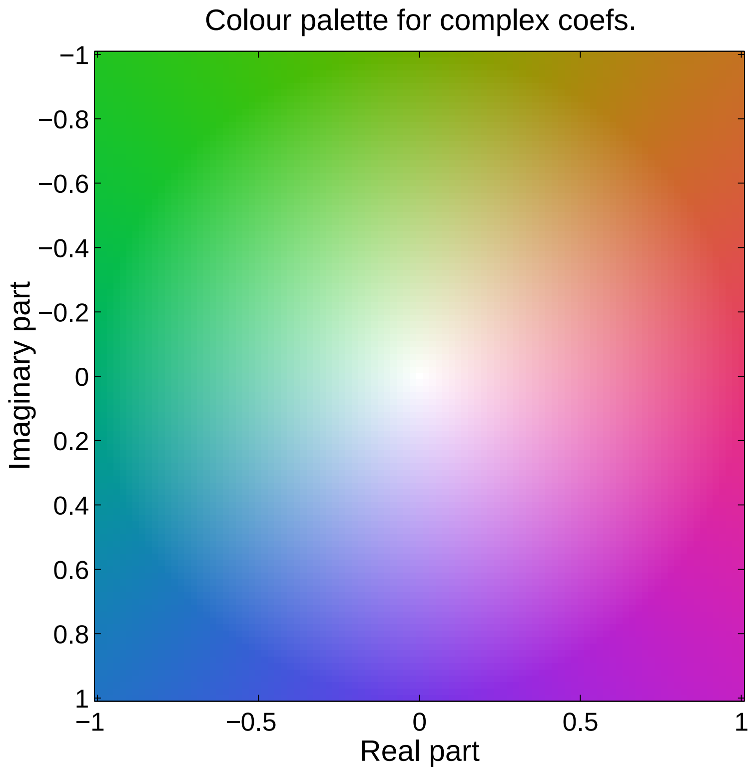

2-D DT ℂWT Decomposition into Subbands

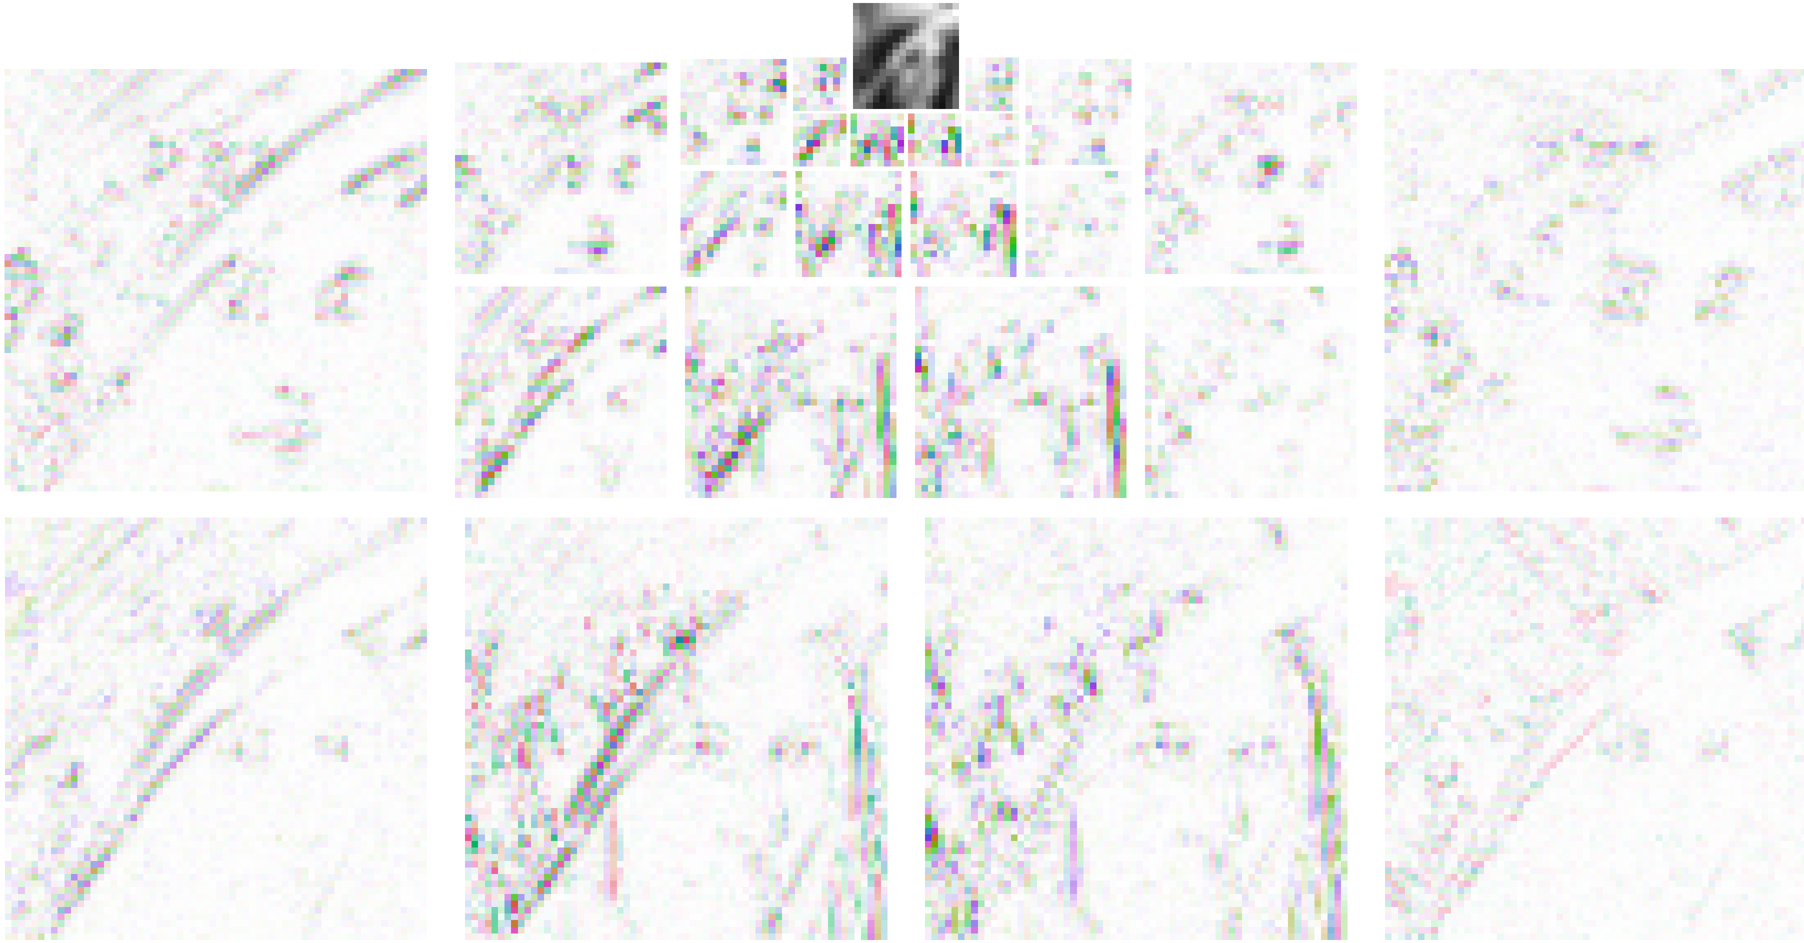

Figure 11: Four-level DT ℂWT decomposition of *Lenna* into 6 subbands per level (only the central 128×128 portion of the image is shown for clarity). A colour-disc palette (see previous slide) is used to display the complex wavelet coefficients.

2-D DT $\mathbb{C}$WT Reconstruction Components from Each Subband

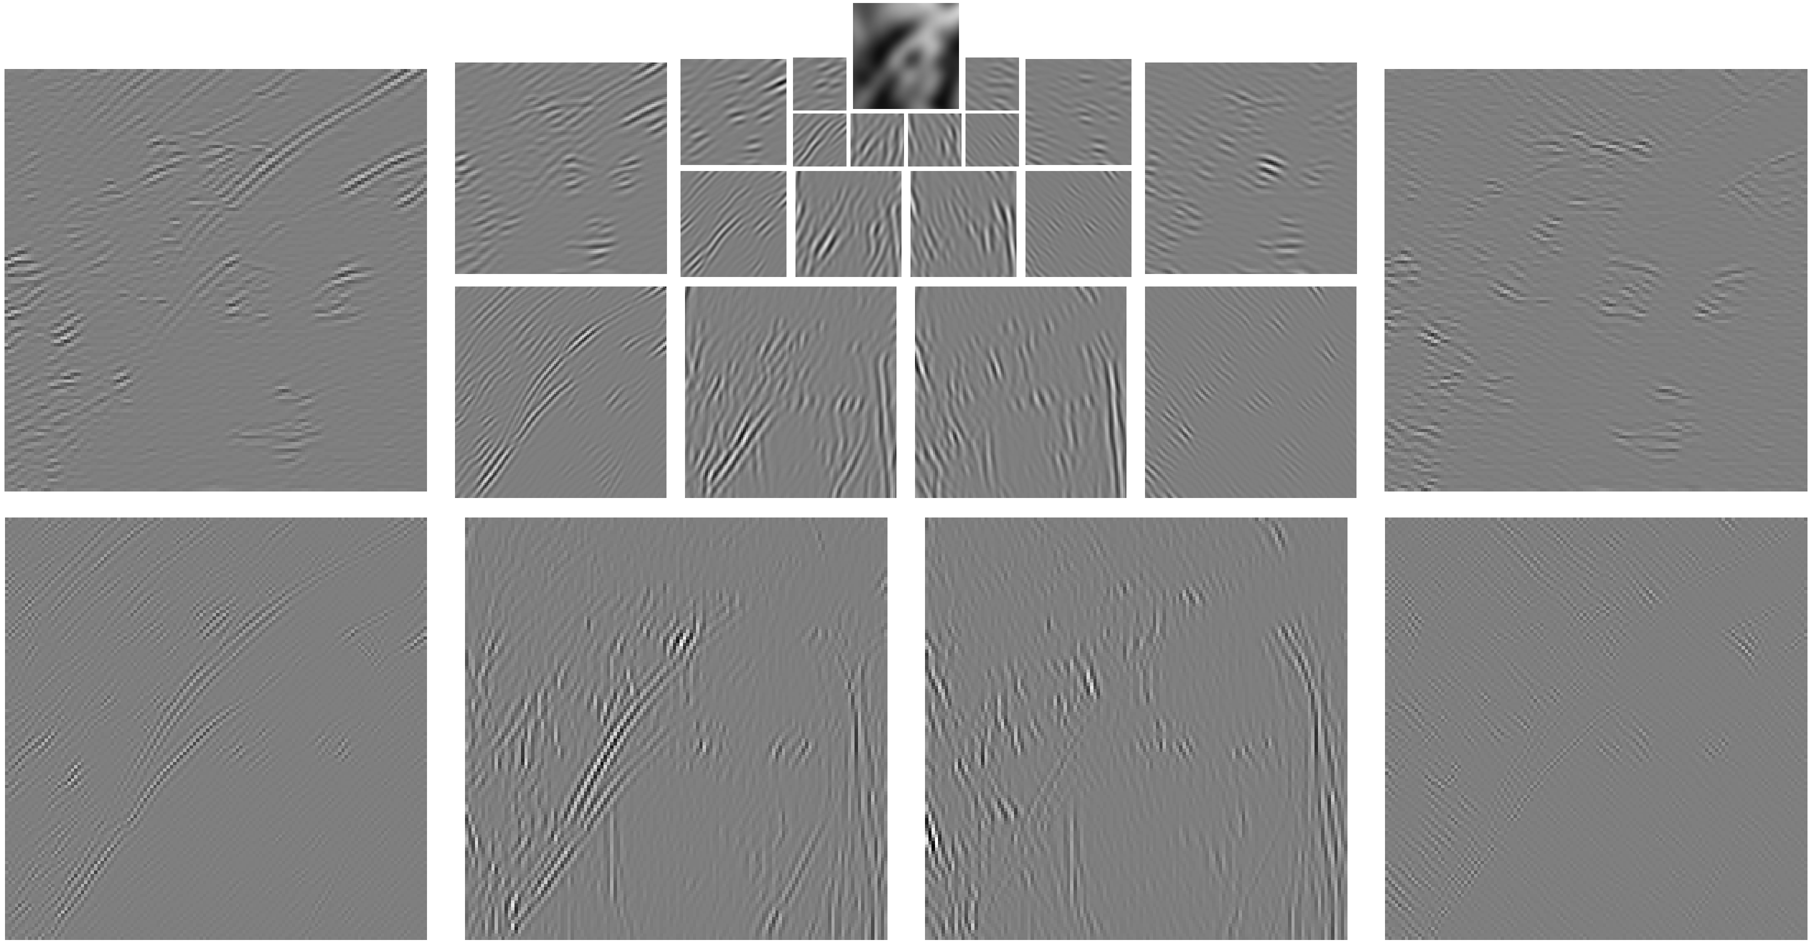

Figure 12: Components from each subband of the reconstructed output image for a 4-level DT $\mathbb{C}$WT decomposition of *Lenna* (central 128×128 portion only).

2-D Shift Invariance of Complex DT ℂWT vs Real DWT

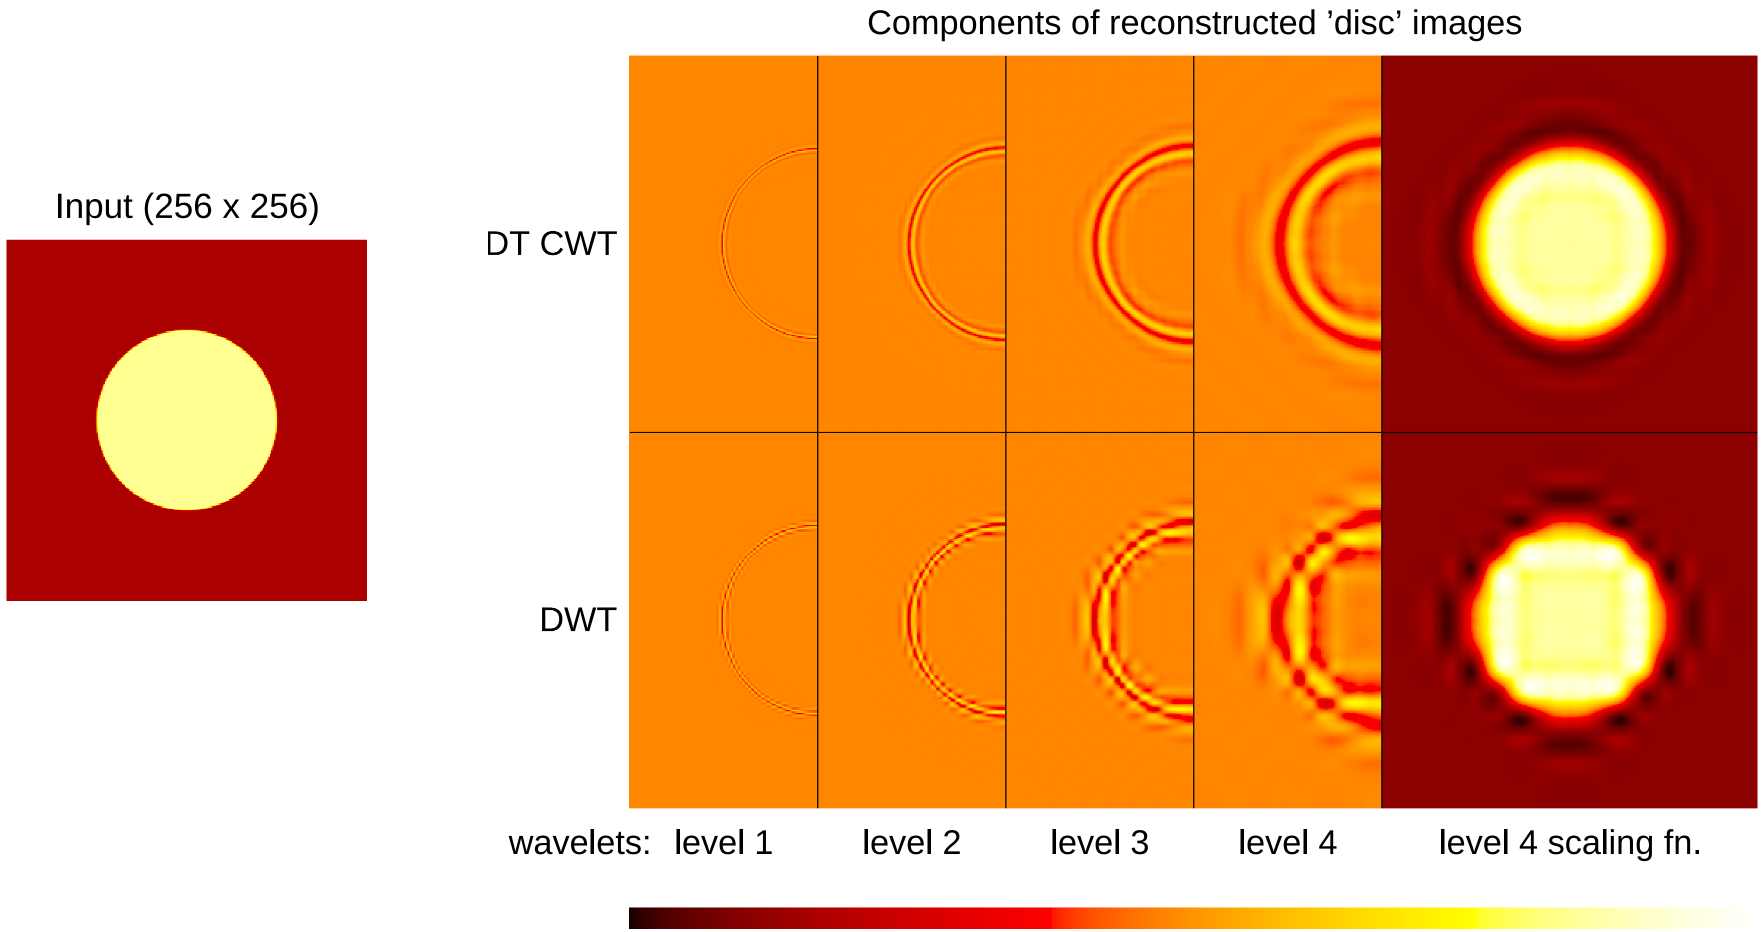

Figure 13: Wavelet and scaling function components at levels 1 to 4 of an image of a light circular disc on a dark background, using the 2-D DT ℂWT (upper row) and 2-D DWT (lower row). Only half of each wavelet image is shown in order to save space.

How do we use the DT $\mathbb{C}$WT in 3-D ?

When the DT $\mathbb{C}$WT is applied to 3-D signals (eg medical MRI or CT datasets), it has the following features:

- It is performed separably, with 2 trees used for the rows, 2 trees for the columns and 2 trees for the slices of the 3-D dataset – yielding an **Octal-Tree** structure (8:1 redundancy).

- The 8 octal-tree components of each coefficient are combined by simple sum and difference operations to yield a **quad of complex coefficients**. These are part of 4 separate subbands in adjacent octants of the 3-D spectrum.

- This produces **28 directionally selective subbands** $(4 \times 8 - 4)$ at each level of the 3-D DT $\mathbb{C}$WT. The subband basis functions are now approximately **planar waves** of the form $e^{i(\omega_1 x+\omega_2 y+\omega_3 z)}$, modulated by a 3-D Gaussian envelope (i.e. **3-D Morlet wavelets**).

- Each subband responds to approximately **flat surfaces** of a particular orientation. There are 7 orientations on each quadrant of a hemisphere.

3D SUBBAND ORIENTATIONS ON ONE QUADRANT OF A HEMISPHERE

3D frequency domain:

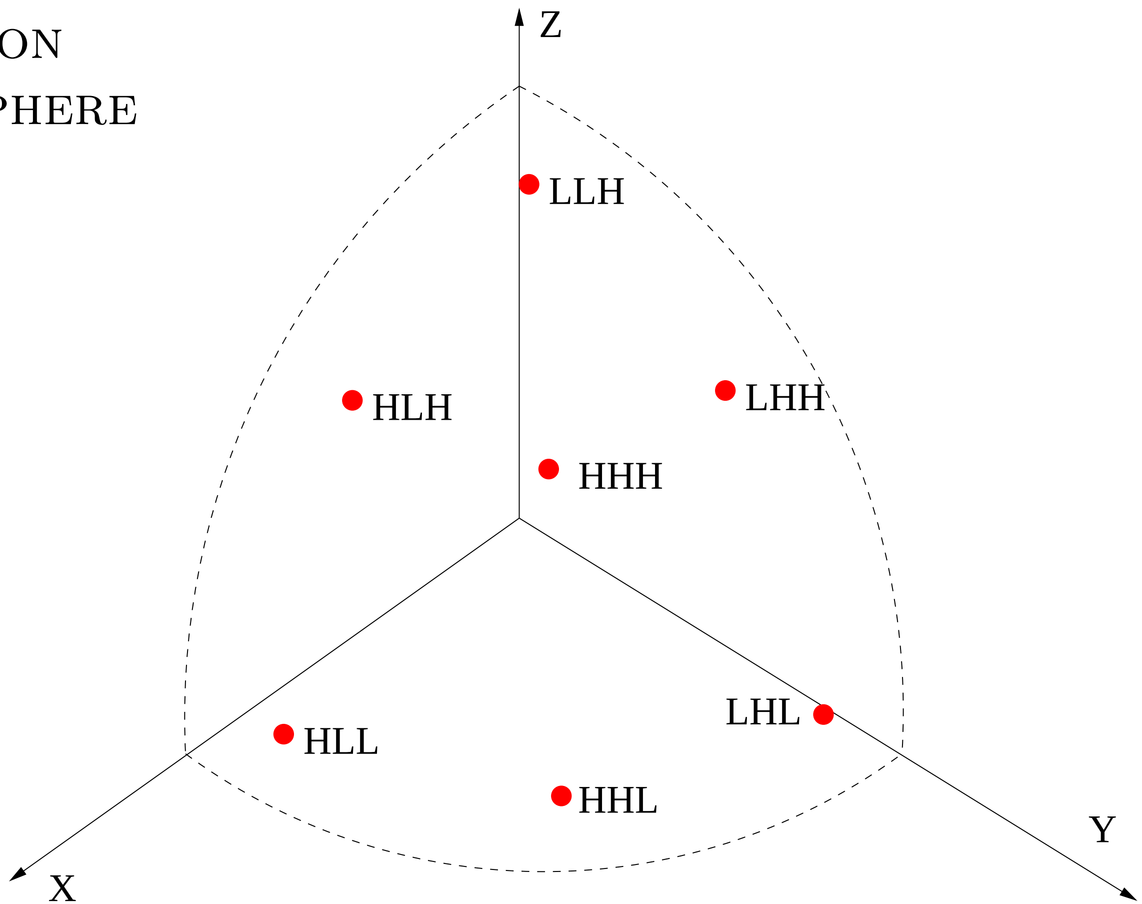

3D Morlet-like basis functions:

$$h_{k1,k2,k3}(x,y,z) \simeq e^{-(x^2+y^2+z^2)/2\sigma^2} \times e^{i(\omega_{k1}\,x + \omega_{k2}\,y + \omega_{k3}\,z)}$$

These are **28 planar waves** (7 per quadrant of a hemisphere) whose orientation depends on $\omega_{k1} \in \{\omega_L, \omega_H\}$ and $\omega_{k2}, \omega_{k3} \in \{\pm\omega_L, \pm\omega_H\}$, where $\omega_H \simeq 3\omega_L$.

Some applications of the DT $\mathbb{C}$WT

- **Motion estimation** [Magarey 98]
- **Motion compensation** and **registration** [Kingsbury 02, Hemmendorff 02]
- **Denoising** [Choi 00, Miller 06]
- **Deconvolution** [Jalobeanu 00, De Rivaz 01, J Ng 07]
- **Texture analysis** [Hatipoglu 99] and **synthesis** [De Rivaz 00]
- **Segmentation** [De Rivaz 00, Shaffrey 02], **classification** [Romberg 00] and **image retrieval** [Kam & T T Ng 00, Shaffrey 03]
- **Watermarking of images** [Loo 00] and **video** [Earl 03]
- **Compression / Coding** [Reeves 03]
- **Seismic analysis** [van Spaendonck & Fernandes 02, Miller 05]
- **Diffusion Tensor MRI visualisation** [Zymnis 04]
- **Object matching & recognition** [Anderson, Fauqueur & Kingsbury 06]
- **Image fusion** [Nikolov & Bull 07] and **object tracking** [Pang & Nelson 08]
- **Sparse image and 3D-data reconstruction** [Zhang 08 & 10]

MOTION ESTIMATION AND IMAGE REGISTRATION

Our proposed algorithm for **robust registration** effectively combines

- **The Dual-Tree Complex Wavelet Transform**
  - Linear phase vs. shift behaviour
  - Easy shiftability of subbands
  - Directional filters select edge-like structures
  - Good denoising of input images

- **Hemmendorf's phase-based parametric method**
  (Hemmendorff et al, IEEE Trans Medical Imaging, Dec 2002)
  - Finds LMS fit of parametric model to edges in images
  - Allows simple filtering of $\mathbf{Q}_{\mathcal{X}}$ to fit more complex motions
  - Integrates well with multiscale DT $\mathbb{C}$WT structure

Registration Algorithm:

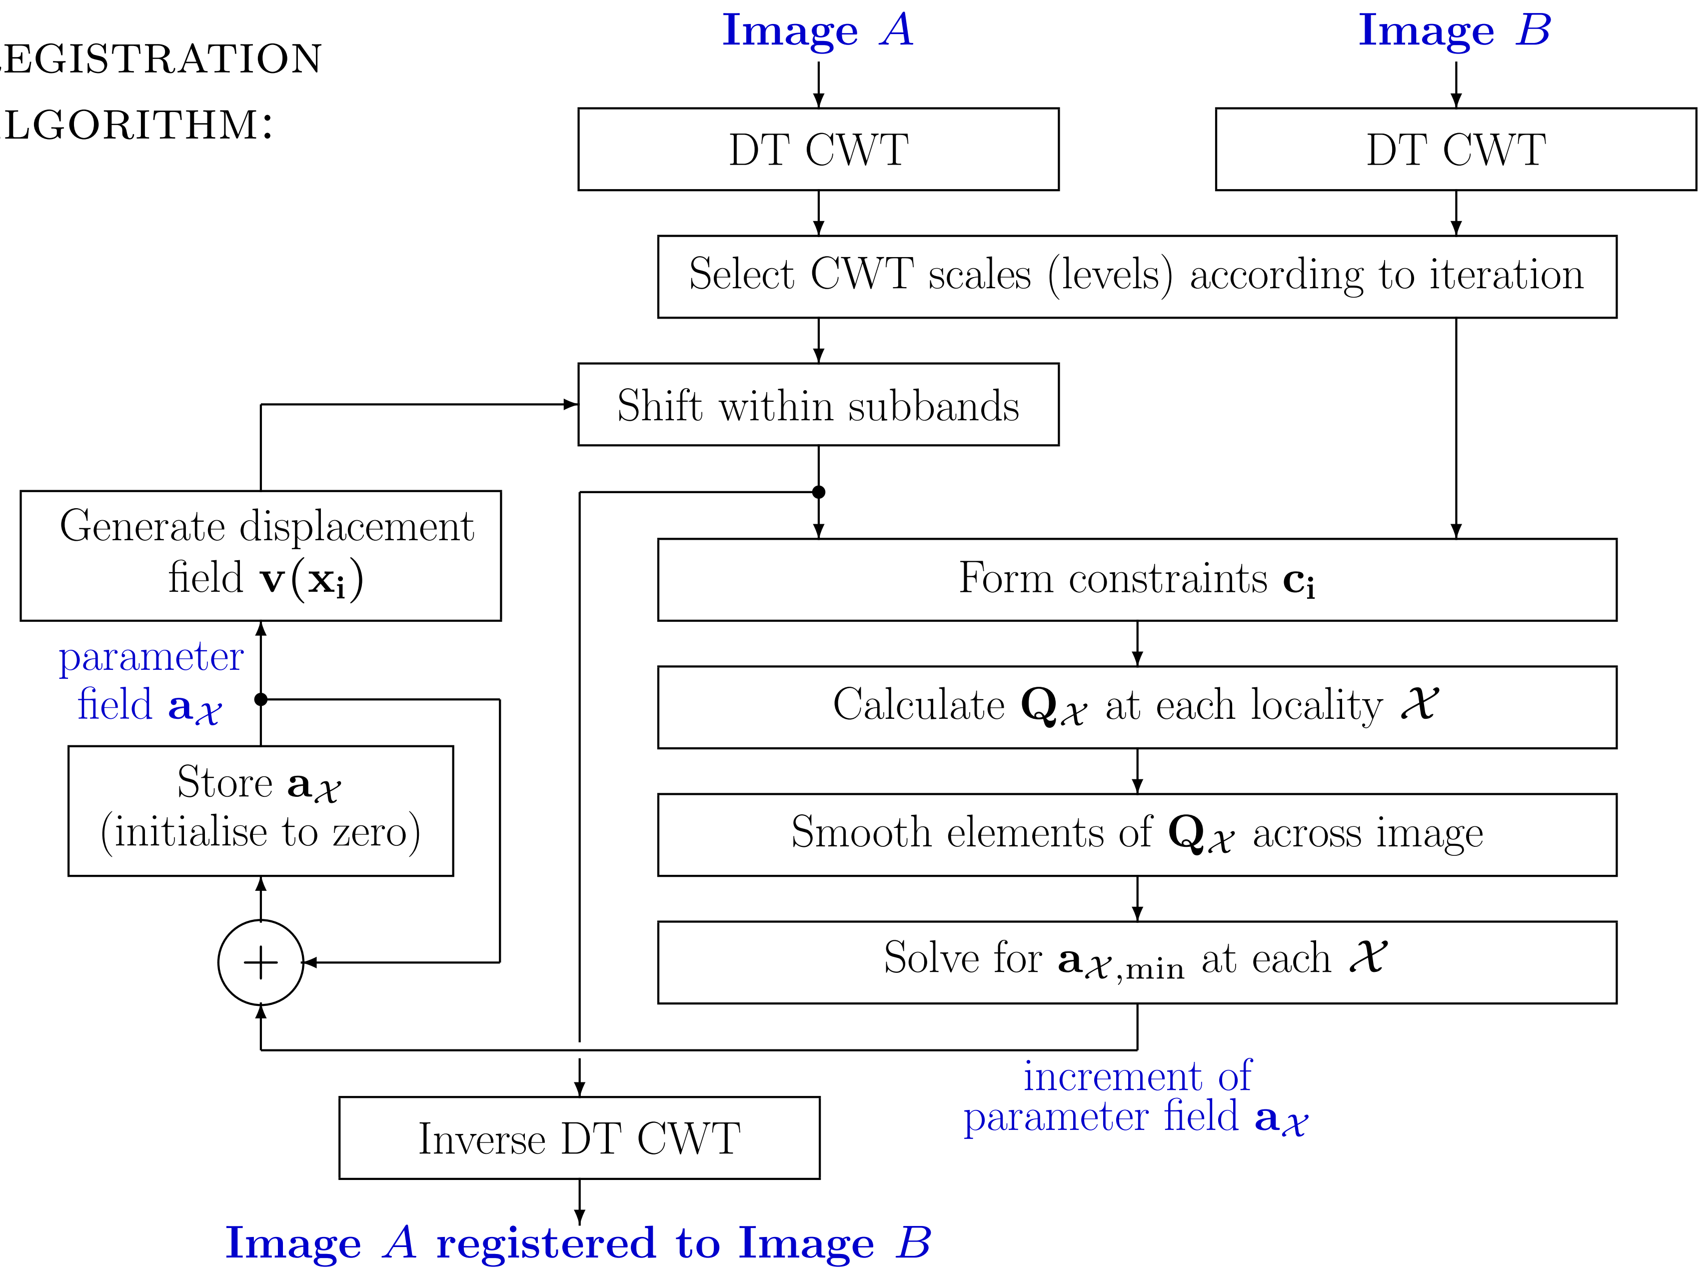

Demonstrations

- Registration of CT scans:
  - Two scans of the abdomen of the same patient, taken at different times with significant differences in position and contrast.
  - Task is to register the two images as well as possible, despite the differences.

- Enhancement of video corrupted by atmospheric turbulence, using registration and complex wavelet fusion across frames:
  - 75 frames of video of a house on a distant hillside, taken through a high-zoom lens with significant turbulence of the intervening atmosphere due to rising hot air (courtesy of ADFL, Canberra).
  - Task is to register each frame to a 'mean' image from the sequence, and then to reconstruct a high-quality still image by fusion of information from the whole registered sequence.

Conclusions

The **Dual-Tree Complex Wavelet Transform** provides:

- Approximate **shift invariance**
- **Directionally selective** filtering in 2 or more dimensions
- **Low redundancy** – only $2^m : 1$ for m-D signals
- **Perfect reconstruction**
- **Orthonormal filters** below level 1, but still giving **linear phase** (conjugate symmetric) complex wavelets
- **Low computation** – order-N; less than 2^m times that of the fully decimated DWT (~ 3.3 times in 2-D, ~ 5.1 times in 3-D)

CONCLUSIONS (cont.)

- A **general purpose multi-resolution front-end**, similar to the multi-scale Gabor-like filters of the **human V1 cortex**, suitable for many image analysis and reconstruction tasks:
  - Enhancement (deconvolution)
  - Denoising
  - Motion / displacement estimation and compensation
  - Texture analysis / synthesis
  - Segmentation and classification
  - Watermarking
  - 3D data enhancement and visualisation
  - Object recognition and image understanding
  - Sparsity-based image & 3D reconstruction

Papers on complex wavelets are available at: **www.eng.cam.ac.uk/~ngk/**

A Matlab DT $\mathbb{C}$WT toolbox is available on request from: **ngk10@cam.ac.uk**