

The Red-Black Wavelet Transform and the Lifting Scheme

Geert Uytterhoeven Adhemar Bultheel

Report TW 318, December 11, 2000



Katholieke Universiteit Leuven

Department of Computer Science

Celestijnenlaan 200A – B-3001 Heverlee (Belgium)

The Red-Black Wavelet Transform and the Lifting Scheme

Geert Uytterhoeven Adhemar Bultheel

Report TW 318, December 11, 2000

Department of Computer Science, K.U.Leuven

Abstract

We recall how the lifting idea can be used to construct non-separable wavelets on the quincunx grid but we also illustrate how this idea leads to the construction of a new kind of second generation wavelets on grids that are alternately hexagonal and triangular. 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.

Keywords : wavelets, lifting, image compression.

AMS(MOS) Classification : 41A30, 68U10, 94A12, 65Y05.

The Red-Black Wavelet Transform and the Lifting Scheme

Geert Uytterhoeven* Adhemar Bultheel†

December 11, 2000

Abstract

We recall how the lifting idea can be used to construct nonseparable wavelets on the quincunx grid but we also illustrate how this idea leads to the construction of a new kind of second generation wavelets on grids that are alternatingly hexagonal and triangular. 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.

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. They allow for the construction of two-dimensional wavelets on e.g. lattices (e.g. on the so called ‘quincunx’ lattice [6, 4]) or on hexagonal grids [10, 8, 7, 2]. Almost all of these are heavily based on Fourier transforms, just like the classical wavelets.

Second generation wavelets designed using the lifting scheme are another option. The lifting scheme is a generic method to create wavelets on intervals, irregular samplings, meshes, manifolds, . . . , without relying on the Fourier transform. The classical one-dimensional wavelets can be constructed using lifting, too. Therefore one has to decompose the filters into lifting steps. However, one can design new transforms directly using the lifting technique. In that case the reconstruction of the underlying filters will require some effort. The lifting idea has been well established now for one dimensional regular grids. For irregular grids some numerical problems have yet to be coped with. For the regular grids in the 2 or more dimensional case, the separable wavelets are easy generalizations of the 1D case. For an increased symmetry however one needs nonseparable wavelets and in that case only little has

*Sony Software Development Center Europe (SDCE), Sint-Stevens-Woluwestraat 55, B-1130 Brussels, Belgium, Geert.Uytterhoeven@sonycom.com, <http://home.tvd.be/cr26864/>

†Department of Computer Science, Katholieke Universiteit Leuven, Celestijnenlaan 200A, B-3001 Heverlee, Belgium, Adhemar.Bultheel@cs.kuleuven.ac.be

been published concerning the lifting scheme in the 2D case. The general ideas however are exposed in the recent paper. [5].

In this text, we present a new kind of second generation wavelets on a hexagonal grid constructed using the lifting scheme. The idea is to apply the Red-Black lifting technique that we described in [15] for a rectangular grid also on hexagonal grids. It is a direct generalization of the construction of the CFD(2,2) (Cohen-Daubechies-Fauveau wavelets with 2 primal and 2 dual vanishing moments) wavelets in the 1-dimensional case.

2 The Lifting Scheme

2.1 Introduction

The lifting scheme 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 the 1D case, these are wavelets that need not be obtained as dilates and translates of one function. They are much more flexible and can be used to define a wavelet basis on an interval or on an irregular grid, or even on a sphere. The lifting scheme can also be used to introduce wavelets without using the concept of Fourier transform. Several introductions to the lifting scheme are available [12, 13, 9, 11].

Second generation wavelets are more general in the sense that all the classical wavelets can be generated by the lifting scheme. In fact, the decomposition of a classical wavelet filter into lifting steps can be easily obtained via the Euclidean algorithm [3].

2.2 Predict and Update

The wavelet transform of a 1D signal is a multiresolution representation of that signal where the wavelets are the basis functions which at each resolution level give a highly decorrelated representation. Thus at each level, 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. In general these filters are coupled if certain conditions are to be fulfilled, like for example perfect reconstruction.

The lifting scheme is an efficient implementation of these filtering operations at each level when computing a discrete wavelet transform. So suppose that the low resolution part of a signal at level $j + 1$ is given and that it consists of a data set which we represent by 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 . Traditionally this is done by separating the set of even samples and the set of odd samples. Such a splitting is sometimes referred to as the lazy wavelet transform. Doing just this of course does not improve our representation of the signal. Therefore, the next step is to recombine these two sets in several subsequent lifting steps which decorrelate the two signals.

Lifting steps usually come in pairs of a primal and a dual lifting step. 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 γ_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 .

The principles that were just explained become very clear in the case of the simple CFD(2, 2) transform that we give below in Section 2.4. The odd samples are predicted by linear interpolation of the neighbouring even samples and are therefore replaced by the interpolation error. The even samples are updated to preserve the mean value of the signal.

To recapitulate, let us consider a simple lifting scheme with only one pair of lifting steps. Then to go from level $j + 1$ to level j one has to perform the following three 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.

2.3 The Inverse Transform

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


2.4 Example: Cohen-Daubechies-Feauveau

One popular family of classical biorthogonal wavelets that fits in the above scheme are the wavelets constructed by Cohen, Daubechies and Feauveau [1]. 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. [16]). 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):

$$\begin{aligned} s_i &\leftarrow x_{2i}, \\ d_i &\leftarrow x_{2i+1} \end{aligned}$$




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:

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



As a result, the signal $s = \{s_k\}$ is a coarse representation of the original signal x , while the signal $d = \{d_k\}$ contains the high frequency information that is lost when going from resolution level $j + 1$ to resolution level j .

The same kind of interpolation scheme can be used for interpolating polynomials of higher degree. These lifting schemes need some special care at the boundaries if one wants the wavelets to live explicitly on the discrete set where the data are defined. A whole family of lifting schemes can be constructed in this way [14], of which the above example is just the simplest possible case.

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 for two or more dimensional data as well. As an illustration, we apply these ideas to a two-dimensional case: a two-dimensional analog of the one-dimensional CDF (2, 2) wavelet. In [15] we used an approach inspired by the well-known Red-Black Gauss-Seidel technique for the iterative solving of linear systems, on a quincunx lattice, and called it the ‘Red-Black wavelet transform’.

However, the principle behind the Red-Black wavelet transform is not restricted to quincunx lattices. It can be extended easily to other grids. In this paper we will apply it to planar hexagonal and triangular grids. Note that it is not limited to planar grids: (semi-)regular grids on a sphere are also possible.

3.1 A Quincunx Grid

Let us first recall the Red-Black technique in the case of a regular rectangular grid. There are two stages: a Horizontal-Vertical lifting and a diagonal lifting.

3.1.1 Horizontal-Vertical lifting

Splitting: The even-odd splitting from the 1D case is now changed into a checkerboard splitting. In accordance with the the Gauss-Seidel Red-Black technique we shall call the squares here also red (the subset λ_j) and black (the subset γ_j). See Figure 1.

Prediction: We have to predict the red values by the neighboring black ones. As for the CFD(2,2) wavelet, we take for the prediction of a black value, the average of the 4 closest red values and we keep the difference. Thus for the red values

$$\text{black}_i \leftarrow \text{black}_i - \frac{1}{4} \left(\sum_{j \in \mathcal{K}(i)} \text{red}_j \right)$$

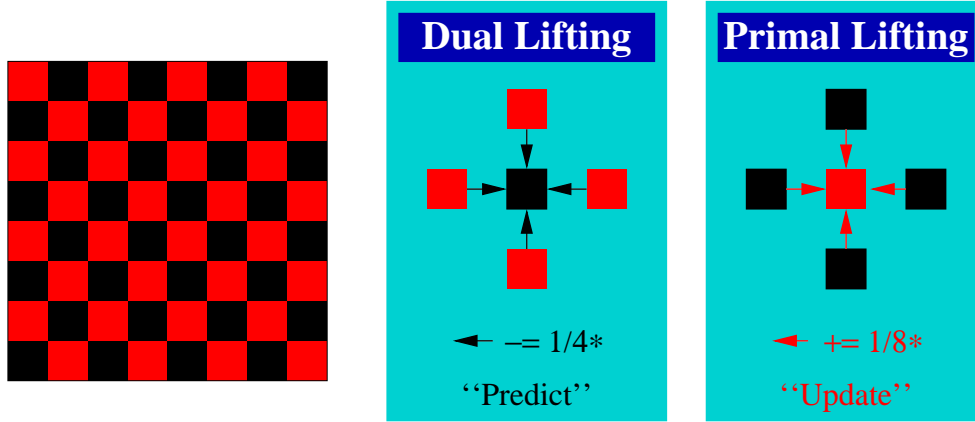


Figure 1: Quincunx Red-Black splitting and lifting steps.

where $\mathcal{K}(i)$ is the index-set of the 4 nearest red neighbours of black_i . Note that x_{kl} is black, i.e., $x_{k,l} \in \gamma_j$ iff $k \bmod 2 \neq l \bmod 2$

Update: For the update step, we arrange that the mean value remains unchanged. This is obtained by choosing the following update for the black values:

$$\text{red}_i \leftarrow \text{red}_i + \frac{1}{8} \left(\sum_{j \in \mathcal{R}(i)} \text{red}_j \right)$$

where $\mathcal{R}(i)$ is the index-set of the 4 nearest black neighbours of red_i . Note that x_{kl} is red, i.e., $x_{k,l} \in \lambda_j$ iff $k \bmod 2 = l \bmod 2$

3.1.2 Diagonal lifting

For the next stage, we have to continue with the red values (λ_j) which we split into blue and yellow ones. We have again a checkerboard but now the values are arranged diagonalwise.

Splitting: The splitting is obvious and illustrated of Figure 2.

Prediction: The yellow value is predicted by the average of the 4 nearest blue ones and again the difference is kept.

$$\text{yellow}_i \leftarrow \text{yellow}_i - \frac{1}{4} \left(\sum_{j \in \mathcal{Y}(i)} \text{blue}_j \right)$$

where $\mathcal{Y}(i)$ is the index-set of the 4 nearest blue neighbours of yellow_i . Note that x_{kl} is yellow, iff $k \bmod 2 = 1 = l \bmod 2$.

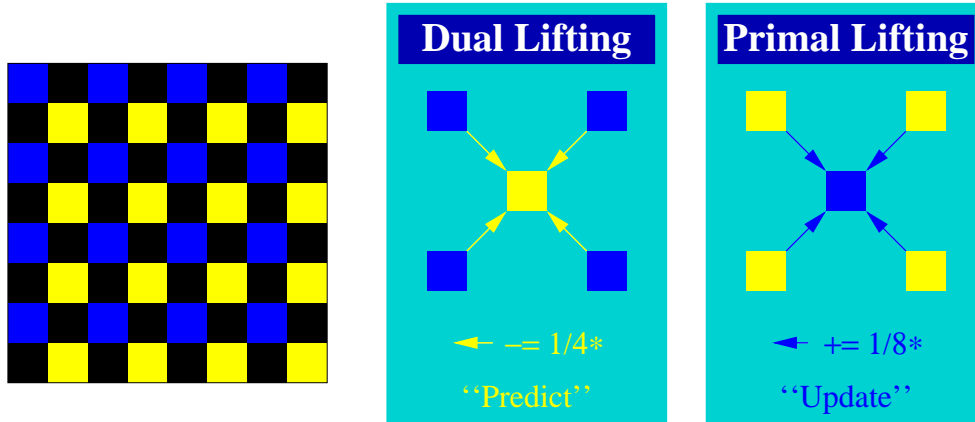


Figure 2: Quincunx Blue-Yellow splitting and lifting steps.

Update: We modify the blue values to keep the average:

$$\text{blue}_i \leftarrow \text{blue}_i + \frac{1}{8} \left(\sum_{j \in \mathcal{B}(i)} \text{yellow}_j \right)$$

where $\mathcal{B}(i)$ is the index-set of the 4 nearest yellow neighbours of blue_i . Note that x_{kl} is blue, iff $k \bmod 2 = 0 = l \bmod 2$.

Now one is back in the Horizontal-Vertical situation and these two stages can be repeated.

3.2 A Regular Hexagonal Grid

The idea, just like in the CDF(2, 2) wavelet, is that we first split the data in two subsets. The quincunx case that we considered above already shows that there is a two-stage algorithm needed to work through because the configuration of the red subset need not be the same as the configuration of the set one started with.

If we start out with the hexagonal configuration, then the splitting into a red and a black subset is rather obvious as shown in Figure 3. Note that the red and black subsets are on triangular, not hexagonal, grids.

Thus we have a red subset λ_j and a black subset γ_j . Then we predict the values of the red subset λ_j from its three immediate neighbors in the black set. In accordance with the CDF (2, 2) lifting scheme, we take for this prediction an average of these three neighbors. The elements in λ_j are replaced by their prediction errors. Next, we have to make an updating step for the data in the red subset γ_j , which is based on the (new) black values, to preserve the mean value of the data.

Splitting: The set λ_{j+1} is split in the red subset λ_j , and the black subset γ_j .




Figure 3: Red-Black splitting and lifting steps.

Prediction: The black samples are predicted using linear interpolation based on the three neighboring red samples:

$$\text{black}_i \leftarrow \text{black}_i - \frac{1}{3} \left(\sum_{j \in \mathcal{K}(i)} \text{red}_j \right),$$

with $\mathcal{K}(i)$ the set of three nearest red neighbors of black_i .

Update: The red samples are updated using the three neighboring black samples to preserve the mean value:

$$\text{red}_i \leftarrow \text{red}_i + \frac{1}{6} \left(\sum_{j \in \mathcal{R}(i)} \text{black}_j \right),$$

with $\mathcal{R}(i)$ the set of three nearest black neighbors of red_i .

As a result, because the averaging operation smooths the data, the values corresponding to the red samples are a low resolution representation of the original grid, while the black samples contain the detail information. The inverse transform is straightforward.

3.3 A Regular Triangular Grid

For the subsequent resolution level $j - 1$, we are left with the (new) red set λ_j from the previous level. However, as mentioned before, these red data are arranged on a triangular grid. Thus we cannot repeat the same operation as on the previous resolution level. Instead

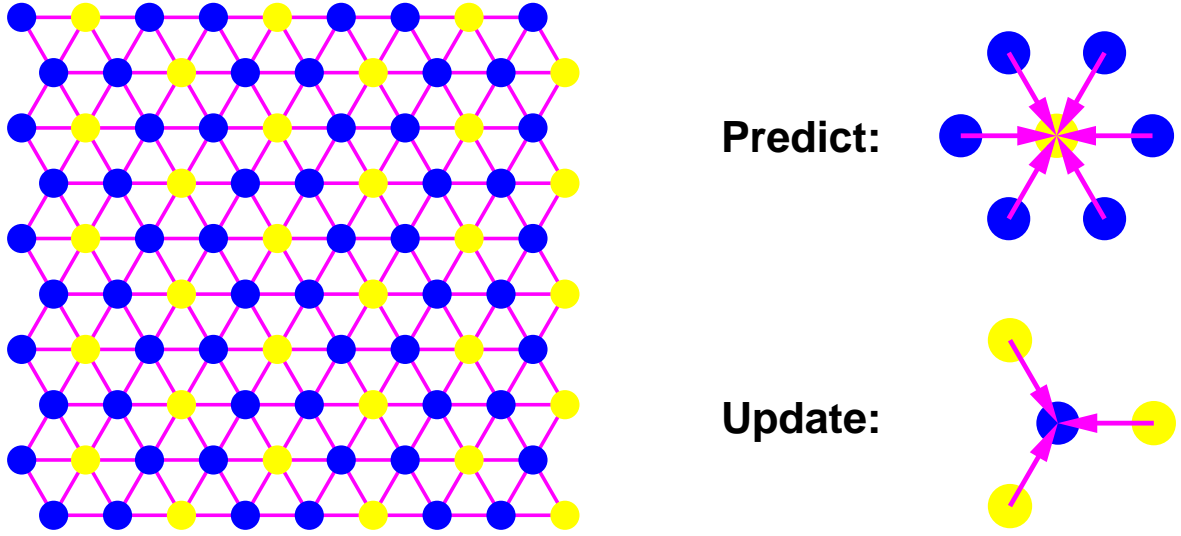


Figure 4: Blue-Yellow splitting and lifting steps.

we chose the following approach. Again we start by splitting λ_j into two subsets: the blue subset λ_{j-1} and the yellow subset γ_{j-1} (see Figure 4). Note that this is not a balanced splitting: the blue subset is twice as large as the yellow subset. The blue subset is on a regular hexagonal grid.

Then we predict the values of the blue subset λ_j from its six immediate neighbors in the yellow set. In accordance with the CDF (2, 2) lifting scheme, we take for this prediction an average of these six neighbors. The elements in λ_j are replaced by their prediction errors. Next, we have to make an updating step for the data in the blue subset γ_j , which is based on the (new) yellow values, to preserve the mean value of the data.

Splitting The red set in Figure 3 (the data λ_j) is redrawn in Figure 4 and is partitioned in the blue (the data λ_{j-1}) and yellow (the data γ_{j-1}) subsets.

Prediction The yellow samples are predicted using linear interpolation based on the six neighboring blue samples:

$$\text{yellow}_i \leftarrow \text{yellow}_i - \frac{1}{9} \left(\sum_{j \in \mathcal{Y}(i)} \text{blue}_j \right),$$

with $\mathcal{Y}(i)$ the set of six nearest blue neighbors of yellow_i .

Update The blue samples are updated using the three neighboring yellow samples to preserve the mean value:

$$\text{blue}_i \leftarrow \text{blue}_i + \frac{1}{6} \left(\sum_{j \in \mathcal{B}(i)} \text{yellow}_j \right),$$

with $\mathcal{B}(i)$ the set of three nearest yellow neighbors of blue_i .

Result: the blue samples (λ_{j-1}) are a low resolution representation of the original grid, while the yellow samples contain the detail information. Again, the inverse transform is straightforward.

3.4 More Resolution Levels

For the next resolution level, we have to decompose the blue subset further, but since this subset is again located on a hexagonal grid, we can do again a red-black update. Note that the hexagonal grid is rotated over 30° degrees, compared to our first step. It is seen that we will have alternately a red-black and a blue-yellow splitting to pass through the subsequent resolution levels.

3.5 In-place Transform

The above transform is 100% in-place and all coefficients will be interleaved. This is one of the powerful features of the lifting scheme: it does not require additional memory to calculate the forward or inverse transform.

3.6 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 steps¹. As a consequence all basis functions are derived from three² functions. These three functions can be obtained by choosing a blue, a yellow or a red sample to take a nonzero value and setting all other coefficients equal to zero. Then the inverse transform is applied to these data, in principle for an infinite number of levels, which results in the basis functions .

Graphs of these basis functions scheme are shown in Figures 5 and 6.

3.7 Properties of the Red-Black Wavelets

Red-Black wavelets are not constructed as tensor-products of classical one-dimensional wavelets.

Because a regular hexagonal grid has more symmetry than a square or rectangular grid, we expect these Red-Black wavelets to be less anisotropic than their counterparts on a rectangular grid [15].

Unfortunately the Blue-Yellow step disturbs the symmetry, because of two reasons:

1. A triangular grid has less symmetry than a square or rectangular grid.

¹Due to the rotation about 90° after the Blue-Yellow step, we have 2 more basis functions, but these are simple dilates and rotates of the 2 other functions

²We do not consider the intermediate basis functions corresponding to the low pass parts after the red-black/blue-yellow lifting steps.




Figure 5: Graphs of the basis functions for the Red-Black wavelet transform: (a) 3D view and (b) top view of the basis function corresponding to the low-pass part of the Red-Black step. The spiral behavior is clearly visible in the top view.




Figure 6: Graphs of the basis functions for the Red-Black wavelet transform: 3D views of the basis functions corresponding to the high-pass parts of the (a) Red-Black and (b) Blue-Yellow steps.

2. We perform a non-balanced splitting in this step: the blue ('low pass') subset is twice as large as the yellow ('high pass') subset.

As a consequence the basis functions are not as symmetric as one would expect. They even show some spiral behavior (cfr. Figure 5b).

4 Relation with other hexagonal wavelets

When compared with existing literature [10, 8, 7, 2, 5] about hexagonal wavelets, our approach is somewhat different.

First we start from a truly hexagonal grid on which the data are given. In [2] for example, the grid is actually triangular (like in our Blue-Yellow step) and the term hexagonal is used because every grid point has 6 nearest neighbours. The grid considered there is in fact given by $\Gamma\mathbb{Z}^2$ with

$$\Gamma = \begin{bmatrix} 1 & -1/2 \\ 0 & \sqrt{3}/2 \end{bmatrix}.$$

This corresponds to our grid for the Blue-Yellow step. The hexagonal grid of the Red-Black step is a subgrid of the same triangular grid but rotated over 30° (or 90° which is the same, given the symmetry of the problem).

The subsampling considered in [2] is separable because given by a diagonal matrix $D = \text{diag}(2, 2)$. Here we have for the Blue-Yellow step a nonseparable subsampling given by the matrix

$$D = \begin{bmatrix} 3 & 1 \\ 0 & -1 \end{bmatrix}.$$

Which means that we have $|\det D| - 1 = 2$ cosets given by shifted versions of the subgrid $D\mathbb{Z}^2$ given by the shifts $\tau_i = (i, 0)^T$, $i = 1, 2$. And hence

$$\Gamma\mathbb{Z}^2 = \bigcup_{i=0}^2 \Gamma(D\mathbb{Z}^2 + \tau_i)$$

where $\tau_0 = (0, 0)^T$. For the Red-Black stage, everything is rotated over 30 degrees, but apart from that, the subsampling is on the red subgrid $D\mathbb{Z}^2$ with now the subsampling matrix

$$D = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}.$$

Which has again $|\det D| - 1 = 2$ cosets with shifts $\tau_i = (i, 0)^T$, $i = 1, 2$. One of these is the black subgrid, the other one collects the centers of the hexagons which do not contain data in our setting.

Because of these rotations over 30 degrees everytime one switches from a Red-Black grid to a Blue-Yellow grid, this has a spiralling effect that can be seen in the graphs of the basis functions.

The explicit construction of the filters is cumbersome and is not needed for implementations reasons. So we leave it to the inspiration of the reader.

5 Conclusion

Red-Black wavelets give an extremely simple computational scheme that generalizes the idea of the CDF(2,2) wavelets to two-dimensional grids. Not only can this idea be used to develop nonseparable wavelets on a rectangular (quincunx) grid, but also on regular hexagonal and triangular grids. The idea is inspired by the well known Red-Black structure used in Gauss-Seidel methods for the iterative solution of equations. It is clear that because of the symmetry in the hexagonal case, the basis functions are less anisotropic than tensor product wavelets and they are more so in the hexagonal case than in the rectangular case.

Note that one is not limited to planar grids: (semi-)regular grids of triangles or hexagons on a sphere are also possible.

6 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.

References

- [1] A. Cohen, I. Daubechies, and J.C. Feauveau. Biorthogonal bases of compactly supported wavelets. *Comm. Pure Appl. Math.*, 45(5):485–560, 1992.
- [2] A. Cohen and J.-M. Schlenker. Compactly supported bidimensional wavelet bases with hexagonal symmetry. *Constr. Approx.*, 9:209–236, 1993.
- [3] I. Daubechies and W. Sweldens. Factoring wavelet transforms into lifting steps. *J. Fourier Anal. Appl.*, 4:245–267, 1998.
- [4] B. Han and R.-Q. Jia. Quincunx fundamental refinable functions and quincunx biorthogonal wavelets. *Math. Comp.*, 2000. To appear.
- [5] J. Kovačević and W. Sweldens. Wavelet families of increasing order in arbitrary dimensions. *IEEE Trans. Image Processing*, 9(3):480–496, March 2000.
- [6] J. Kovačević and M. Vetterli. Nonseparable multidimensional perfect reconstruction filter banks and wavelet bases for \mathbf{R}^n . 38(2):533–555, March 1992.
- [7] A. Laine and S. Schuler. Hexagonal QMF banks and wavelets. In M. Akay, editor, *Time-Frequency and Wavelet Transforms in Biomedical Engineering*. IEEE Press, 1997.
- [8] A. Laine, S. Schuler, W. Huda, and B. Steinbach. Hexagonal wavelet processing of digital mammography. In M.H. Loew, editor, *Medical Imaging 1993: Image Processing*, volume 1898, pages 559–573. Spie proceedings, 1997.

- [9] P. Schröder and W. Sweldens. Spherical wavelets: efficiently representing functions on the sphere. In *Computer Graphics Proceedings (SIGGRAPH95)*, pages 161–172. ACM Siggraph, 1995.
- [10] E.P. Simoncelli and E.H. Adelson. Non-separable extensions of quadrature mirror filters to multiple dimensions. *Proc. IEEE*, 78(4):652–663, 1990.
- [11] W. Sweldens. The lifting scheme: A new philosophy in biorthogonal wavelet constructions. In A. F. Laine and M. Unser, editors, *Wavelet Applications in Signal and Image Processing III*, pages 68–79, 68-79, 1995. *Proc. SPIE* 2569.
- [12] W. Sweldens. The lifting scheme: A custom design construction of biorthogonal wavelets. *Appl. Comput. Harmonic Anal.*, 3(2):186–200, 1996.
- [13] W. Sweldens. The lifting scheme: A construction of second generation wavelets. *SIAM J. Math. Anal.*, 29:511–548, 1998.
- [14] W. Sweldens and P. Schröder. Building your own wavelets at home. In *Wavelets in Computer Graphics, ACM SIGGRAPH Course Notes*. ACM, 1996.
- [15] G. Uytterhoeven and A. Bultheel. The Red-Black wavelet transform. In M. Moonen, editor, *Signal Processing Symposium (IEEE Benelux)*, pages 191–194. IEEE Benelux Signal Processing Chapter, 1998.
- [16] 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.