The Red-Black Wavelet Transform

Geert Uytterhoeven *Adhemar Bultheel*

Department of Computer Science
Katholieke Universiteit Leuven
Celestijnenlaan 200A, B-3001 Heverlee, Belgium
Phone +32-16-327080 Fax +32-16-327996
E-mail Geert.Uytterhoeven@cs.kuleuven.ac.be

Abstract

We present a new kind of second generation wavelets on a rectangular grid. These wavelets are constructed using a 2D lifting scheme which is based on a red-black blocking scheme. Compared to classical tensor product wavelets on the same grid, these new wavelets show less anisotropy. The performance of the new wavelets was compared to tensor product wavelets in an image denoising application.

1 INTRODUCTION

Classical one-dimensional wavelet transforms can be extended to more dimensions using tensor products, yielding a separable multi-dimensional transform. A disadvantage of this technique is the introduction of an anisotropy in the wavelet decomposition. In the two-dimensional case, a tensor product wavelet transform will favor horizontal, vertical and diagonal features of the original data. Other features are not that easily detected.

Non-separable wavelets can provide a solution to this. In this paper, we present a new kind of second generation wavelets on a rectangular grid — more specifically, on a quincunx lattice — constructed using the lifting scheme.

2 THE LIFTING SCHEME

The lifting scheme (e.g. [5]) is an algorithm to calculate wavelet transforms in an efficient way. It found its roots in a method to improve a given wavelet transform to obtain some specific properties. Later it was extended to a generic method to create so-called 'Second Generation' wavelets. In this paper, we design a new family of second generation wavelets on a rectangular grid.

One of the great advantages of the lifting scheme realization of a wavelet transform is that it decomposes the wavelet filters into extremely simple elementary steps, and each of these steps is very easily invertible. As a result, the inverse transform can always be obtained immediately from the forward transform.

In practice, discrete signals are represented by integers. Another advantage of the lifting scheme is that a discrete wavelet transform computation using the lifting scheme can easily be converted to a transform that maps integers to integers, while retaining the perfect reconstruction property[1, 6].

2.1 Predict and Update

At each resolution level of a wavelet transform, the signal is split into a high pass and a low pass part and the low pass part is split again etc. These high pass and low pass parts are obtained by applying corresponding wavelet filters.

The lifting scheme is an efficient implementation of these filtering operations. Suppose that the low resolution part of a signal at level $j+1$ consists of a data set λ_{j+1}. This set is transformed into two other sets at level j: the low resolution part λ_j and the high resolution part γ_j. This is obtained first by just splitting the data set λ_{j+1} into two separate data subsets λ_j and γ_j (e.g. the even samples and the odd samples). Such a splitting is sometimes referred to as the *lazy wavelet transform*. Then these two sets are recombined in several subsequent pairs of lifting steps which decorrelate the two signals.

A *dual lifting* step can be seen as a prediction: the data γ_j are 'predicted' from the data in the subset λ_j. When the signals are still highly correlated, then such a prediction will usually be very good, and thus we do not have to keep this information in both signals. That is why we can keep λ_j and store only the part of γ_j that is not predictable (the prediction error). Thus γ_j is replaced by $\gamma_j - \mathcal{P}(\lambda_j)$ where $\mathcal{P}$ represents the prediction operator. This is the real decorrelating step.

However, the new representation has lost certain basic properties, which one usually wants to keep, like for example the mean value of the signal. To restore this property, one needs a *primal lifting* step, whereby the set λ_j is updated with data computed from the (new) subset γ_j. Thus λ_j is replaced by $\lambda_j + \mathcal{U}(\gamma_j)$ with $\mathcal{U}$ some updating operator.

In general, several such lifting steps can be applied in sequence to go from level $j+1$ to level j.

To recapitulate, let us consider a simple lifting scheme with only one pair of lifting steps.

Splitting (*lazy wavelet transform*) Partition the data set

λ_{j+1} into two distinct data sets λ_j and γ_j.

Prediction (*dual lifting*) Predict the data in the set γ_j by the data set λ_j: $\gamma_j \leftarrow \gamma_j - \mathcal{P}(\lambda_j)$.

Update (*primal lifting*) Update the data in the set λ_j by the data in set γ_j: $\lambda_j \leftarrow \lambda_j + \mathcal{U}(\gamma_j)$.

These steps can be repeated by iteration on the λ_j, creating a multi-level transform or multi-resolution decomposition.

The inversion rules are obvious: revert the order of the operations, invert the signs in the lifting steps, and replace the splitting step by a merging step. Thus, inverting the three step procedure above results in

Inverse update $\lambda_j \leftarrow \lambda_j - \mathcal{U}(\gamma_j)$,

Inverse prediction $\gamma_j \leftarrow \gamma_j + \mathcal{P}(\lambda_j)$,

Merge $\lambda_{j+1} \leftarrow \lambda_j \cup \gamma_j$.

2.2 Integer Transforms

In practice, discrete signals are represented by integers. Doing the filtering operations on these numbers however will transform them in rational or real numbers because the filter coefficients need not be integers. To obtain an efficient implementation of the discrete wavelet transform, it is of great practical importance that the wavelet transform is represented by a set of integers as well.

We shall round the intermediates of each lifting step to integers (for example the nearest integer) and indicate this operation by square braces. Thus, we actually compute rounded values:

$$\gamma_j \leftarrow \gamma_j - [\mathcal{P}(\lambda_j)]\,; \lambda_j \leftarrow \lambda_j + [\mathcal{U}(\gamma_j)]\,.$$

It is not difficult to verify that each step of the lifting scheme with rounding is perfectly invertible and thus the whole signal is perfectly reconstructable, whatever the rounding rule we use, on condition of course that the rounding is deterministic.

2.3 Example: Cohen-Daubechies-Feauveau

The popular family of classical biorthogonal wavelets constructed by Cohen, Daubechies and Feauveau [2] fits in the above scheme. Especially its member with two vanishing moments for both the primal and dual wavelet (hence named CDF $(2, 2)$ wavelet) is widely used.

Thanks to the lifting scheme, the accompanying wavelet transform can be implemented in an efficient way (see e.g. [6]). From the second generation viewpoint, one transform step of a discrete signal $x = \{x_k\}$ looks like:

Splitting Split the signal x (i.e. λ_{j+1}) into even samples (i.e. λ_j) and odd samples (i.e. γ_j):

$$s_i \leftarrow x_{2i}; d_i \leftarrow x_{2i+1}.$$

Prediction Predict the odd samples using linear interpolation:

$$d_i \leftarrow d_i - \frac{1}{2}(s_i + s_{i+1}).$$

Update Update the even samples to preserve the mean value of the samples:

$$s_i \leftarrow s_i + \frac{1}{4}(d_{i-1} + d_i).$$

Note that this transform works on one-dimensional data. For two-dimensional data, it can be applied row and columnwise, resulting in a tensor product wavelet transform. This is a separable transform.

3 THE RED-BLACK WAVELET TRANSFORM

In the previous section, we presented the principle of constructing wavelet transforms by the lifting scheme in a very general setting. It was not assumed that the data sets were one-dimensional for example. The same principle and the same description can be applied to two or more dimensional data as well. As an illustration, we apply these ideas to the two-dimensional case in its simplest possible appearance: a two-dimensional analogue of the one-dimensional CDF (2, 2) wavelet. We assume that the data are given on a regular rectangular grid. It should be clear form our previous exposition, that this is about the simplest nontrivial case that can be conceived, and that much more complex schemes are possible. However, because in the one-dimensional case, the CDF (2, 2) wavelets are widely used, we are convinced that the two-dimensional analogue that we present below has the same potentials as the the CDF (2, 2) wavelet.

3.1 A Two-Step Method

Our approach is inspired by the well-known Red-Black Gauss-Seidel technique for the iterative solving of linear systems. In the literature [4], the kind of lattice we use is known as a quincunx lattice. However, we prefer the name Red-Black wavelet transform because it is simpler and it is more appropriate to describe the splitting step in the lifting scheme.

The idea is that we first split the data in two subsets. The even/odd splitting from the one-dimensional case is replaced by a checkerboard splitting, with red and black squares. Thus we have a red subset λ_j and a black subset γ_j. Then we predict the values of the black subset γ_j from its immediate neighbors in the red set. In accordance with the CDF (2, 2) lifting scheme, we take for this prediction an average of the four neighbors. The elements in γ_j are replaced by their prediction errors. Next, we have to make an updating step for the data in the red subset λ_j, which is based on the (new) black values, to preserve the mean value of the data.

For the subsequent resolution level $j-1$, we are left with the (new) red set λ_j from the previous level. However, it is clear by looking at Figure 1, that these data are arranged along diagonals on the grid. Thus we cannot repeat the same operation on a horizontal/vertical basis as on the previous resolution level, but we can do it on the diagonals instead. Again we start by splitting λ_j into two

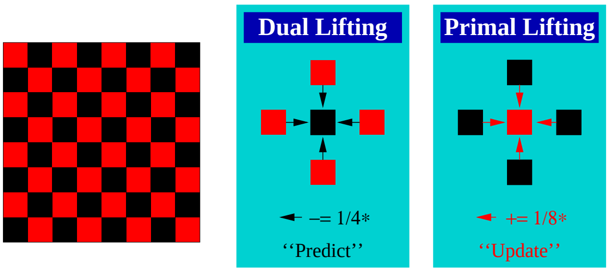

Figure 1: Horizontal/vertical lifting in the odd steps: *Red-Black*.

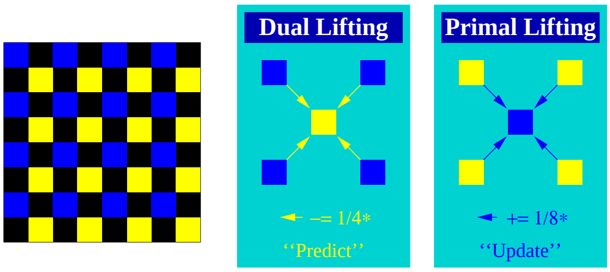

Figure 2: Diagonal lifting in the even steps: *Blue-Yellow*.

subsets: the blue subset λ_{j-1} and the yellow subset γ_{j-1} (see Figure 2). The same kind of averaging and updating steps are then computed for the blue-yellow squares. For the next resolution level, we have to decompose the blue subset further, but since this subset is again arranged in a horizontal/vertical manner, we can do again a red-black update. It is seen that we will have alternatingly a red-black and a blue-yellow splitting to pass through the subsequent resolution levels. Therefore we say that the Red-Black wavelet transform is a 'two level' transform.
We give now the formulas which are based on these ideas in a more schematic presentation. Since we are working on a rectangular grid, the input data set λ_{j+1} is an image, i.e. a rectangular matrix $X = x_{i,j}$.

3.1.1 *Horizontal/Vertical Lifting*

Splitting: The image λ_{j+1} is split in the *red* squares λ_j, and the *black* squares γ_j (cfr. Figure 1).

Prediction: The black squares are predicted using linear interpolation based on the 4 neighboring red squares:

$$x_{i,j} \leftarrow x_{i,j} - \frac{1}{4}\left(x_{i-1,j} + x_{i,j-1} + x_{i,j+1} + x_{i+1,j}\right),$$

for $i \bmod 2 \neq j \bmod 2$.

Update: The red squares are updated using the black squares to preserve the mean value:

$$x_{i,j} \leftarrow x_{i,j} + \frac{1}{8}\left(x_{i-1,j} + x_{i,j-1} + x_{i,j+1} + x_{i+1,j}\right),$$

for $i \bmod 2 = j \bmod 2$.

The inverse transform is straightforward.

3.1.2 *Diagonal Lifting*

For this step we consider the checkerboard to be rotated over $45°$.

Splitting The *red* squares in Figure 1 (the data λ_j) are partitioned in the *blue* (the data λ_{j-1}) and *yellow* (the data γ_{j-1}) squares in Figure 2.

Prediction The yellow squares are predicted using linear interpolation based on the 4 neighboring blue squares:

$$x_{i,j} \leftarrow x_{i,j} - \frac{1}{4}\left(\begin{array}{c} x_{i-1,j-1} + x_{i-1,j+1} + \\ x_{i+1,j-1} + x_{i+1,j+1} \end{array}\right),$$

for $i \bmod 2 = 1$ and $j \bmod 2 = 1$.

Update The blue squares are updated using the yellow squares to preserve the mean value:

$$x_{i,j} \leftarrow x_{i,j} + \frac{1}{8}\left(\begin{array}{c} x_{i-1,j-1} + x_{i-1,j+1} + \\ x_{i+1,j-1} + x_{i+1,j+1} \end{array}\right),$$

for $i \bmod 2 = 0$ and $j \bmod 2 = 0$.

Again, the inverse transform is straightforward.

The next step (horizontal/vertical lifting, again) will be performed on the blue squares only, yielding a multi-resolution decomposition.

3.2 Integer Wavelet Transform

Because the Red-Black wavelet transform is based on the lifting scheme, a version that maps integers to integers is immediately deduced: all computations giving a floating point number are rounded to an integer, e.g.

$$\mu \leftarrow \mu - \frac{1}{4}\left(\rho + \sigma + \tau + \upsilon\right)$$

becomes

$$\mu \leftarrow \mu - \left[\frac{1}{4}\left(\rho + \sigma + \tau + \upsilon\right)\right]$$

and similarly for all the other steps.

3.3 Basis Functions

In a classical one-dimensional wavelet transform, all basis functions are derived from two functions — the mother and the father wavelet — by dilation and translation. In the case of one-dimensional bi-orthogonal wavelets, there is also a dual basis. However, in our construction we use different wavelet transforms for the even and odd steps. As a consequence all basis functions are derived from three[1] functions:

$\psi_{\text{low}}(x,y)$ Basis function corresponding to the low pass part of the diagonal lifting step.

$\psi_{\text{mid}}(x,y)$ Basis function corresponding to the high pass part of the diagonal lifting step.

$\psi_{\text{high}}(x,y)$ Basis function corresponding to the high pass part of the horizontal/vertical lifting step.

Graphs of these basis functions (for the integer variant of the lifting scheme) are shown in Figure 3.

[1] We do not consider the intermediate basis functions corresponding to the low pass parts after the horizontal/vertical lifting steps.

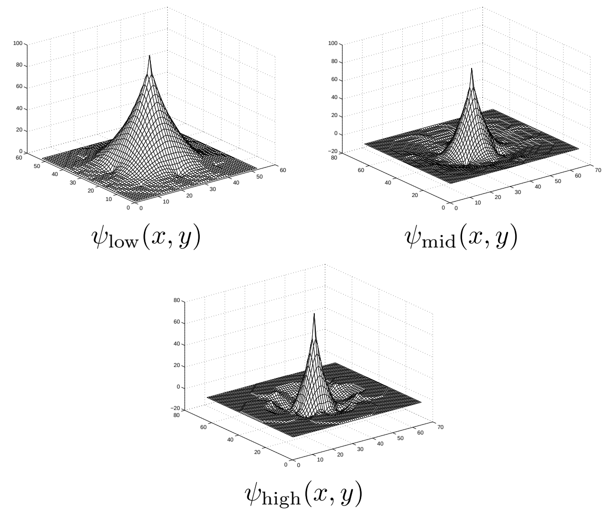

Figure 3: Graphs of the basis functions for the integer Red-Black wavelet transform.

3.4 Other Members of the Family

As we already mentioned at the beginning of this section, we used in the description above very simple prediction and update 'molecules'. The molecules have the form of a cross: a normal cross for the odd steps and a diagonal cross for the even steps. These give us wavelet functions with 2 vanishing moments, cfr. the Cohen-Daubechies-Feauveau biorthogonal wavelets. By using larger molecules and more advanced prediction and update formulas, one can create smoother basis functions with larger supports.

4 PROPERTIES OF THE RED-BLACK WAVELET TRANSFORM

The Red-Black wavelet transform has the following properties:

2D: The Red-Black wavelet transform is a non-separable true two-dimensional transform.

Isotropy: It is less anisotropic than tensor product wavelets on the same rectangular grid.
The edges which are manifest in the figures of the tensor product wavelets show the pronounced vertical/horizontal orientation, while this is not visible in the figures of the Red-Black wavelets.

Symmetry: All basis functions have four symmetry axes: horizontal, vertical, and diagonal.

Biorthogonality: The splitting in red and black or blue and yellow pixels creates trivial biorthogonal basis functions. Since lifting steps preserve the biorthogonality properties the transform as a whole has biorthogonal basis functions, too [5].

Smoothness: The wavelet functions (both primal and dual) have two vanishing moments: constant, linear and bi-linear functions can be represented exactly by the scaling functions.

Implementation: Both the forward and inverse transform are very simple to implement and have fast execution times ($\mathcal{O}(n)$, with n the number of elements in the grid), due to the underlying lifting technology.

5 CONCLUSION

Red-Black wavelets are less anisotropic than tensor product wavelets. Compared to a similar tensor product wavelet, image denoising using Red-Black wavelets performs better for lines that are not horizontal, vertical or diagonal. Besides image processing, other possible applications are situated in partial differential equations.
The 'two level' principle of the Red-Black wavelet transform is not restricted to a quincunx lattice. It can also be extended to triangular or hexagonal regular grids: an update and predict pair on a triangular grid will yield a hexagonal low pass grid, and vice versa. Note that one is not limited to planar grids: (semi-)regular grids of triangles or hexagons on a sphere are also possible.

Acknowledgements

This paper presents research results of the Belgian Programme on Interuniversity Poles of Attraction, initiated by the Belgian State, Prime Minister's Office for Science, Technology and Culture. The scientific responsibility rests with its authors. This research is also supported by the Flemish Information Technology Action Program ('Vlaams Actieprogramma Informatietechnologie'), project number ITA/950244.
After finishing this work, we learned that Kovačević and Sweldens have completed a similar construction[3].

References

[1] R. Calderbank, I. Daubechies, W. Sweldens, and B.-L. Yeo. Losless image compression using integer to integer wavelet transforms. In *International Conference on Image Processing (ICIP), Vol. I*, pages 596–599. IEEE Press, 1997.

[2] A. Cohen, I. Daubechies, and J.C. Feauveau. Biorthogonal bases of compactly supported wavelets. *Comm. Pure Appl. Math.*, 45(5):485–560, 1992.

[3] J. Kovačević and W. Sweldens. Wavelet families of increasing order in arbitrary dimensions. *IEEE Trans. Image Processing*, 9(3):480–496, March 2000.

[4] J. Kovačević and M. Vetterli. Nonseparable multidimensional perfect reconstruction filter banks and wavelet bases for $\mathbf{R}^n$. *IEEE Trans. Inf. Th.*, 38(2):533–555, March 1992.

[5] W. Sweldens. The lifting scheme: A construction of second generation wavelets. *SIAM J. Math. Anal.*, 27(2), 1997.

[6] G. Uytterhoeven, F. Van Wulpen, M. Jansen, D. Roose, and A. Bultheel. WAILI: Wavelets with Integer Lifting. Technical Report TW262, Department of Computer Science, K.U. Leuven, July 1997.