

A Wavelet Toolbox for Large Scale Image Processing

Geert Uytterhoeven Dirk Roose Adhemar Bultheel

Report TW 283, August 1998



Katholieke Universiteit Leuven
Department of Computer Science
Celestijnenlaan 200A – B-3001 Heverlee (Belgium)

A Wavelet Toolbox for Large Scale Image Processing

Geert Uytterhoeven Dirk Roose Adhemar Bultheel

Report TW 283, August 1998

Department of Computer Science, K.U.Leuven

Abstract

The wavelet transform has proven to be a valuable tool for image processing applications, like image compression and noise reduction. In this paper we present a scheme to process very large images that do not fit in the memory of a single computer, based on the software library WAILI (Wavelets with Integer Lifting). Such images are divided into blocks that are processed quasi independently, allowing efficient parallel programming. The blocking is almost completely transparent to the user.

Keywords : wavelets, lifting, image compression, image processing, tiling.
AMS(MOS) Classification : 41A30, 68U10, 94A12, 65Y05.

A Wavelet Toolbox for Large Scale Image Processing*

Geert Uytterhoeven[‡]

Dirk Roose[†]

Adhemar Bultheel[†]

August 31, 1998

Abstract

The wavelet transform has proven to be a valuable tool for image processing applications, like image compression and noise reduction. In this paper we present a scheme to process very large images that do not fit in the memory of a single computer, based on the software library WAILI (Wavelets with Integer Lifting). Such images are divided into blocks that are processed quasi independently, allowing efficient parallel programming. The blocking is almost completely transparent to the user.

1 Introduction

Wavelet transforms have good decorrelating properties. Wavelets are localized in both the spatial domain and the frequency domain, and are based on a multi-resolution analysis. This makes wavelets successful for image compression.

In application areas like Geographical Information Systems (GIS), image sizes are measured in gigabytes (GB) and even terabytes (TB). For example, an aerial color image of Belgium, where each pixel corresponds to one square meter, merely consumes 90 GB of storage. It is obvious that such images need to be stored and handled in a compressed form. Wavelets can provide for this since wavelet-based techniques offer better compression rates than other techniques. Moreover, many image processing operations can be done in the wavelet domain. Some of them can even be done better in the wavelet domain than in the spatial domain.

We have developed the software library WAILI, which allows the combination of image compression and processing in one package. In this paper we describe the extension of WAILI to block-based processing, which allows for a parallel implementation.

*Submitted to the special track on *Parallel Computing in Image Processing, Video Processing, and Multimedia* of the *International Conference on Parallel Computation (Fourth International Conference of the Austrian Center for Parallel Computation (ACPC'99))*.

[†]Department of Computer Science, Katholieke Universiteit Leuven, Celestijnenlaan 200A, B-3001 Heverlee, Belgium

[‡]Geert.Uytterhoeven@cs.kuleuven.ac.be, <http://www.cs.kuleuven.ac.be/~geert/>

2 Wavelets and the Lifting Scheme

2.1 Predict and Update

The wavelet transform of a 1D signal is a multi-resolution representation of that signal where the wavelets are the basis functions which at each resolution level give a highly decorrelated representation.

At each resolution 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.

The lifting scheme is an efficient implementation of these filtering operations. Several introductions to the lifting scheme are available [8, 9, 5, 7]. 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. For smooth signals, the prediction error is small. 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, while the transform should still be invertible. This is easily achieved within the lifting framework.

We 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:

$$\begin{aligned}\gamma_j &\leftarrow \gamma_j - [\mathcal{P}(\lambda_j)], \\ \lambda_j &\leftarrow \lambda_j + [\mathcal{U}(\gamma_j)].\end{aligned}$$

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 [1, 11].

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 [11]). 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:

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




Figure 1: The two-dimensional wavelet transform: iteration on the LL subband, showing the wavelet coefficients in “Mallat” ordering.

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 .

A whole family of lifting schemes can be constructed in this way [10], of which the above example is just the simplest possible case.

Since lifting steps use coefficients from regions near the coefficients that are updated, one needs special care at the boundaries. This is also important for a parallel implementation.

Note that this transform works on one-dimensional data. For two-dimensional data, like images, it can be applied row and columnwise, resulting in a tensor product wavelet transform. Instead of with 2 subbands (low pass and high pass), one will end up with 4 subbands: LL — low pass in both the horizontal and vertical directions, LH — low pass in the vertical, high pass in the horizontal direction, HL and HH. When iterated on the LL subband, the result is a multi-resolution decomposition as shown in Figure 1. This ordering of the subbands at the different resolution levels is called the “Mallat” ordering [3].

3 WAILI: Wavelets with Integer Lifting

WAILI is a software library — written in C++ — providing wavelet transforms and wavelet-based operations on two-dimensional images [11]. An image is a rectangular matrix containing pixels. Color images contain 3 channels, one for each primary color component. Hence the pixels of color images are three-valued. Popular color encoding schemes are RGB (red, green, blue), YUV and Lab.

WAILI implements various integer wavelet transforms, using lifting technology. It includes some image processing operations, such as:

- Crop and merge on wavelet transformed images. This allows to cut a rectangular subimage out of a large wavelet transformed image, or to replace a rectangular area in a large wavelet transformed image.
- Noise reduction, using thresholding based on generalized cross validation.

WAILI is available in source form¹ for research purposes.

¹<http://www.cs.kuleuven.ac.be/~wavelets/>

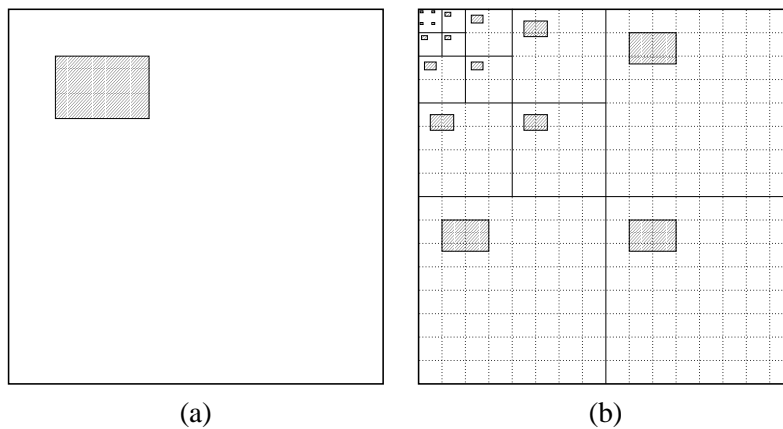


Figure 2: Locality in (a) space and in (b) the wavelet domain.

4 Very Large Images and Tiling

We define an image to be *very large* if it does not fit in the memory of the computer system that is processing it. Note that by this definition, whether an image is very large or not depends not only on the image but also on the computer system.

4.1 Tiling

In order to process very large images, they have to be divided in blocks (tiling).

Fortunately wavelets are localized in both the spatial and frequency domains: they have a finite support in both domains. This means that one needs only a limited subset of the input image to calculate a subset of the wavelet coefficients. The same is true for the inverse transform. From Figure 2, it is clear that only a fraction of the blocks in the right part of the figure are needed to reconstruct the rectangle in the left part.

Since not all blocks have to be present in main memory at the same time, this allows for the successful processing of very large images on machines with a single central processing unit (CPU). Moreover, the processing of the blocks can be distributed among multiple machines and/or CPUs.

We choose to separately divide each subband in blocks. Blocks are square and block sizes must be powers of two (except near the borders of the image), but block sizes may differ among resolution levels.

Since we want to process both complete images and subimages that are part of a large image, the upper left corner of a (sub)image is not always located at the origin $(0, 0)$, but may have arbitrary coordinates (see Figure 3). To facilitate the wavelet transforms, the internal boundaries of the blocks must always be aligned to the current block size. Hence, not only the blocks at the right and bottom borders of the image may be smaller than the block size, but also at the left and upper borders.

The whole wavelet decomposition is divided in blocks like shown in fig. 4. Within each channel, the black arrowed tree forms a “Mallat” ordering (cfr. Figure 1). A block within an image is thus uniquely defined by 5 indices:

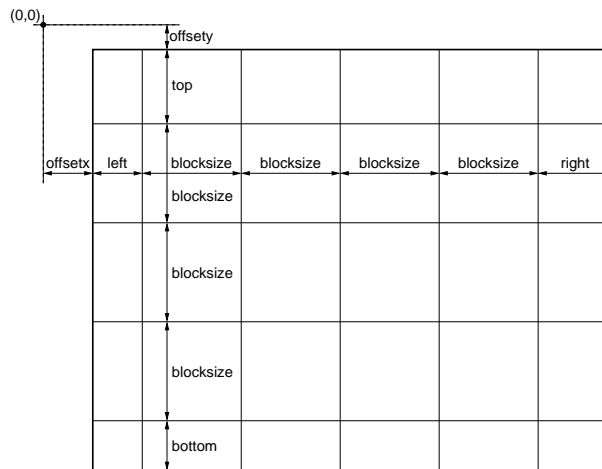


Figure 3: Division of a subband in blocks. The margins (top, bottom, left and right) are chosen such that the coordinates of the block boundaries are always multiples of the block size.

channel Color images consist of multiple channels.

level Resolution level.

subband LL, LH, HL or HH.

column and **row** Internal to the subband.

To implement wavelet transforms on tiled images, one has to consider the following steps:

Splitting step (lazy wavelet) To split a block in “even” and “odd” coefficients, in both the horizontal and vertical directions, one needs that block only. So this step is trivial to parallelize.

Lifting steps Lifting steps need a few coefficients from adjacent blocks to calculate the lifting operation on the border coefficients. Thus in a parallel implementation there must be communication of the border coefficients (overlap regions). Since in a primal lifting step the “even” blocks are updated by the “odd” blocks, and vice versa in a dual lifting step, these updates can be done independently. Hence, this step can be parallelized easily. The parallel efficiency grows with increasing block size due to the “perimeter effect”.

4.2 Block Management

We extended WAILI to support very large images that are divided in blocks. This is nearly completely transparent to the user: all traditional operations can be performed, just like on small non-blocked images.

The “block manager” is responsible for managing the blocks of which a very large image is composed. To access a part of the (transformed) image, the block manager is asked for blocks and retrieves them, e.g. from disk or via a network. If a block is modified, it has to be updated on the

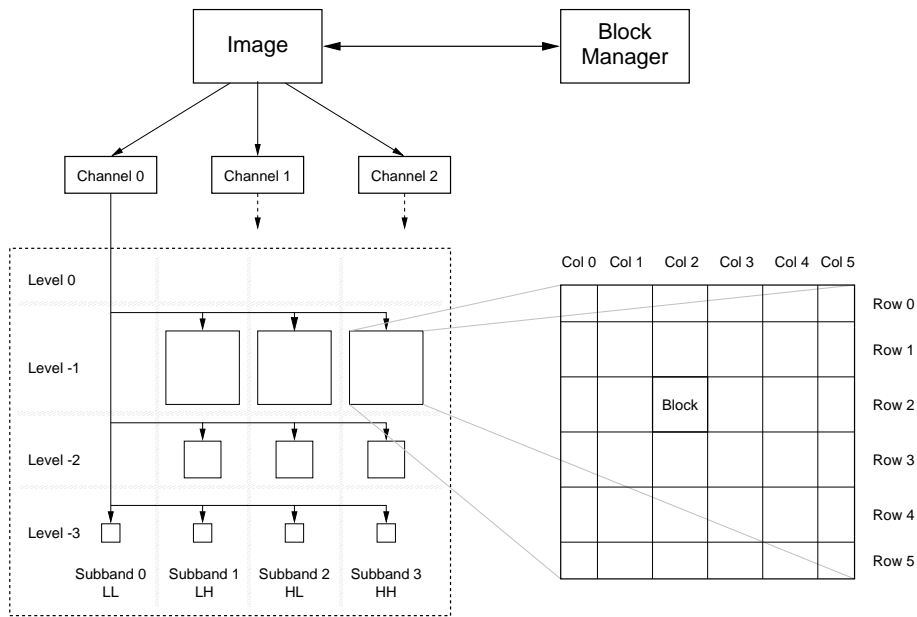


Figure 4: Subdivision of an image into channels, levels, subbands and blocks.

storage device. Blocks are passed to the compression module to remove redundancy among the wavelet coefficients within each blocks. The block manager for sequential processing also keeps a cache of recently used blocks to speed up processing. An overview of the complete system is given in Figure 5.

To obtain a parallel version, we only need to adapt the block manager and introduce communication of the border coefficients during the lifting steps.

4.3 Block-based Denoising of Images

WAILI uses soft-thresholding to reduce additive, stationary noise. The threshold is selected at each transform level and for each subband separately, using generalized cross validation [11]. Since this is an asymptotical method, it only performs well for sufficiently large subbands.

For wavelet transformed images that are divided into blocks, the denoising technique is applied independently on each block of each subband. Because only the highpass subbands are processed and blocks belonging to lowpass subbands are left alone, the borders between the blocks do not become visible after the denoising step.

Experiments showed that one gets better results, compared to the case where complete subbands are denoised, because now the threshold is more adapted to the local properties of the image. Thus denoising is not only easily parallelizable, it also yields a better image quality.

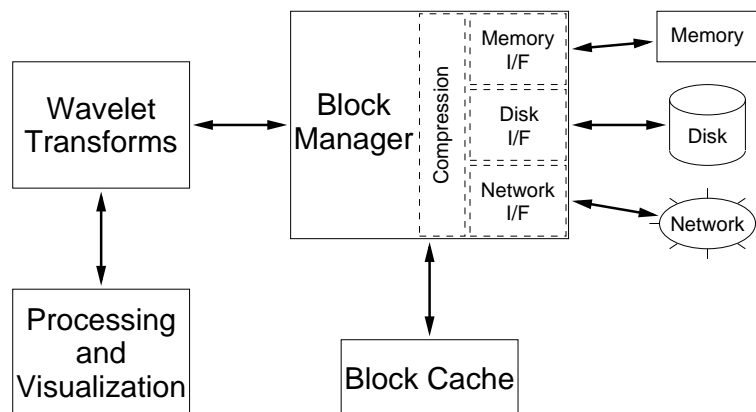


Figure 5: Design of the large scale image processing system.

5 Image Compression

Due to their decorrelating properties, wavelets are well suited for the compression of images and signals. If a sufficiently smooth image is transformed, most high pass coefficients will be zero or very small.

Nearly all modern image compression algorithms based on wavelets use a technique known as “zerotree encoding”, or a variant of it [6, 4]. Most image compression methods are targeted at “small” images. In that case the individual subbands are too small to get a sufficient compression gain from compressing them separately, hence zerotree encoding combines the subbands to exploit the correlation that is still present among spatially related coefficients at different resolution levels. If a coefficient at a low resolution level is small (or even zero), chances are high that the corresponding coefficients at higher resolution levels are also small (or zero) (see Figure 6).

For very large images, each zerotree contains coefficients from a multitude of blocks. Thus zerotree encoding will cause a block access pattern that is not feasible and the application of zerotree encoding is limited. A possible solution is to use a two-level encoding strategy:

1. For the lowest resolution levels, zerotree encoding can still be used because the subbands at these levels contain only one block.
2. For higher resolution levels, a block-based encoding can be used. At these levels the subbands can be very large, larger than the block size. If the block size is sufficiently large, each block will contain enough coefficients to be well compressible. Note that these subbands are high pass subbands, with many small or zero coefficients and (typically) only a few large coefficients. Since all blocks can be compressed independently, this step is trivially parallelizable.

Further research will be done to investigate this.

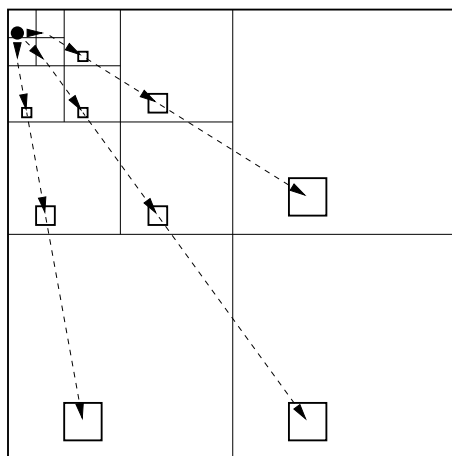


Figure 6: Zerotree compression: a tree of coefficients at different resolution levels that correspond to one low pass coefficient.

6 Conclusion

In this paper, we have described a strategy to perform wavelet-based image compression and processing on very large images. The images are split in blocks, which can be handled nearly independently.

The (integer) wavelet transform is implemented using the lifting scheme, which can be parallelized easily and efficiently, since only communication of pixels at the borders of the blocks is required.

The method is implemented as an extension to the software library WAILI. Blocks are accessed and handled under supervision of a block manager. Although the software has not been parallelized yet, the limited communication and the existing block manager allow an easy and efficient parallelization.

The division in blocks also has a positive influence on the performance of our denoising algorithm. Because denoising works independently on the blocks, it adapts better to the local properties of the image.

7 Acknowledgements

This paper presents research results of the Flemish Information Technology Action Program (‘Vlaams Actieprogramma Informatietechnologie’), project number ITA/950244. The scientific responsibility rests with its authors.

References

- [1] R. Calderbank, I. Daubechies, W. Sweldens, and B.-L. Yeo. Wavelet transforms that map integers to integers. *Appl. Comput. Harmon. Anal.*, 5(3):332–369, 1998.

- [2] A. Cohen, I. Daubechies, and J. Feauveau. Bi-orthogonal bases of compactly supported wavelets. *Comm. Pure Appl. Math.*, 45:485–560, 1992.
- [3] S. G. Mallat. Multifrequency channel decompositions of images and wavelet models. *IEEE Trans. Acoust. Speech Signal Process.*, 37(12):2091–2110, 1989.
- [4] A. Said and W. A. Pearlman. Image compression using the spatial-orientation tree. In *Proc. IEEE Int. Symp. Circuits and Syst., Chicago, IL*, pages 279–282, May 1993.
- [5] P. Schröder and W. Sweldens. Spherical wavelets: Efficiently representing functions on the sphere. *Computer Graphics, (SIGGRAPH '95 Proceedings)*, 1995.
- [6] J. M. Shapiro. Embedded image coding using zerotrees of wavelet coefficients. 41(12):3445–3462, 1993. *IEEE Trans. Signal Process.*
- [7] 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. *Proc. SPIE 2569*, 1995.
- [8] W. Sweldens. The lifting scheme: A custom-design construction of biorthogonal wavelets. *Appl. Comput. Harmon. Anal.*, 3(2):186–200, 1996.
- [9] W. Sweldens. The lifting scheme: A construction of second generation wavelets. *SIAM J. Math. Anal.*, 29(2), 1997.
- [10] W. Sweldens and P. Schröder. Building your own wavelets at home. In *Wavelets in Computer Graphics, ACM SIGGRAPH 1996 Course Notes*. ACM Press, 1996.
- [11] G. Uytterhoeven, F. Van Wulpen, M. Jansen, D. Roose, and A. Bultheel. WAILI: Wavelets with Integer Lifting. TW Report 262, Department of Computer Science, Katholieke Universiteit Leuven, Belgium, July 1997.