Presented at SPIE Photonics East: Digital Storage and Archiving Systems, 25-26 October 1995, pp. 324-335

An Approach to Efficient Storage, Retrieval, and Browsing of Large Scale Image Databases

Norbert Strobel, Sanjit K. Mitra, and B.S. Manjunath

Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106

ABSTRACT:

This paper suggests a wavelet transform based multiresolution approach as a viable solution to the problems of storage, retrieval and browsing in a large image database. We also investigate the performance of an optimal uniform mean square quantizer in representing all transform coefficients to ensure that the disk space necessary for storing a multiresolution representation does not exceed that of the original image. In addition, popular wavelet filters are compared with respect to their reconstruction performance and computational complexity. We conclude that, for our application, the Haar wavelet filters offer an appropriate compromise between reconstruction performance and computational efforts.

Keywords: digital library, storage, retrieval, browsing, multiresolution representations, discrete wavelet transform, uniform quantization, filter comparison, Haar filters

1 INTRODUCTION

There are three major problems associated with a digital library containing many, potentially very large-size images: storage, browsing, and retrieval. As these problems are interdependend, an integrated approach is needed. From a practical point of view, it can be assumed that database users require rapid browsing through the data available to find the images they need which then should be instantly retrievable. Database providers, on the other hand, face the problem of how to implement these features given the resources available.

One obvious limitation is finite storage space. While not necessarily a pressing issue when dealing with small images, even hard drives with capacities of up to one Terrabyte cannot accommodate too many Landsat satellite images. They typically consist of seven bands in the range of 5000 by 5000 pixels and require 175 MB of storage space. A hard drive with one Gigabyte, as, for example, available for personal computers, can just accommodate five or six of such images leaving barely enough room for additional software. Intimately tied to image storage is browsing which can be classified into three different types: textual queries, image content based search,[1] and visual browsing. In this paper we restrict ourselves to the last type. While textual queries and image content based retrieval are methods to narrow down the search space, visual browsing is used in the last phase to obtain an overview of the data found.[2] To this end very small size subimages, called thumbnails, representing original images are often sufficient. Based on the thumbnails delivered, the user should be able to make a decision about

which images to discard and which interesting examples to further investigate. Visual browsing should support what humans can do best, namely exploring.

Unfortunately, due to the enormous size of satellite imagery the standard approach of storing multiple copies of the same image at different resolutions as, e.g., employed by Chabot,[3] takes up a significant amount of resources. In addition, communication bandwidth is another major problem for remote databases. For example, when using the internet's file transfer protocol (ftp) it can take between one and up to two hours to retrieve just one Landsat image.

Considering the complexity of storage, browsing and retrieval when dealing with large-size image databases, we propose a multiresolution approach based on the discrete wavelet transform. Its properties nicely support browsing and retrieval requirements as explained below.

The paper is organized as follows: First, we give a brief introduction to multiresolution representations. Classical pyramid schemes are reviewed and the discrete wavelet transform is explained. Second, we investigate a simple solution to the major shortcoming of the discrete wavelet transform which is the increase in bit precision needed to retain perfect reconstruction of the original image from the transform coefficients. Third, we evaluate different popular biorthogonal wavelet filters with respect to their reconstruction performance and computational complexity. Finally, a discussion of the results is provided.

2 MULTIRESOLUTION REPRESENTATIONS

2.1 The Basic Concept

A multiresolution representation of an image creates a pyramid structure consisting of multiple replicas of that image at different resolutions. A typical example is shown in Fig. 1.

The multiresolution representation at the highest level usually provides a crude approximation of the original image which can be used as a thumbnail for visual browsing. Successively reconstructed copies at finer resolutions are then obtained by incorporating information about additional details present at that and lower levels. In addition, multiresolution decompositions inherently support progressive data transmission since the thumbnail can be sent first and successively be refined when sending additional data. In fact, a well-designed multiresolution-based retrieval system can cut communication costs at least in half.[4] Furthermore, any image database may be connected to many different display devices. Database providers, therefore, are well advised to use a multiresolution technique which effectively supports different hardware configurations (multiuser environment).[5]

To summarize, multiresolution techniques which assist browsing among thumbnails and support progressive image retrieval are powerful tools, especially for searching databases containing large images. We review two basic multiresolution concepts: pyramid structures and subband decomposition. We consider the one-dimensional case to illustrate basic multiresolution techniques.

2.2 Pyramid Structures

The most obvious procedure to generate a multiresolution decomposition is to repeatedly down-sample signals as shown in Fig. 2. To generate a lower resolution approximation $r_1(n)$ of the original signal $r_0(n)$, every other sample is kept. Note that subscripts refer to decomposition levels indicated in Fig. 1. Those data elements which were discarded in the upper branch are retained as $d_1(n)$ signal. This procedure can be iterated on $r_1(n)$ to obtain even coarser approximations. For example, after K iterations we obtain the approximation signal $r_K(n)$

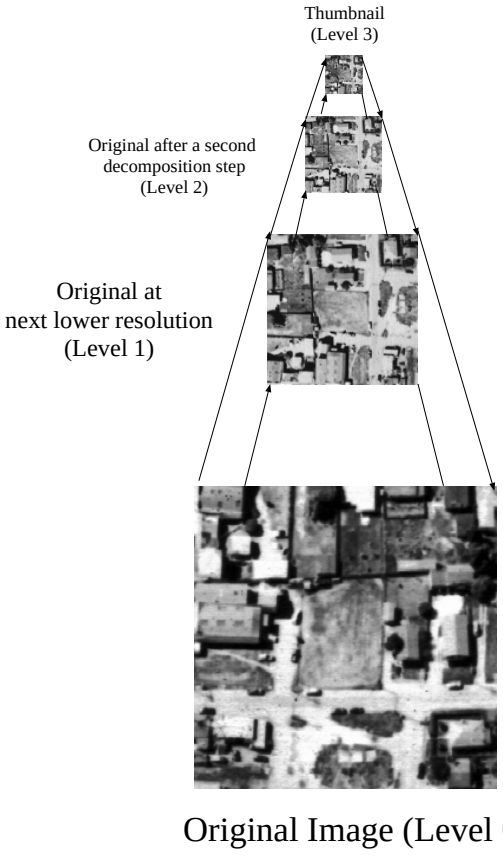

Figure 1: Multiresolution representation as, for example, obtained by a discrete wavelet transform with three levels

and detail signals $d_i(n)$, $i = 1, 2, \ldots K$. Note that each down-sampled signal comprises only half as many data points as the input to the down-sampler since only every second sample is kept. Although the total amount of data stored is equal to that of the original signal and even though it can be recovered exactly, there are serious disadvantages to this approach. First, simple subsampling introduces aliasing artifacts which can significantly impair the information content of image approximations at higher levels of this pyramid. Thus, the use of subsampling without any prefiltering may render coarse image approximations unsuitable for display.[5] Second, there is no guarantee that the remaining pixels are good representatives of the areas from which they are taken.

Another method to create a multiresolution pyramid is based on a prediction/residual approach[6] as shown in Fig. 3 for a single stage decomposition. To create a prediction pyramid from an input signal $r_0(n)$, a coarse approximation $r_1(n)$ of the original input is derived first. This can be achieved by passing it through a low-pass filter $h_0(n)$ followed by down-sampling. A prediction of the original is obtained by first up-sampling and then interpolating the coarse (and smaller) approximation $r_1(n)$ with the lowpass filter $g_0(n)$. Subtracting the prediction signal from the original yields the prediction error $d_1(n)$. Iterating the pyramid decomposition on the approximation data $r_1(n)$ leads to successively coarser representations $r_2(n)$, $r_3(n)$, ... and corresponding prediction errors $d_2(n)$, $d_3(n)$, ... The approximation signal at the highest level and all detail signals are stored. Together they allow a perfect reconstruction of the original data in the absence of any quantization. The main disadvantage of the prediction/residual approach is its overcomplete signal expansion. There are 25 – 33% more data elements to store depending on the number of iterations.

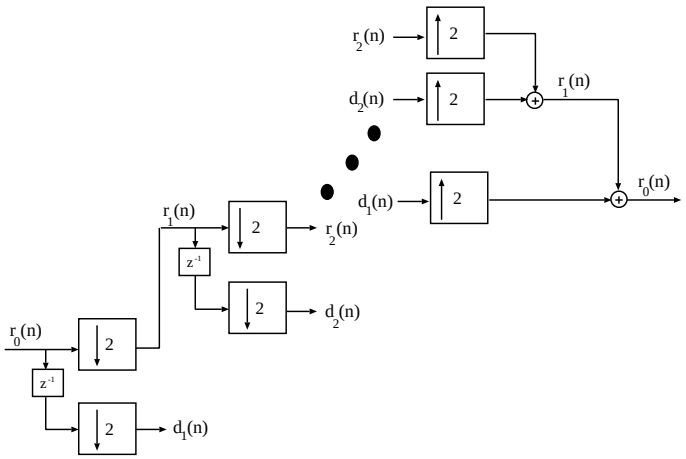

Figure 2: Multiresolution representation by direct down-sampling

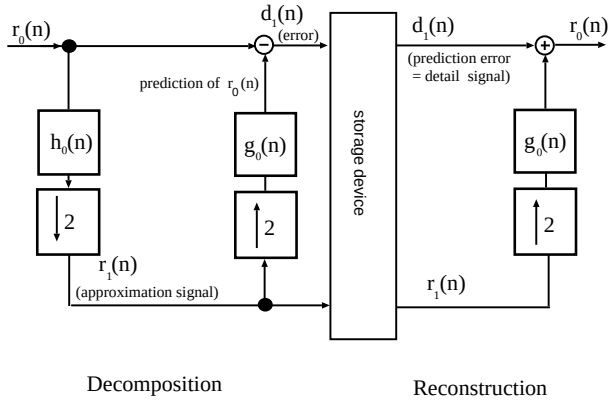

Figure 3: Multiresolution representation by prediction pyramid

2.3 Subband Decomposition and the Discrete Wavelet Transform

Similar to the prediction pyramid, subband decomposition provides a hierarchical data structure. However, there is no increase in the number of transform elements to store; in addition, reconstruction at successive resolutions (as shown in Fig. 1) is extremely simple. This makes the subband decomposition very attractive for image database applications.

In general, an M-channel filter bank could be used to generate lower resolution approximations for browsing. This method is illustrated in the block diagram of Fig. 4. There we derive a coarser approximation $r_1(n)$ from the input signal $r_0(n)$ by passing it through a lowpass filter $h_0(n)$ first. Then we critically decimate (down-sample with a factor of M) the resulting sequence obtaining $r_1(n)$. The detail signals corresponding to the i-th channel, called $d_i(n)$, follow similarly by processing $r_0(n)$ with the filter $h_i(n)$ and then down-sampling the result. The filter bank depicted in Fig. 4 allows perfect reconstruction of $r_0(n)$ from approximation and detail signals $r_1(n)$ and $d_i(n)$, $i = 1, 2, \ldots, M-1$, respectively, if some constraints on the filter coefficients are satisfied.[7]

Although an M-band filter bank could be used for computing multiresolution representations, i.e. arbitrarily

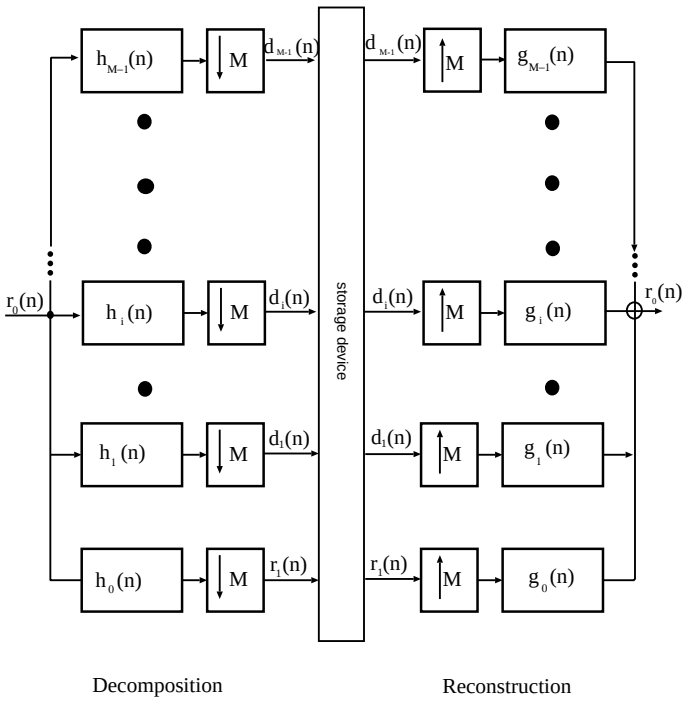

Figure 4: M-channel filter bank with analysis filters $h_i(n)$ and synthesis filters $g_i(n)$

many channels are possible, an iterated two-channel filter bank as shown in Fig. 5 is preferred for convenience. Not only does it generate the smallest number of bands necessary to implement a multiresolution representation but it also serves as a basic architecture to implement the discrete wavelet transform.[8] Iterating such a two-channel filter bank K–times on the approximation signal $r_1(n)$ results in an octave-band splitting scheme with successively coarser and smaller approximation signals $r_i(n)$ and additional details or wavelet coefficients $d_i(n)$. For perfect reconstruction we store the thumbnail representation $r_K(n)$ and all detail signals $d_i(n)$, $i = 1, 2, \ldots, K$. Note that the amount of data points kept is equal to the number of samples in the original signal.

The inverse wavelet transform is shown on the right hand side of Fig. 5. To obtain $r_{K-1}(n)$, for example, we recombine detail and approximations signals $d_K(n)$, $r_K(n)$ by first up-sampling each, then interpolating and finally adding them both.

In the preceding paragraphs all theory and examples were given for one-dimensional signals only. Images, however, are two-dimensional. To apply the subband or discrete wavelet transform to 2-D signal we make use of separate row and column processing.[9] This means that first the rows and then the columns are filtered, or vice versa. The separable implementation yields four different spatially oriented outputs (frequency channels) for each transform iteration as shown in Fig. 6. They are assumed to correspond to important (LL band) and less significant visual information (LH, HL, and HH bands):

1. LH band: obtained by first lowpass filtering along the rows and then applying a highpass filter along the columns; gives the vertical high frequencies (horizontal edges).
2. HL band : obtained by first highpass filtering along the rows followed by lowpass filtering along the columns; contains horizontal high frequencies (vertical edges).
3. HH band: resulting from highpass filtering in both directions; highlights diagonal frequencies

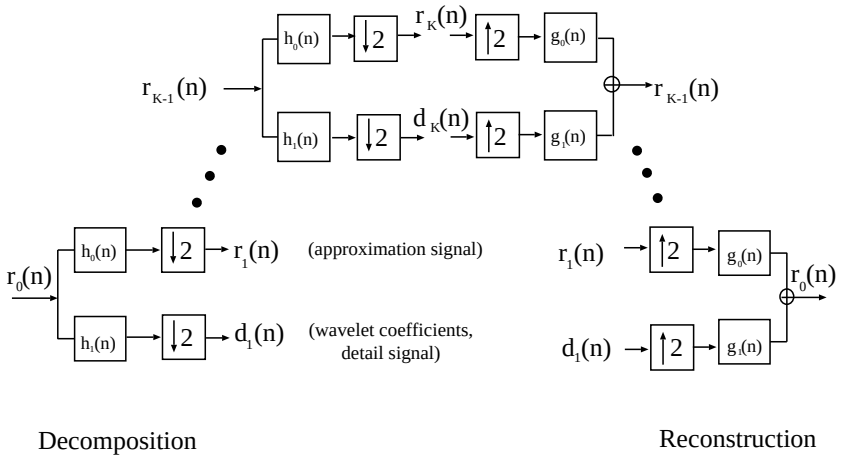

Figure 5: Multiresolution Representation by octave-band subband or discrete wavelet transform

4. LL band: based on lowpass filtering along rows and columns; results in a coarse approximation of the original image.

For typical applications in a geographical image database we expect between three and six iterations on the LL band which lead to LH, HL, and HH bands at successively higher scales (smaller resolutions) and one remaining thumbnail at the highest scale. This subband can be considered a small-size approximation of the original.

The filter bank implementation as shown in Fig. 5 allows forward and inverse wavelet transforms in $O(N)$ time (N = image width $\times$ image height). The numerical complexity of this so-called *fast* wavelet transform, hence, only increases linearly with data size.

The goal of any multiresolution browsing scheme is to provide replicas of the original image at different resolutions. To avoid phase distortions in coarser image approximations the filters employed must have linear phase. To ensure fast computation the filters should also be short. Unfortunately, linear phase and orthogonality in two-band finite response (FIR) systems are incompatible properties.[10] The only symmetric exact reconstruction filters are those corresponding to the Haar basis. To obtain linear phase FIR filters, the orthonormality property is usually relaxed leading to wavelet transforms with biorthogonal bases.[8]

3 QUANTIZATION OF TRANSFORM COEFFICIENTS

3.1 The Need for Quantization

Although the number of wavelet transform coefficients equals the number of pixels in the original image, storage requirements usually differ as the wavelet transform converts image intensities usually stored in some integer format (e.g., one byte/pixel) into floating point numbers. As a result, transform data has to be stored with higher precision, if perfect reconstruction is desired. Such an increase in data storage and transmission costs clearly offsets many advantages of a multiresolution approach, thus, requiring some data reduction after transformation. Ideally we would like to be able to support multiresolution browsing and provide (near) lossless

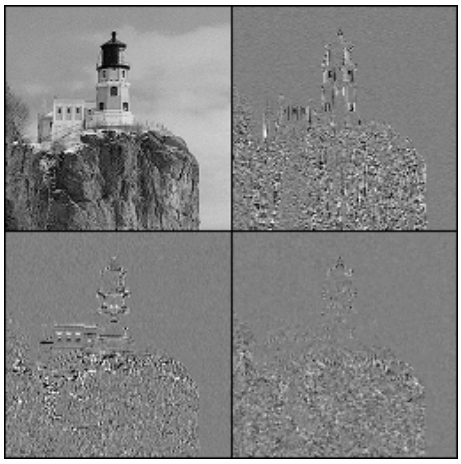

Figure 6: Wavelet decomposition: One forward wavelet transform iteration was performed. The image in the upper left corner shows the thumbnail obtained (LL band). The other bands refer to spatially oriented channels. The HL band is displayed in the upper right corner, the HH band is in the lower right, and the LH band is in the lower left corner.

image reconstruction in the shortest time possible. This is a classical problem in still image compression which can be defined in terms of optimizing transmission rate with respect to image distortion.[11] In an interactive multiresolution browsing environment, however, good rate/distortion performance is only one requirement. Even more important is fast data decoding and image reconstruction. Users can be expected to demand a fast response time at least in the early stages of a browsing session when primarily dealing with thumbnails. Note that due to quantization effects, subband or wavelet coding typically results in a lossy reconstruction unless a final residual error is included.

3.2 Optimal Uniform Quantization

To ensure uncomplicated and fast decoding all subbands in our investigation are allocated (suboptimal) eight bits per pixel and uniformly quantized. As we concentrate on images stored with one byte/pixel this constraint effectively enforces that transform coefficients take up the same space as the original uncompressed image.

We decided to design an optimal uniform mean square quantizer for all subbands except the LL band regardless of their statistical properties.[12] We would like to remark that uniform quantization is asymptotically optimal only when combined with variable length coding. For fixed length codes, as in our case however, nonuniform quantization would achieve higher rate/distortion performances. The LL band and all successively reconstructed image approximations were quantized by thresholding and rounding to the nearest integer within the range $[0, 255]$. This successive requantization step introduces additional errors and, thus, impairs the final reconstruction performance. Requantizing LL bands, however, simplifies the file handling at the client side since only one file (approximate representation) is used for both display and further reconstructions instead of two copies (clipped approximation together with inverse transform result).

The quantizer design for the LH, HL, and HH bands at different scales follows the algorithm suggested by Max.[13] We assume a symmetric input distribution of transform coefficients $p(x)$ and need to find a quantizer bin

width r such that the following quantization distortion $D(r)$ is minimized:

$$D(r) = \int_{\frac{-r}{2}}^{\frac{r}{2}} x^2 p(x)dx + 2\sum_{i=1}^{K-1} \int_{(i-0.5)r}^{(i+0.5)r} (x - ir)^2 p(x)dx + 2\int_{(K-0.5)r}^{\infty} (x - Kr)^2 p(x)dx \quad (1)$$

Equation (1) computes the distortion of a symmetric midtread quantizer with an odd number of $2K + 1$ reconstruction levels ($K = 127$ for 8 bits). We favor a symmetric midtread quantizer as it ensures a reconstruction level at zero which eliminates some of the picture noise.[14] The optimal bin width r_{min} which minimizes Eq. (1) presents a compromise between granular and overload distortion as shown in Fig. 7.

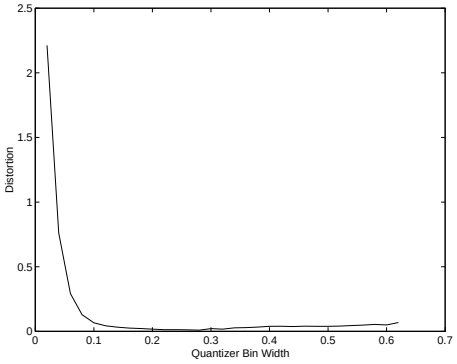

Figure 7: Quantizer Distortion: The minimum distortion occurs for a quantizer bin width r_{min} of 0.28. To the left, i.e., for smaller bin widths, overload distortion dominates, while to the right, i.e., for larger bin widths, granular distortion steadily increases.

4 FILTER COMPARISONS

4.1 Selection of Wavelet Filters

The wavelet filters considered in this paper have either been previously evaluated and recommended due to a superior performance,[15] suggested because of a possible fast (multiplierless) implementation,[16] or proposed for a new still image compression standard (CREW),[17].[18] Many of the filters recommended belong to the class of spline filters.[19] The CREW filter is a "reversible" implementation of a traditional wavelet filter suggested by Le Gall and Tabatabai.[20] The examples evaluated have average filter length smaller than 16 since low computational complexity of the inverse discrete wavelet transform is of major importance in an interactive environment.

4.2 Evaluation Criteria: Filter Length, PSNR, Error Entropy

The filter evaluation was primarily concerned with the trade-off between reconstruction performance of individual filter candidates versus their computational requirements. The figures of merit used are the filter lengths and the peak-signal-to-noise ratio (PSNR). The longer the filters, the more multiplications and additions are necessary to compute one filter output sample. The higher the PSNR, defined for 8-bit images as

$$PSNR = 20\log_{10}(\frac{255}{RMSE}), \quad (2)$$

the smaller is the reconstruction error. The parameter $RMSE$ in Eq. (2) refers to the root-mean-squared error, which for an N x N image $f(i,j)$ and its reconstructed equivalent $\hat{f}(i,j)$, is given by:

$$RMSE = \sqrt{\frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N}[f(i,j)-\hat{f}(i,j)]^2}. \tag{3}$$

We additionally included the zeroth-order entropy, H_e, of the error image $e(i,j) := f(i,j) - \hat{f}(i,j)$ defined as:

$$H_e := -\sum_{i=1}^{N}\sum_{j=1}^{N} p(e(i,j))\log_2(p(e(i,j)). \tag{4}$$

The variable $p(e(i,j))$ in Eq. (4) denotes the probability of occurrence of error value $e(i,j)$.

The error entropy H_e provides a measure of how much information still is contained in the error image, i.e, it gives an estimate for how long the average error codeword length would be, if a lossless variable length encoding scheme was chosen. In other words, if perfect reconstruction is desired by means of lossy plus lossless residual encoding, then the error entropy supplies an estimate for the additional information still needed after the last lossy decoding step.

4.3 Results

The results are listed in Table 1. There the filter coefficients start with the zero coefficient and comprise the causal part of the linear phase filter. Those related to negative indices follow by symmetry. Note that the coefficients listed are down-scaled by a factor of $\sqrt{2}$. We would like to remark that the factor $\sqrt{2}$ is commonly incorporated into the coefficients $h_0(n)$, $h_1(n)$, $g_0(n)$, and $g_1(n)$. However, it is more efficiently implemented by installing a multiplier with a factor 2 in the synthesis stage.

The PSNR and error entropies listed in Table 1 are average values over ten test images of size 512 by 512. Three transforms iterations were performed on: Lena, Clown, Baboon, Sailboat, Barbara, Miramar, Channel, Landsat, LAX, and Forest.

5 DISCUSSION

Taking a look at Table 1 confirms Daubechies' popular 9/7 filter as a preferable candidate for image coding applications where low distortion (high PSNR) is of prime importance. However, within the limits of our primitive quantization scheme, Haar filter and the binomial 2/6 filter appear quite competitive. In fact, due to the extremely low computational requirements of the Haar filter, we chose it as our prime candidate for further investigations. Since it can be implemented solely based on shift and addition/subtraction operations, the Haar filter guarantees very fast reconstruction at the client side. Its low computational complexity and simple structure are also very attractive properties for parallel implementations. Almost all filters yielded reconstructed images which appeared perceptually lossless. We would, however, like to remark that the results for filters no. 9 and 10 have to be read carefully as they significantly suffered from the successive requantization of the LL bands at each resolution level. It should be pointed out that this is a major difference of our browsing scheme compared to the usual way of performing any inverse subband transform and it can, indeed, become a major source or errors. In practical applications, however, successive requantization of the LL band seems more appropriate compared to the option of keeping two images at the client side. Especially at intermediate resolutions we are more concerned with image retrieval speed than quality. In addition, it appears to be more important that a user can get more, less accurate

images, than fewer, more exact replicas still conveying little more visual browsing information.

An excellent candidate for an application where LL bands must not take up more than eight bits per pixel would be the reversible wavelet implementation.[17] However, the bit precision required to exactly preserve HL, LH, and HH increases with the number of forward transformation iterations and, hence, straight quantization with eight bits/pixel introduces quantization errors there. To avoid any misunderstandings, it must be pointed out that the CREW system proposed as a new image compression standard is capable of handling this and many other problems by employing a very elegant and powerful coding scheme loosely inspired by zero-tree quantization methods,[18,25]

Again, note that the merit of our preliminary quantization system is its simplicity. More sophisticated, state-of-the-art wavelet coding algorithms can obtain much higher rate/distortion rates. Shapiro,[26] for example, achieved PSNR values of more than 35 dB with a rate of only one bit per pixel using, however, a significantly more sophisticated (zero tree) coding scheme.

If combined lossy, lossless data transmission is a design objective, than Table 1 reveals that the zeroth-order error entropy can be assumed to be of the order of two bits. Hence, after performing the last lossy inverse transform step, an amount of information approximately equal to one quarter of the original data remains to be send to obtain an exact replica of the original image.

6 ACKNOWLEDGEMENTS

This research is part of an ongoing Alexandria digital library project being carried out at the University of California, Santa Barbara under NSF Grant Number IRI 94-11330. We thank the Alexandria Web team for their constructive criticisms. The authors also acknowledge the assistance of Daniel Wu who helped implementing the wavelet transform algorithm.

7 REFERENCES

[1] M. Flickner, H. Sawhney, W. Niblack, *et al.*, "Query by image and video content: The QBIC system," *IEEE Computer Magazine*, vol. 28, pp. 23 – 32, September 1995.

[2] J. Cha and S. Lee, "Complex icon browsing method for retrieval of multimedia objects," in *Proc. IASTED/ISMM 1994*, (Honolulu, Hawaii), pp. 10–13, 1994.

[3] V. E. Ogle and M. Stonebraker, "Chabot: Retrieval from a relational database of images," *IEEE Computer Magazine*, vol. 28, pp. 40 – 48, September 1995.

[4] A. Ortega, Z. Zhang, and M. Vetterli, "Modelling and optimization of a multiresolution remote image retrieval system," in *Proc. SPIE 1994*, vol. 2186, (San Jose, California), pp. 185–196, 1994.

[5] M. Rabbani and P. W. Jones, *Digital Image Compression Techniques.* Bellingham, WA: SPIE, 1991.

[6] B. Burt and E. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Transactions on Communications*, vol. 31, pp. 532 – 540, 1983.

[7] P. Vaidyanathan, *Multirate Systems and Filter Banks.* Englewood Cliffs, NJ: Prentice Hall, 1992.

[8] M. Vetterli and J. Kovacevic, *Wavelets and Subband Coding.* Englewood Cliffs, NJ: Prentice Hall, 1995.

[9] J. Woods and S. O'Neil, "Subband coding of images," *IEEE Transactions on Acoustics, Speech adn Signal Processing*, vol. 34, pp. 1278 – 1288, 1986.

[10] M. Smith and D. Barnwell, "Exact reconstruction for tree-structured subband coders," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 34, pp. 434 – 441, 1986.

[11] A. N. Netravali and B. G. Haskell, *Digital Pictures.* New York, NY: Plenum Press, 1995.

[12] P. Westerink, J. Biemond, and D. Boekee, "Subband coding of color images," in *Subband Image Coding* (J. Woods, ed.), Boston, Kluwer Academic Publishers, 1991.

[13] J. Max, "Quantizing for minimum distortion," *IRE Transactions on Information Theory*, pp. 7 – 12, March 1960.

[14] H. Gharavi and A. Tabatabai, "Sub-band coding of digital images using two-dimensional quadrature mirror filtering," in *Proceedings of the SPIE*, vol. 707, (Cambridge, Masachusetts), pp. 51 – 61, 1986.

[15] J. Villasenor, B. Belzer, and J. Liao, "Wavelet filter evaluation for image compression," *IEEE Transactions on Image Processing*, vol. 4, pp. 1053 – 1060, August 1995.

[16] R. Moorhead and Z. Zhu, "Signal processing aspects of scientific visualization," *IEEE Signal Processing Magazine*, vol. 12, pp. 20 – 41, September 1995.

[17] A. Zandi, J. Allen, E. Schwartz, and M. Boliek, "CREW: compression with reversible embedded wavelets," in *Proc. IEEE Data Compression Conference*, (Snowbird, Utah), pp. 212–221, March 1995.

[18] M. Boliek and A. Zandi, "Coding of still pictures." Initial Proposal, June 1995. RICOH California Research Center.

[19] P. Sriram and M. Marcellin, "Image coding using wavelet transforms and entropy-constrained trellis-coded quantization," *IEEE Trans. on Image Processing*, vol. 4, pp. 725 – 733, June 1995.

[20] D. L. Gall and A. Tabatabai, "Sub-band coding of digital images using symmetric short kernel filters and arithmetic coding techniques," in *Proc. International Conference on Acoustics Speech and Signal Processing*, pp. 761– 765, New York, IEEE, 1988.

[21] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, "Image coding using wavelet transform," *IEEE Trans. Image Processing*, vol. 1, pp. 205–220, April 1992.

[22] R. DeVore, B. Jawerth, and B. Lucier, "Image compression through wavelet transforms coding," *IEEE Trans. on Information Theory*, vol. 38, pp. 719 – 746, March 1992.

[23] E. Simoncelli, W. Freeman, E. Adelson, and D. Heeger, "Shiftable multiscale transforms," *IEEE Transactions on Information Theory*, vol. 38, pp. 587 – 607, 1992.

[24] E. Majani, "Biorthogonal wavelets for image compression," in *Proceedings of the SPIE*, vol. 2308, (Chicago, Illinois), pp. 478 – 488, 1994.

[25] A. Lewis and G. Knowles, "Image compression using the 2-D wavelet transform," *IEEE Trans. on Image Processing*, vol. 1, pp. 244 – 250, April 1992.

[26] J. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. on Signal Processing: Special Issue on Wavelets and Signal Processing*, vol. 41, pp. 3345 – 3462, December 1993.

Table 1: Biorthogonal Filters for the Discrete Wavelet Transform

		L [a]	Filter Coefficients [b]	PSNR [c]	H_e [d]	Comments
1	$h_0/\sqrt{2}$	9	.602948, .266864, −.078223, −.016864, .0267490	47.58	1.99	Ref.[21]
	$g_0/\sqrt{2}$	7	.557544, .295636, −.028772, −.045636			
2	$h_0/\sqrt{2}$	2	.5, .5	46.91	2.11	Haar Filter[22]
	$g_0/\sqrt{2}$	2	.5, .5			
3	$h_0/\sqrt{2}$	2	.5, .5	46.47	2.24	Ref.[15]
	$g_0/\sqrt{2}$	6	.5000, .0625, −.0625			
4	RTS	2	$r(n) = \lfloor \frac{x(2n)+x(2n+1)}{2} \rfloor$	44.75	2.55	Ref.[17]
		6	$d(n) = \lfloor \frac{-r(n)+4(x(2n+2)-x(2n+3))+r(n+2)}{4} \rfloor$			
5	$h_0/\sqrt{2}$	5	.750, .250, −.125	43.81	2.24	Ref.[23]
	$g_0/\sqrt{2}$	3	.50, .25			
6	$h_0/\sqrt{2}$	6	.557544, .0337282, −.091272	43.38	2.20	Ref.[24]
	$g_0/\sqrt{2}$	10	.434907, .094320, −.047544, .004942, .013374			
7	$h_0/\sqrt{2}$	9	.7031250, .2968750, −.1250000, −.0468750, .0234375	41.95	2.44	Ref.[21]
	$g_0/\sqrt{2}$	3	.50, .25			
8	$h_0/\sqrt{2}$	8	.703125, −.109375, −.140625, .046875	34.34	3.19	Ref.[16]
	$g_0/\sqrt{2}$	4	.375, .125			
9	$h_0/\sqrt{2}$	4	.750, −.250	30.19	3.83	Ref.[16]
	$g_0/\sqrt{2}$	4	.375, .125			

[a] Filter length of $h_0(n)$ and $g_0(n)$, respectively

[b] The filter coefficients start with the zero coefficient; negative indices follow by symmetry. They are shown downscaled by a factor of $\sqrt{2}$ compared, e.g., to the results listed by Ref.,[15] to clearly reveal potential for computationally efficient implementations.

[c] The PSNR is listed in dB and states the average result obtained from a set of ten test images.

[d] H_e is the average *error* entropy of the ten corresponding error images.