IMPROVED ESTIMATOR FOR RECOVERING HIGH FREQUENCY COMPONENTS FROM COMPRESSED IMAGE DATA

FIELD OF THE INVENTION

This invention relates to digital signal processing techniques in general and in particular to the use of digital signal processing techniques for compression and decompression of data and for the reliable recovery of high frequency components discarded during compression of the data.

BACKGROUND

The wavelet transform has been shown to be better than traditional Fourier methods for representing the different spatial frequencies (edges and textures) in images. This is important for image compression since it is well known that humans are more sensitive to low and medium frequencies than they are to high frequencies. As a result, many compression schemes quantize lower frequency coefficients more finely than they do higher frequency coefficients.

It is also well known that detail barely perceptible in a given frequency sub-band must be at least four times as intense in the next higher frequency sub-band to be perceivable. The transformed image is divided into contiguous blocks of coefficients, e.g., 4×4, 8×8, etc., and each block is matched against a dictionary or code book of blocks belonging to a sufficiently large training set. A pointer to the closest match in the dictionary is saved in place of a transform block thus achieving compression. This method requires a complex algorithm to build an efficient dictionary for describing an ensemble of data. The disadvantage in this approach is that it is not general enough for any type of data.

After quantization of the transform coefficients, the last step of conventional image compression schemes is to apply a lossless coding techniques such as Huffman coding or arithmetic coding. The disadvantage with conventional approaches is that quantization results in artifacts in the decompressed images that are very visible at medium to high compression ratios. As a result, compression ratios resulting in reproduced images of acceptable quality are rather modest.

Broadly, it is an object of the present invention to overcome the most serious limitations of current approaches to data compression by taking advantage of the mathematics of wavelet theory and of the frequency sensitive characteristics of human perception.

This and other objects of the present invention will become apparent to those skilled in the art from the following description of the invention.

SUMMARY OF THE INVENTION

It is well known that human visual perception relies most heavily on the low frequency components of an image. Consequently, it is advantageous, when compressing an image, to discard those portions of an image which contain high frequencies and which are therefore less important for human perception. This permits a file representative of the image to require less storage space.

Although the high frequency coefficients are less important for human visual perception, their presence nevertheless enhances the overall clarity of the image. High frequency components are associated with edges in an image, hence their presence tends to sharpen edges and delineate boundaries more precisely. For this reason, it is advantageous to incorporate the high frequency components back into the image before displaying it. However, since the high frequency components were presumably discarded in the interest of reducing storage requirements for the image, they must somehow be estimated.

The invention comprises a wavelet transform based method and system for estimating missing high frequency components of an image based on those frequency components which are present in the image. As mentioned above, these components may be missing because they were discarded, typically by being set to zero, during compression. However, the method of the invention can also be used where the high frequency coefficients were never present in the image to begin with. For example, in the course of image enlargement, it is desirable to assign a pixel value to the gaps that form between pixels as the image grows. The method of the invention can also estimate these values based on existing pixel values, thereby enhancing the quality of the enlarged image.

Application of the method of the invention to the enlargement of an original image starts with the assumption that the original image includes the low frequency components of the wavelet transform of an enlarged image which is four times larger than the original image. The high frequency coefficients of the enlarged image are then estimated from the low frequency coefficients in the original image. This is followed by the application of the inverse wavelet transform. The result is an image which is not only four times larger but is also of enhanced quality because the resolution has been doubled. This process can be performed repeatedly to obtain successively larger images.

Application of the method of the invention to the compression of an image and to the subsequent estimation of the discarded high frequency coefficients begins with applying the wavelet transform to the original image and discarding some or all of the high frequency coefficients. The wavelet transform is then applied to the remaining

coefficients of the wavelet transform of the original image. In a preferred practice of the invention, this process optionally continues for three additional levels of transformation. At this level, only those high frequency coefficients most important for perception are kept. These must then be efficiently encoded, typically by a lossless arithmetic encoding algorithm.

For decompression, the compressed file is decoded to recreate the wavelet transform up to the number of levels performed in the compression step described above. The inverse wavelet transform produces an image which is for example ¼ of the original image size. This image is then enlarged using the method of the invention as described above in connection with image enlargement. This results in a full size reproduced image.

This approach results in higher levels of compression and reproduced image quality than is possible without the use of the novel expansion technique.

For data representing a video signal, the high frequency coefficients of the first two levels of the wavelet transform of each frame are discarded resulting in frames that are 1/16 of the original size. Thus, at 30 frames per second, the amount of data to be processed per second is 480 times smaller than with state-of-the-art approaches (e.g., MPEG-1, MPEG-2, H.231, H.234). This results in extremely high compression ratios of up to 10,000:1, enough to make possible the real-time transmission of full-size, full-color, full-motion video through telephone lines using 28.8 Kbps modems. For decompression, the two-level expansion of each reproduced frame at 1/16 the original size, produces a full-size video sequence of high quality.

For data representing an acoustic signal, the ability to estimate the high frequency coefficients of the wavelet transform of the one-dimensional sound signal from the low frequency coefficients also results in high levels of compression and improved signal quality.

In some applications, processing speed is critical and compression ratio is of secondary importance. For example, in scanners, printers and photocopiers, the ability to perform compression in real-time is important. These applications call for an embodiment of the invention in which the Haar wavelet transform is used to compress the image on a line-by-line or block-by-block basis. Because of the simplification introduced by the use of the Haar wavelet transform, convolution of the input signal by the filter coefficients reduces to the multiplication of the sum of two adjacent pixel values by the filter coefficients followed by a one pixel shift. This procedure is one that can readily be implemented in hardware with its concomitant increase in performance.

To use the retained low frequency components of the wavelet transform to recover the missing high frequency components, the system according to the invention

optionally passes the retained frequency components to a low-pass synthesis filter which is biorthogonal to the low-pass analysis filter used to generate the wavelet transform of the image. This results in the generation of the low-frequency sub-band of the original image. The retained frequency components are also passed to an estimation system which generates an estimate of the high-frequency sub-band of the image. The low frequency sub-band and the high frequency sub-band are then combined at a combining stage to form the original image.

In one optional embodiment, the estimation system includes an estimation filter having a transfer function derived from the wavelet transforms high and low frequency analysis filters and their corresponding biorthogonal synthesis filters. The output of this estimation filter is then filtered by the wavelet transforms high-frequency synthesis filter before being combined, at the combining stage, with the output of the low-frequency synthesis filter.

In another optional embodiment of the invention, the output of the estimation filter is used as a starting estimate which is iteratively refined at a refining stage. The preferred iterative method is a conjugate gradient method. Preferably, those high frequency coefficients which were retained rather than discarded are clamped at their known values during successive iterations of the process executed by the refining stage. The output of the refining stage is then filtered by the wavelet transforms high-frequency synthesis filter before being combined, at the combining stage, with the output of the low-frequency synthesis filter.

In yet another embodiment of the invention, used in the case in which there are no known high frequency coefficients, the estimation system includes an estimation filter having a transfer function derived from the wavelet transforms high and low frequency analysis filters and their corresponding biorthogonal synthesis filters. The output of this estimation filter is then combined, at the combining stage, with the output of the low-frequency synthesis filter. In this case, there is no need to filter the output of the estimation filter with the wavelet transforms high-frequency synthesis filter before passing it to the combining stage.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features and advantages of the invention will be apparent from the following description and the accompanying drawings, in which like reference characters refer to the same parts throughout the different views.

FIG. 1 shows a system for evaluating the wavelet transform of an original image, compressing the wavelet transform after discarding is high frequency components, and

then recovering the discarded high frequency components to generate a reconstructed image;

FIG. 2 shows the data recovery stage of FIG. 1;

FIG. 3 shows an embodiment of the data recovery stage shown in FIG. 1 in which the estimation system processes only the low frequency coefficients to obtain an estimate of the original high frequency coefficients;

FIG. 4 shows a prior art process for reconstructing an image based on the high frequency and low frequency components of its wavelet transform;

FIG. 5 shows the data recovery stage of FIG. 1 in which the estimation system relies on known high frequency coefficients to iteratively converge on an estimate of the missing high frequency coefficient;

FIG. 6 shows the data recovery stage of FIG. 1 in which uses a single filter is used to estimate high frequency components with low frequency components;

FIG. 7 shows an image to be processed by the system of FIG. 1;

FIG. 8 shows the image of FIG. 7 after application of the high and low frequency analysis filters to the rows of the image;

FIG. 9 shows the image of FIG. 8 after application of the high and low frequency analysis filters to the columns of the image;

FIG. 10 shows the image of FIG. 7 after three levels of the transformation described in FIGS. 8 and 9;

FIG. 11a shows an image such as that shown in FIG. 7 with its wavelet transform coefficients separated into three sub-bands;

FIG. 11b shows two alternative paths to reconstructing the left image *L* from quadrants *a* and *b*; and

FIG. 11c shows two alternative paths to reconstructing the complete image *I* using the two sub-bands *L* and *R*.

DETAILED DESCRIPTION

System Overview

Referring to FIG. 1, a system **29** incorporating the invention includes a wavelet transform stage **17** having a low frequency analysis filter **10** and a high frequency analysis filter **14** for evaluating the wavelet transform of an original image **11**. The original image **11** can be provided by a scanner, a digital camera, a photocopy machine or any other device generating a digital signal representative of an image. The source of the original image **11** can also be another wavelet transform stage having a pair of analysis filters identical to the high and low frequency synthesis filters **10, 14** shown in FIG. 1.

The wavelet transform generated by the wavelet transform stage **17** includes a low frequency portion **12** generated by convolving the original image with the low frequency analysis filter **10** and a high frequency portion **18a** generated by convolving the original image through the high frequency analysis filter **14**.

Because human visual perception depends primarily on the low frequency components of the image, it is desirable, when compressing an image, to disregard selected high frequency components of the image. This is accomplished by a thresholding stage **16** which compares the high frequency portion **18a** of the original image **11** with a preselected threshold value. If the wavelet transform coefficient associated with a portion of the image falls on one side of the threshold, that coefficient is disregarded, typically by setting it to zero. If it falls on the other side, that coefficient is retained, typically by passing it through unchanged. This thresholding process results in a diminished high frequency portion **18**.

The diminished high-frequency portion **18** and the low frequency portion **12** of the original image **11** are compressed in a conventional manner at a compression stage **13**. The compressed image can then be stored or transmitted in its compressed form.

To recover the original image, the compressed image must be decompressed and the disregarded high frequency portions of the original image must be estimated. The decompression step is performed in a conventional fashion and is not shown in FIG. 1.. The process of estimating the disregarded high frequency portions to generate a faithful reproduction of the original image **11a** takes place in the data recovery stage **19** which is described below in connection with FIGS. 2-6.

Evaluation of wavelet transform

The wavelet transform of an original image **11** is obtained by first convolving each row of the image with two orthogonal filters: a high-pass filter **14** and a low-pass filter **10** as depicted in FIG. 1. These filters, which are also referred to as high and low frequency analysis filters respectively, are obtained from the coefficients of the scaling function defining the wavelet basis used in the transformation. Thus, the effect of the wavelet transform is to determine the high frequency and low frequency spatial energy distribution of the image. The low frequency spatial distribution of the original image **11** is represented by the low-pass filtered image **12** and the high frequency spatial distribution is represented by the shown in a high-pass filtered image **18**.

FIG. 1 illustrates a wavelet transform stage **17** for performing one step in the performance of successive wavelet transforms. Each row of the image **11** is convolved with a low-pass filter **10**. However, the convolution proceeds by shifting by two pixels

rather than by a single pixel. This results in a low-pass filtered image **12** half as wide as the original image **11**.

Each row of the image **11** is also convolved with a high pass filter **14** biorthogonal to the low-pass filter **10**. The convolution again proceeds by shifting two pixels at a time rather than by a single pixel. This shift results in a high-pass filtered image **18a** that is likewise half as wide as the original image.

The high-frequency wavelet transform coefficients generating the high-pass filtered image **18** are then compared with a preselected threshold at a thresholding stage **16**. Those coefficients falling to one side of this threshold are set to zero. The remaining coefficients are left unchanged. The reduced set of high frequency wavelet transform coefficients thus formed generates a post-threshold high-pass filtered image **18** half as wide as the original image **11** and containing those high frequency components that are important for human perception.

Application of the above wavelet transform to the rows of the original image **11** shown in FIG. 7 results in the two images shown in FIG. 8. The low-pass filtered image, shown on the left hand side of FIG. 8 is clearly recognizable as a distorted version of the image in FIG. 7. The high-pass filtered image, shown on the right-hand side of FIG. 8, is, however, only barely recognizable. This is because, as mentioned above, it is the low frequency components of an image that are most important for human visual perception.

To complete the first level of the wavelet transform of the original image **11**, the wavelet transform stage **17** convolves the high and low frequency analysis filters **10, 14** with the columns of the original image **11**. The resulting four images, shown in FIG. 9, represent the wavelet transform coefficients present in four frequency bands. The lowest frequencies are in the upper-left image, the highest are in the lower-right image, and intermediate frequencies are in the remaining two images. That the upper-left image **16** is the most recognizable is also no coincidence since, as stated above, it is the low-frequency components of the image that are most important for human perception.

Although not explicitly shown in FIG. 1, it is readily apparent that the foregoing process can be repeated several times resulting in a multiple frequency sub-band decomposition. FIG. 9 shows an original image and FIG. 10 shows its wavelet transform carried out to four levels of transformation. The many black areas in FIG. 5 represent wavelet transform coefficients that are either zero or very small. The abundance of such coefficients is useful in providing a more compact representation of the original image **11**.

The significance of such a transformation from the standpoint of compression lies in the fact that information important for human perception is concentrated in the frequency sub-bands near the upper left corner of FIG. 10. The other coefficients can be

discarded without a major impact on the perceived quality of the reproduced image obtained by applying the inverse wavelet transform to the remaining coefficients.

The choice of filter coefficients for the analysis filters **10, 14** determines the wavelet basis used in the transform. This basis must be well localized in both spatial and frequency domains and, in order to avoid redundancy that hinders compression, it must constitute a biorthogonal set.

In one aspect of the invention, the filter coefficients for the analysis filters **10, 14** are chosen to implement the Haar wavelet transform. This is accomplished by choosing the filter coefficients for the high pass filter **14** to be 0.5 and –0.5 and the filter coefficients for the low pass filter **10** to be 0.5 and 0.5. In this embodiment, convolution of the original image **11** can be accomplished efficiently on a scan line-by-scan line basis by halving the sum (or the difference in the case of the high-pass filter), of two adjacent pixel values, shifting one pixel, and repeating the process. Although the compression ultimately achieved by the use of the Haar wavelet transform is not optimal, the foregoing convolution algorithm is amenable to economical implementation in hardware. For applications in which real-time compression is critical, for example in photocopiers, scanners, digital cameras, or printers, the loss in compression ratio is more than offset by the increased throughput resulting from a hardware implementation of the convolution step.

The foregoing hardware implementation can occur either in the course of compression or decompression. For example, in a scanner, it may be more convenient to implement the convolution in hardware at the scanner and to decompress the resulting image at the host in software. In a printer, it may be more convenient to perform the compression in software at the host and to implement the decompression step in hardware at the printer. In a photocopier, which can be thought of as a scanner and printer working together, one can save memory and improve performance by implementing both the compression and the decompression in hardware.

Establishing the threshold for coefficients

The threshold stage **16** incorporates a preselected threshold for determining whether or not a particular high frequency component is to be kept. The selection of this threshold requires consideration of the frequency dependent characteristics of human perception to determine what transform coefficients to keep in order to achieve a particular compression ratio.

It is well known that human vision is more sensitive to low and medium spatial frequencies than it is to high frequencies. This is why the coefficients associated with the

upper left image of FIG. 9 namely the low frequency coefficients, are the most important.

It is also well known that if a detail (edge) is barely visible at a particular frequency, its contrast with respect to grey (intensity) must be quadrupled when the frequency is doubled for that detail to remain distinguishable. Since in the dyadic wavelet transform each frequency sub-band width is one half the width of the next higher frequency sub-band, we can ignore those coefficients of the higher frequency sub-band having absolute values less than four times the maximum absolute value of the coefficients that were discarded in the next lower frequency sub-band. This establishes the proper relationship between the thresholds **16** used to discard coefficients in all the frequency sub-bands.

The combination of these two principles leads to the following scheme to achieve high compression with minimal loss of important observable information:

(1) Keep all the coefficients of the lowest frequency sub-band.

(2) Establish a threshold for the next sub-band as a fraction of the absolute value of the largest coefficient of the sub-band.

(3) Multiply this threshold by 4 for each higher sub-band.

(4) Discard coefficients having absolute values below the corresponding thresholds.

The following values are used to compute the actual threshold values using the maximum coefficient values in each sub-band. These values are for four levels of transformation:

Level 4 : 0.01
Level 3 : 0.04
Level 2 : 0.16
Level 1 : 0.64

Using the above values, the method discards all high-pass coefficients of level 4 that are below 1% of the maximum absolute value of the coefficients of level 4. The other thresholds are 4% of the maximum for level 3, 16% of the maximum for level 2, and 64% of the maximum for level 1.

The above values normally result in low to moderate compression, i.e. 30:1 to 40:1, with the quality of the reconstructed image being almost identical to that of the original. For higher compression, the above values are increased proportionately subject to the constraint that the ratio between consecutive levels should be 4 and that no value

can be greater than 1.0. At high compression ratios, this results in the elimination of all the coefficients of levels 1 and 2 but with a very slow and graceful degradation of the quality of the reconstructed image after decompression.

Encoding the Coefficients

Once the wavelet transform of the original image **11** has been decomposed into the desired number of sub-bands as described above, the compression stage **13** encodes, using the fewest number of bits, the remaining wavelet transform coefficients associated with each sub-band. Two values must be encoded: the location within the sub-band, and the value (including the sign) of each wavelet transform coefficient.

A coefficient's location within the sub-band is expressed as the distance in rows to either the previous non-zero coefficient or, in the case of the first non-zero coefficient of the sub-band, the distance to the upper left corner of the sub-band. These location values must be encoded exactly.

A coefficient's values can be encoded efficiently by dividing the interval between the maximum and minimum threshold values into quantization bins. If the number of quantization bins is large enough, given the difference between the maximum and minimum absolute values at each level, the quantization error will not be noticeable. A preferred value for the number of quantization bins in this embodiment of the invention is 32.

The above discussion applies to both the luminance and chrominance components of an image. However, humans are much less sensitive to changes in the chrominance than they are to changes in luminance. As a result, only the lowest frequency coefficients of the chrominance components need be kept. This permits very high compression for the chrominance components. This will only work, i.e., give good quality colors after the decompression, when it is combined with the unique wavelet transform reconstruction procedure described below.

Once all the coefficients are encoded into a binary file, the next step is to apply a lossless coding scheme such as arithmetic coding to obtain the final compressed binary file. Note that if no coefficients from the highest levels are kept, the size of the compressed file can be reduced significantly. No codebook or dictionary is involved in the above scheme. Consequently the scheme is general and applicable to any kind of data.

Decompression

The first step in decompression of the compressed image is to arithmetically decode the binary compressed file. Then the coefficient values and locations are

calculated and the wavelet transform of the original data, in which most if not all coefficients of the higher frequency sub-bands are zero, is recreated.

Recovery of discarded coefficients

In order to improve the quality of the reconstructed data, all the coefficients that were discarded for compression must be estimated. The coefficients to be estimated include the high frequency coefficients that were disregarded at the thresholding stage **16** as well as low frequency coefficients generated by the convolution of an analysis filter **10, 14** with one or more such coefficients. The remainder of this section describes the data recovery stage **19** for accomplishing this task.

FIG. 2 shows a data recovery stage **19** according to the invention having a low frequency synthesis filter **22** which corresponds to the low frequency analysis filter **10** of the wavelet transform stage **17** and an estimation system **38**. The low frequency wavelet transform coefficients **12** are convolved with the low frequency analysis filter **10** and the result of the convolution is passed to a combining stage **23**.

The low frequency wavelet transform coefficients **12** are also passed to an estimation system **28** which estimates the discarded wavelet transform coefficients from the low frequency wavelet transform coefficients. In one embodiment of the invention, to be described in connection with FIG. 5, the estimation system also uses the high frequency wavelet transform coefficients **18** that exceeded the threshold at the thresholding stage **16**.

Wavelets are functions generated from a single function Ψ by dilations and translations.

(1) $$\Psi_n^j(x) = 2^{-1/2}\Psi(2^{-j}x - n)$$

where j corresponds to the level of the transform, and hence governs the dilation, and n governs the translation.

The basic idea of the wavelet transform is to represent an arbitrary function f as a superposition of wavelets.

(2) $$f = \sum_{j,n} a_n^j(f)$$

Since the Ψ_n^j constitute an orthonormal basis, the wavelet transform coefficients are given by the inner product of the arbitrary function f and the wavelet basis functions:

(3) $$a_n^j(f) = \left\langle \Psi_n^j, f \right\rangle$$

In a multiresolution analysis, one really has two functions: a mother wavelet Ψ and a scaling function φ. Like the mother wavelet, the scaling function φ generates a family of dilated and translated versions of itself:

(4) $$\varphi_n^j(x) = 2^{-j/2}\varphi(2^{-j}x - n)$$

When compressing data files representative of images, it is important to preserve symmetry. As a result, the requirement of an orthonormal basis is relaxed and biorthogonal wavelet sets are used. In this case, the Ψ_n^j no longer constitute an orthonormal basis, hence the computation of the coefficients a_n^j is carried out via the dual basis

(5) $$a_n^j(f) = \left\langle \tilde{\Psi}_n^j, f \right\rangle$$

where $\tilde{\Psi}$ is a function associated with the corresponding synthesis filter coefficients defined below.

When f is given in sampled form, one can take these samples as the coefficients x_n^j for sub-band j. The coefficients for sub-band $j+1$ are then given by the convolution sums:

(6a) $$x_n^{j+1} = \sum_k h_{(2n-k)} x_k^j$$

(6b) $$c_n^{j+1} = \sum_k g_{(2n-k)} x_k^j$$

This describes a sub-band algorithm with:

(7a) $$h_n = 2^{1/2} \int \varphi(x-n)\varphi(x)dx$$

representing a low pass filter and

(7b) $$g_\ell = (-1)^\ell h_{-\ell+1}$$

representing a high pass filter. Consequently, the exact reconstruction is given by:

$$(8) \qquad x_\ell^j = \sum_n \left[\tilde{h}_{2n-\ell} x_n^{j+1} + \tilde{g}_{2n-\ell} c_n^{j+1} \right]$$

The relation between the different biorthogonal filters is given by:

$$(9a) \qquad \tilde{g}_n = (-1)^n h_{-n+1}$$

$$(9b) \qquad g_n = (-1)^n \tilde{h}_{-n+1}$$

$$(9c) \qquad \sum_n h_n \tilde{h}_{n+2k} = \delta_{k,0}$$

where h_n and g_n represent the low-pass analysis filter and the high-pass analysis filter respectively, and $\tilde{h}_n$ and $\tilde{g}_n$ represent the corresponding synthesis filters.

We now turn to a matrix modified formulation of the one-dimensional biorthogonal wavelet transform. Using the above impulse responses h_n and g_n we can define the circular convolution operators at resolution 2^{-j}:$\mathbf{H}^j, \mathbf{G}^j, \tilde{\mathbf{H}}^j, \tilde{\mathbf{G}}^j$. These four matrices are circulant and symmetric.

The fundamental matrix relation for exactly reconstructing the data at resolution 2^{-j} is

$$(10) \qquad \tilde{\mathbf{H}}^j \mathbf{H}^j + \tilde{\mathbf{G}}^j \mathbf{G}^j = \mathbf{I}^j$$

where $\mathbf{I}^j$ is the identity matrix.

Let $\underline{\mathbf{x}}^{j+1}$ be a vector of low frequency wavelet transform coefficients at scale $2^{-(j+1)}$ and let $\underline{\mathbf{c}}_x^{j+1}$ be the vector of associated wavelet coefficients. We have, in augmented vector form:

$$(11) \qquad \begin{bmatrix} \underline{\mathbf{x}}^{j+1} \\ \underline{\mathbf{c}}_x^{j+1} \end{bmatrix} = \begin{bmatrix} \mathbf{H}^j & 0 \\ 0 & \mathbf{G}^j \end{bmatrix} \begin{bmatrix} \underline{\mathbf{x}}^j \\ \underline{\mathbf{x}}^j \end{bmatrix}$$

where $\underline{\mathbf{x}}^{j+1}$ is the smoothed vector obtained from $\underline{\mathbf{x}}^j$ and the wavelet coefficients $\underline{\mathbf{c}}_x^{j+1}$ contain information lost in the transition between scales 2^{-j} and $2^{-(j+1)}$.

(12) $\underline{\mathbf{x}}^{j} = \begin{bmatrix} \tilde{\mathbf{H}}^{j} & \tilde{\mathbf{G}}^{j} \end{bmatrix} \begin{bmatrix} \underline{\mathbf{x}}^{j+1} \\ \underline{\mathbf{c}}_{x}^{j+1} \end{bmatrix}$

A prior art system for implementing equation (12), depicted schematically in FIG. 4, provides a way to recover $\underline{\mathbf{x}}^{j}$ given $\underline{\mathbf{x}}^{j+1}$ and $\underline{\mathbf{c}}_{x}^{j+1}$. FIG. 4 shows low-pass filtered wavelet transform coefficients representative of the low frequency image **12** being passed through a low-frequency synthesis filter **22** corresponding to the low-pass filter **10**. Similarly, high-pass filtered wavelet transform coefficients representative of the high-frequency image **18** are passed through a high-frequency synthesis filter **21**. The outputs of both synthesis filters **21, 22** are combined at a combining stage **23** to generate the original image **11**. However, if some of the high frequency coefficients $\underline{\mathbf{c}}_{x}^{j+1}$ have been discarded, then $\underline{\mathbf{x}}^{j}$ will lack details that would have been provided by the missing $\underline{\mathbf{c}}_{x}^{j+1}$.

Since, from equation (11), $\underline{\mathbf{x}}^{j+1} = \mathbf{H}^{j}\underline{\mathbf{x}}^{j}$ we can, in principle, recover $\underline{\mathbf{x}}^{j}$ from $\underline{\mathbf{x}}^{j+1}$ merely by inverting $\mathbf{H}^{j}$. However, this is generally not practical both because of the presence of inaccuracies in $\underline{\mathbf{x}}^{j+1}$ and because $\mathbf{H}^{j}$ is generally an ill-conditioned matrix. As a result, the above problem is ill-posed and there is, in general, no unique solution.

If we discard the high frequency coefficients, $\underline{\mathbf{c}}_{x}^{j+1}$, then equation (12) reduces to $\underline{\mathbf{x}}^{j} = \tilde{\mathbf{H}}^{j}\underline{\mathbf{x}}^{j+1}$. This results in $\underline{\mathbf{y}}^{j}$, a blurred approximation of $\underline{\mathbf{x}}^{j}$ From equation (12), we have:

(13) $\underline{\mathbf{y}}^{j} = \tilde{\mathbf{H}}^{j}\underline{\mathbf{x}}^{j+1}$

which, when combined with equation (11), reduces to

(13a) $\tilde{\mathbf{H}}^{j}\underline{\mathbf{x}}^{j+1} = \tilde{\mathbf{H}}^{j}\mathbf{H}^{j}\underline{\mathbf{x}}^{j}$

During decompression, the $\underline{\mathbf{x}}^{j+1}$ (transformed rows or columns of level $j+1$) are known and the problem is to determine the $\underline{\mathbf{x}}^{j}$ of the next higher level. Equation (13a) can be written

(14) $\underline{\mathbf{x}}^{j+1} = \mathbf{H}^{j}\underline{\mathbf{x}}^{j}$

This can be thought of as an image restoration problem in which the image defined by the vector $\underline{\mathbf{x}}^j$ has been blurred by the operator $\mathbf{H}^j$, which, due to its low pass nature, is an ill-conditioned matrix.

Regularization, as disclosed in "Méthodes de résolution des problèmes mal-posés" by A.N. Tikhanov and V.Y. Arsenin, Moscow, Edition MIR, hereinafter incorporated by reference, is a method used to solve ill-posed problems of this type. This method is analogous to a constrained least squares minimization technique.

A solution for this type of problem is found by minimizing the following Lagrangian function:

$$(15) \quad \mathbf{J}(\underline{\mathbf{x}}^j,\alpha) = \left|\underline{\mathbf{x}}^{j+1} - \mathbf{H}^j\underline{\mathbf{x}}^j\right|^2 + \alpha\left|\mathbf{G}^j\underline{\mathbf{x}}^j\right|^2$$

where $\mathbf{G}^j$ is the regularization operator and α is a positive scalar such that $\alpha\rightarrow 0$ as the accuracy of $\underline{\mathbf{x}}^{j+1}$ increases.

It is also known from regularization theory that if $\mathbf{H}^j$ acts as a low-pass filter, $\mathbf{G}^j$ must be a high-pass filter. In other words, since $\mathbf{H}^j$ is the low-pass filter matrix of the biorthogonal wavelet transform, $\mathbf{G}^j$, must be the corresponding high-pass filter matrix.

Equation (15) may be also written with respect to the estimated wavelet transform coefficients $\underline{\mathbf{c}}_x^{j+1}$ and $\underline{\hat{\mathbf{x}}}^{j+1}$.

$$(16) \quad \mathbf{J}(\underline{\mathbf{x}}^j,\alpha) = \left|\underline{\mathbf{x}}^{j+1} - \underline{\hat{\mathbf{x}}}^{j+1}\right|^2 + \alpha\left|\underline{\mathbf{c}}_x^{j+1}\right|^2$$

Using the exact reconstruction matrix relation shown in Equation 10, we get:

$$(16a) \quad \underline{\mathbf{x}}^{j+1} = \tilde{\mathbf{H}}^j\mathbf{H}^j\underline{\mathbf{x}}^{j+1} + \tilde{\mathbf{G}}^j\mathbf{G}^j\underline{\mathbf{x}}^{j+1}$$

Also, we can write

$$(16b) \quad \underline{\hat{\mathbf{x}}}^{j+1} = \mathbf{H}^j\underline{\mathbf{x}}^j = \mathbf{H}^j\left(\mathbf{H}^j\underline{\mathbf{x}}^{j+1} + \tilde{\mathbf{G}}^j\underline{\mathbf{c}}_x^{j+1}\right)$$

Then, subtracting (16b) from (16a) gives:

$$(16c) \quad \underline{\mathbf{x}}^{j+1} - \underline{\hat{\mathbf{x}}}^{j+1} = \mathbf{G}^j\tilde{\mathbf{G}}^j\underline{\mathbf{x}}^{j+1} - \mathbf{H}^j\tilde{\mathbf{G}}^j\underline{\mathbf{c}}_x^{j+1}$$

Substituting (16c) into (16) results in:

$$(17)\quad \mathbf{J}(\underline{\mathbf{c}}_x^{j+1},\alpha)=\left|\mathbf{G}^j\tilde{\mathbf{G}}^j\underline{\mathbf{x}}^{j+1}-\mathbf{H}^j\tilde{\mathbf{G}}^j\underline{\mathbf{c}}_x^{j+1}\right|^2+\alpha\left|\underline{\mathbf{c}}_x^{j+1}\right|^2$$

By setting the derivative of **J** with respect to $\underline{\mathbf{c}}_x^{j+1}$ equal to zero, we obtain the following estimate for the high frequency coefficients $\underline{\mathbf{c}}_x^{j+1}$:

$$(18)\quad \underline{\mathbf{c}}^{j+1}=\mathbf{M}\underline{\mathbf{x}}^{j+1}$$

where the estimation matrix **M** is given by

$$(19)\quad \mathbf{M}=\left(\alpha\mathbf{I}^j+\tilde{\mathbf{G}}^{jT}\mathbf{H}^{jT}\mathbf{H}^j\tilde{\mathbf{G}}^j\right)^{-1}\tilde{\mathbf{G}}^{jT}\mathbf{H}^{jT}\mathbf{G}^j\tilde{\mathbf{G}}^j$$

in which "**T**" refers to the matrix transpose.

A data recovery stage **19** for implementing equations (18) and (19), shown in FIG. 3, provides a way to estimate the high frequency components $\underline{\mathbf{c}}_x^{j+1}$ of the image using only the low frequency components $\underline{\mathbf{x}}^{j+1}$ and matrices derived from the known properties of the two orthogonal filters: $\mathbf{G}^j$, $\tilde{\mathbf{G}}^j$, $\mathbf{H}^j$, and $\tilde{\mathbf{H}}^j$. FIG. 3 shows the wavelet transform coefficients representative of the low-frequency image **12** being passed through a low-frequency synthesis filter **22** as they were in FIG. 4. However, unlike the system of FIG. 4, these same wavelet transform coefficients are used to estimate the high-frequency coefficients representative of the high-frequency image **18**. This is accomplished by passing the low-frequency coefficients to an estimation system **28** consisting of an estimation filter **24** followed by a high-frequency synthesis filter **21**. The matrix **M** associated with the estimation filter **24** can be precalculated for the selected biorthogonal wavelet set. The output of the estimation filter **24** is then filtered by the high-frequency synthesis filter **21**. The outputs of both synthesis filters **21, 22** are then combined at the combining stage **23** as was the case in FIG. 7. The combining stage output represents an estimate of the original image **11a**.

Refining the estimate of the high-frequency coefficients

The estimation system **28** of FIG. 3 provides a good initial estimate of $\underline{\mathbf{c}}_x^{j+1}$, the missing wavelet transform high-frequency coefficients. In another aspect of the invention, this estimate can be further refined by an iterative conjugate gradient algorithm using the above initial estimate of $\underline{\mathbf{c}}_x^{j+1}$ and an initial search direction given by the gradient vector $\nabla\mathbf{J}\left(\underline{\mathbf{c}}_x^{j+1},\alpha\right)$.

The search for the global minimum of **J** is greatly helped by clamping the known values of the vector $\underline{\mathbf{c}}_x^{j+1}$. FIG. 5 shows a data recovery stage **19** incorporating this clamping function.

The illustrated data recovery stage **19** is similar to that depicted in FIG. 3 with the exception that the estimation system **28** includes a refinement stage **25** implementing the conjugate gradient method interposed between the estimation filter **24** and the high-frequency synthesis filter **21**. The refinement stage **25** accepts known values of the high-frequency wavelet transform coefficients **18** and clamps then at those values throughout the iterations of the conjugate gradient algorithm.

If there are no known values of $\underline{\mathbf{c}}_x^{j+1}$ at a given row or column at a given wavelet transform level, then the actual values of the inverse wavelet transform, i.e., the $\mathbf{x}^j$ can be calculated directly without first calculating the $\underline{\mathbf{c}}_x^{j+1}$

Since

$$(20) \quad \underline{\mathbf{x}}^j = \begin{bmatrix} \tilde{\mathbf{H}}^j & \tilde{\mathbf{G}}^j \end{bmatrix} \begin{bmatrix} \underline{\mathbf{x}}^{j+1} \\ \underline{\mathbf{c}}_x^{j+1} \end{bmatrix}$$

we can rewrite equation (20) as

$$(21) \quad \underline{\mathbf{x}}^j = \begin{bmatrix} \tilde{\mathbf{H}}^j & \mathbf{T} \end{bmatrix} \begin{bmatrix} \underline{\mathbf{x}}^{j+1} \\ \underline{\mathbf{x}}^{j+1} \end{bmatrix}$$

where the matrix **T** is given by

$$(22) \quad \mathbf{T} = \tilde{\mathbf{G}}^j \left(\alpha \mathbf{I}^j + \tilde{\mathbf{G}}^{jT} \mathbf{H}^{jT} \mathbf{H}^j \tilde{\mathbf{G}}^j \right)^{-1} \tilde{\mathbf{G}}^{jT} \mathbf{H}^{jT} \mathbf{G}^j \tilde{\mathbf{G}}^j$$

FIG. 6 depicts a data recovery stage **19** for implementing the foregoing method. The illustrated data recovery stage **19** is similar to that shown in FIG. 3 with the exceptions that the high-pass synthesis filter **21** is no longer necessary and that the estimation filter **26** incorporates the matrix **T** given in Equation (22). To reduce computation time, **T** is precalculated for a given biorthogonal wavelet set. In the system of FIG. 6, the output of the estimation filter **26** is passed directly to the combining stage **23**.

The data recovery stage **19** of FIG. 6 is particularly useful for enlarging an image. In such a case, enlarging an image can be accomplished by inserting additional

wavelet transform coefficients between known wavelet transform coefficients. The values of these additional wavelet transform coefficients can be estimated using the data recovery stage **19** of FIG. 6.

The above expression for $\underline{x}^j$ is an initial estimate because of the parameter α. The actual vector (row or column) is again obtained by the conjugate gradient algorithm as described in connection with FIG. 5.

The decompression procedure is illustrated in FIG. 11 for one level of the wavelet transform of data representing an image. In FIG. 11a, quadrant *a* represents the low frequency sub-band and quadrant *b* and half *R* represent the higher frequency sub-bands in increasing order.

FIG. 11b shows the process of recovering the left side *L* of a given transform level. If *b* is empty, i.e., if there are no known high frequency coefficients, matrices **T** and $\tilde{\mathbf{H}}$ are used to compute the columns of *L* directly, one by one. If the most important coefficients of "*b*" are known, then matrix **M** is used to compute an initial estimate of a given column. This estimate is refined by the conjugate gradient method with clamping of the known coefficients to obtain a complete set *b'* of high frequency coefficients. The inverse wavelet transform on *a* and *b'* gives the left side *L*. By processing the left and right sides, *L* and *R*, respectively, by rows, we obtain the reproduction of the entire level which is either the low frequency component of the next level or the final decompressed image I if the level is 1, as shown in FIG. 11c.

This reconstruction process is applied to the luminance and chrominance components the only difference being that no clamping is normally required for the chrominance components since adequate estimates of the high frequency coefficients can be obtained from the low frequency coefficients alone. This approach results in higher compression and higher quality reproduced images than any other known method.

It will thus be seen that the invention efficiently attains the objects set forth above. Since certain changes may be made in the above constructions without departing from the scope of the invention, it is intended that all matter contained in the above description or shown in the accompanying drawings be interpreted as illustrative and not in a limiting sense.

It is also to be understood that the following claims are intended to cover all generic and specific features of the invention described herein, and all statements of the scope of the invention which as a matter of language might be said to fall there between.

Having described the invention, what is claimed as new and secured by Letters Patent is:

This Sheet Has Been Replaced





CLAIMS

1. A method for recovering information missing from a digital signal generated by evaluating the wavelet transform of a source image, said digital signal having a low frequency portion generated by filtering said source image with a low-frequency analysis filter and a high frequency portion, said method comprising the steps of:

filtering said low-frequency portion with a low-frequency synthesis filter corresponding to said low-frequency analysis filter, thereby generating a low-frequency sub-band of said source image;

processing said low-frequency portion with an estimation system, thereby generating an estimation system output representative of a high-frequency sub-band of said source image; and

combining said high-frequency sub-band with said low-frequency sub-band, thereby generating an estimate of said source image.

2. The method of claim 1 wherein said high frequency portion of said digital signal is generated by filtering said source image through a high-frequency analysis filter biorthogonal to said low-frequency analysis filter, and a portion of said high frequency portion is disregarded, thereby generating a diminished high-frequency portion and wherein said processing step further comprises the steps of:

filtering said low frequency portion with an estimation filter, thereby generating an estimation filter output; and

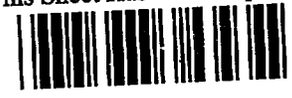
filtering said estimation filter output through a high-frequency synthesis filter corresponding to a high-frequency analysis filter biorthogonal to said low frequency analysis filter.

3. The method of claim 2 wherein said first filtering step comprises the step of filtering said low frequency portion with a filter having a matrix transfer function **M** given by

$$\mathbf{M} = \left(\alpha \mathbf{I}^j + \tilde{\mathbf{G}}^{jT}\mathbf{H}^{jT}\mathbf{H}^j\tilde{\mathbf{G}}^j\right)^{-1}\tilde{\mathbf{G}}^{jT}\mathbf{H}^{jT}\mathbf{G}^j\tilde{\mathbf{G}}^j$$

wherein $\mathbf{G}^j$ represents said high-frequency analysis filter, $\tilde{\mathbf{G}}^j$ represents said high-frequency synthesis filter, $\mathbf{H}^j$ represents a low-frequency analysis filter, $\tilde{\mathbf{H}}^j$ represents said corresponding low-frequency synthesis filter, $\mathbf{I}^j$ represents an identity matrix, and α is a number.

This Sheet Has Been Replaced



4. The method of claim 1 wherein said processing step comprises the step of filtering said low frequency portion with a filter having a matrix transfer function **T** given by

$$\mathbf{T} = \tilde{\mathbf{G}}^j\left(\alpha\mathbf{I}^j + \tilde{\mathbf{G}}^{jT}\mathbf{H}^{jT}\mathbf{H}^j\tilde{\mathbf{G}}^j\right)^{-1}\tilde{\mathbf{G}}^{jT}\mathbf{H}^{jT}\mathbf{G}^j\tilde{\mathbf{G}}^j$$

wherein $\mathbf{H}^j$ represents said low-frequency analysis filter, $\tilde{\mathbf{H}}^j$ represents said low-frequency synthesis filter, $\mathbf{G}^j$ represents a high-frequency analysis filter corresponding to said low frequency analysis filter, $\tilde{\mathbf{G}}^j$ represents a synthesis filter corresponding to $\mathbf{G}^j$, $\mathbf{I}^j$ represents an identity matrix, and α is a number.

5. The method of claim 1 further comprising the step of iteratively refining said estimate of said source image.

6. The method of claim 5 wherein said refining step comprises the steps of

supplying an initial estimate for initiating a conjugate gradient search; and

initiating a conjugate gradient search for an optimal estimate of said high frequency portion of said digital signal.

7. The method of claim 6 further comprising the step of clamping said estimation system output at values determined by said high frequency portion of said digital signal.

8. A system for recovering information missing from a digital signal generated by evaluating the wavelet transform of a source image, said digital signal having a low frequency portion generated by filtering said source image with a low-frequency analysis filter and a high frequency portion, said system comprising:

a low-frequency synthesis filter corresponding to said low-frequency analysis filter for filtering said low-frequency portion, to generate a low-frequency sub-band of said source image;

an estimation system for processing said low-frequency portion, thereby generating an estimation system output representative of a high-frequency sub-band of said source image; and

means for combining said high-frequency sub-band with said low-frequency sub-band, thereby generating an estimate of said source image.

9. The system of claim 8 wherein said high frequency portion of said digital signal is generated by filtering said source signal through a high-frequency analysis filter biorthogonal to said low-frequency analysis filter, and a portion of said high

frequency portion is disregarded, thereby generating a diminished high-frequency portion and wherein said estimation system further comprises:

an estimation filter for filtering said low frequency portion, thereby generating an estimation filter output; and

a high-frequency synthesis filter for filtering said estimation filter output, said high-frequency synthesis filter corresponding to a high-frequency analysis filter biorthogonal to said low frequency analysis filter.

10. The system of claim **9** wherein said estimation filter comprises means for filtering said low frequency portion with a filter having a matrix transfer function **M** given by

$$\mathbf{M} = \left(\alpha \mathbf{I}^j + \tilde{\mathbf{G}}^{jT}\mathbf{H}^{jT}\mathbf{H}^j\tilde{\mathbf{G}}^j\right)^{-1}\tilde{\mathbf{G}}^{jT}\mathbf{H}^{jT}\mathbf{G}^j\tilde{\mathbf{G}}^j$$

wherein $\mathbf{G}^j$ represents said high-frequency analysis filter, $\tilde{\mathbf{G}}^j$ represents said high-frequency synthesis filter, $\mathbf{H}^j$ represents a low-frequency analysis filter, $\tilde{\mathbf{H}}^j$ represents said corresponding low-frequency synthesis filter, $\mathbf{I}^j$ represents an identity matrix, and α is a number.

11. The system of claim **8** wherein said estimation system comprises means for filtering said low frequency portion with a filter having a matrix transfer function **T** given by

$$\mathbf{T} = \tilde{\mathbf{G}}^j\left(\alpha \mathbf{I}^j + \tilde{\mathbf{G}}^{jT}\mathbf{H}^{jT}\mathbf{H}^j\tilde{\mathbf{G}}^j\right)^{-1}\tilde{\mathbf{G}}^{jT}\mathbf{H}^{jT}\mathbf{G}^j\tilde{\mathbf{G}}^j$$

wherein $\mathbf{H}^j$ represents said low-frequency analysis filter, $\tilde{\mathbf{H}}^j$ represents said low-frequency synthesis filter, $\mathbf{G}^j$ represents a high-frequency analysis filter corresponding to said low frequency analysis filter, $\tilde{\mathbf{G}}^j$ represents a synthesis filter corresponding to $\mathbf{G}^j$, $\mathbf{I}^j$ represents an identity matrix, and α is a number.

12. The system of claim **8** further comprising means for iteratively refining said estimate of said source image.

13. The system of claim **12** wherein said refining means comprises:

means for supplying an initial estimate for initiating a conjugate gradient search; and

means for initiating a conjugate gradient search for an optimal estimate of said high frequency portion of said digital signal.

This Sheet Has Been Replaced

14. The system of claim 13 further comprising means for clamping said estimation system output at values determined by said high frequency portion of said digital signal.

15. A method of decompressing a compressed image, said compressed image containing both high frequency components and low frequency components, to obtain a high integrity reproduction of an original image, said method comprising the steps of:

performing a wavelet transform operation on the original image to generate a wavelet transformed image, said wavelet transformed image containing both high frequency components and low frequency components,

comparing the high frequency components to a threshold value to disregard selected high frequency components,

compressing the wavelet transformed image to generate the compressed image, and

decompressing the compressed image by

filtering the low frequency components of the compressed image,

estimating the selected high frequency components disregarded in the comparing step, and

combining the filtered low frequency components and the estimated high frequency components to construct the decompressed reproduction of the original image with high integrity.

16. The method of claim 15 wherein the original image includes high frequency components and low frequency components, and wherein said step of performing a wavelet transform operation on the original image comprises the steps of

filtering the high frequency components of the original image, and

filtering the low frequency components of the original image.

17. The method of claim 16 further comprising the step of performing the high frequency filtering step and the low frequency filtering step in parallel.

18. The method of claim 15 wherein said step of comparing the high frequency components comprises the step of retaining those components having values above the threshold value and disregarding those components having values below the threshold value.



27. The method of claim 15 wherein said step of performing a wavelet transform operation comprises the step of

performing a Harr wavelet transform by

setting a first high frequency filtering coefficients to about 0.5 and setting a second high frequency filtering coefficient to about -0.5, and

setting first and second low frequency filtering coefficients to be about 0.5.

28. The method of claim 15 wherein said comparing step comprises the step of selecting the threshold value as a function of the number of iterations of the wavelet transform.

29. The method of claim 15 wherein said step of estimating the disregarded high frequency components of the image comprises the step of employing a regularization operator.

30. The method of claim 15 wherein said step of estimating the disregarded high frequency components of the image comprises the step of using an estimator filter to estimate the high frequency components using only the low frequency components.

31. The method of claim 15 wherein said step of estimating the disregarded high frequency components of the image comprises the step of refining an initial estimate of the high frequency components.

32. The method of claim 15 wherein said step of estimating the disregarded high frequency components of the image comprises the step of refining the estimated high frequency components with an iterative conjugate gradient.

33. The method of claim 32 wherein said step of refining further comprises the steps of

providing an initial estimate of the disregarded high frequency components,

clamping known values of the high frequency components, and

iteratively refining the initial estimate.

34. A method of decompressing a compressed image, said compressed image containing both high frequency components and low frequency components, to obtain the original image, wherein the original image is compresses by a wavelet transform to generate a wavelet transformed image containing both high frequency

This Sheet Has Been Replaced





19. The method of claim 15, wherein said estimating step comprises the steps of

performing a matrix operation on the low frequency components to obtain an estimation filter output, and

performing a filtering operation on the estimation filter output to generate an estimate of the high frequency components.

20. The method of claim 15 wherein said estimating step comprises the step of

performing a matrix operation on the low frequency components to generate the corresponding high frequency components.

21. The method of claim 15 wherein said comparing step comprises the step of retaining selected high frequency components that meet the threshold value.

22. The method of claim 21 further comprising the step of combining the low frequency components with the retained high frequency components to generate the wavelet transformed image.

23. The method of claim 16 wherein said step of filtering the high frequency components of the original image comprises the steps of

performing a convolving filtering operation on the original image wherein during the convolution, the pixels of the image are shifted by two to produce a filtered image half the size of the original image.

24. The method of claim 23 wherein said step of filtering the low frequency components of the original image comprises the steps of

performing a convolving filtering operation on the original image, wherein during the convolution, the pixels of the image are shifted by two to produce a filtered image half the size of the original image.

25. The method of claim 16 wherein said low frequency filtering step and said high frequency filtering step use mutually biorthogonal filters.

26. The method of claim 15 wherein said step of comparing the high frequency components comprises the step of setting the value of those components below the threshold value generally to zero.

This Sheet Has Been Replaced



components, at least a portion of which are discarded, and low frequency components, said method comprising the steps

providing the compressed image, and

decompressing the compressed image by

filtering the low frequency components of the compressed image,

estimating the discarded high frequency components, and

combining the filtered low frequency components and the estimated high frequency components to construct a decompressed reproduction of the original image with high integrity and accuracy.

35. A system for decompressing a compressed image, said compressed image containing both high frequency components and low frequency components, to obtain a high integrity reproduction of an original image, said system comprising:

means for performing a wavelet transform operation on the original image to generate a wavelet transformed image, said wavelet transformed image containing both high frequency components and low frequency components,

means for comparing the high frequency components to a threshold value to disregard selected high frequency components,

means for compressing the wavelet transformed image to generate the compressed image, and

means for decompressing the compressed image, said decompressing means including

means for filtering the low frequency components of the compressed image,

means for estimating the selected high frequency components disregarded by the comparing means, and

means for combining the filtered low frequency components and the estimated high frequency components to construct the decompressed reproduction of the original image with high integrity.

36. The system of claim **35** wherein the original image includes high frequency components and low frequency components, and wherein said means for performing a wavelet transform operation on the original image comprises:

This Sheet Has Been Replaced

means for filtering the high frequency components of the original image, and

means for filtering the low frequency components of the original image.

37. The system of claim **36** further comprising parallel filtering means for performing the high frequency filtering step and the low frequency filtering step in parallel.

38. The system of claim **35** wherein said high frequency comparing means comprises means for retaining those components having values above the selected threshold and disregarding those components having values below the threshold value.

39. The system of claim **35**, wherein said estimating means comprises

means for performing a matrix operation on the low frequency components to obtain an estimation output, and

means for performing a filtering operation on the estimation filter output to generate an estimate of the high frequency components.

40. The system of claim **35** wherein said estimating means comprises means for performing a matrix operation on the low frequency components to generate the corresponding high frequency components.

41. The system of claim **35** wherein said comparing means comprises means for of retaining selected high frequency components that meet the threshold value.

42. The system of claim **41** further comprising the means for combining the low frequency components with the retained high frequency components to generate the wavelet transformed image.

43. The system of claim **36** wherein said high frequency filtering means comprises means for performing a convolving filtering operation on the high frequency components of the original image, said convolving means including means for shifting pixels of the image by two to produce a filtered image half the size of the original image.

44. The system of claim **43** wherein said low frequency filtering means comprises means for performing a convolving filtering operation on the low frequency components of

the original image, said convolving means including means for shifting pixels of the image by two to produce a filtered image half the size of the original image.

45. The system of claim 36 wherein said low frequency filtering means and said high frequency filtering means include mutually biorthogonal filters.

46. The system of claim 35 wherein said threshold comparing means comprises means for setting the value of those components below the threshold value generally to zero.

47. The system of claim 35 wherein said step of performing a wavelet transform operation comprises:

means for performing a Harr wavelet transform, said Harr transform means including

high frequency filtering means having a first high frequency filtering coefficient set to about 0.5 and a second high frequency filtering coefficient set to about -0.5, and

low frequency filtering means having a first and second low frequency filtering coefficients set to about 0.5.

48. The system of claim 1 wherein said comparing means comprises means for selecting the threshold value as a function of the number of iterations of the wavelet transform.

49. The system of claim 35 wherein said estimating means comprises means for employing a regularization operator.

50. The system of claim 35 wherein said estimating means comprises means for estimating the high frequency components using only the low frequency components.

51. The system of claim 35 wherein said estimating means comprises means for refining the estimated high frequency components.

52. The system of claim 35 wherein said estimating means comprises means for refining the estimated high frequency components with an iterative conjugate gradient.

53. The system of claim 52 wherein said refining means further comprises:

This Sheet Has Been Replaced

means for providing an initial estimate of the disregarded high frequency components, and

means for clamping known values of the high frequency components, and

means for iteratively refining the initial estimate.

54. A system for decompressing a compressed image, said compressed image containing both high frequency components and low frequency components, to obtain the original image, wherein the original image is compresses by a wavelet transform operation to generate a wavelet transformed image containing both high frequency components, at least a portion of which are discarded, and low frequency components, said system comprising:

means for providing the compressed image, and

means for decompressing the compressed image, said decompressing means including

means for filtering the low frequency components of the compressed image,

means for estimating the discarded high frequency components, and

means for combining the filtered low frequency components and the estimated high frequency components to construct a decompressed reproduction of the original image with high integrity and accuracy.

55. An apparatus for recovering discarded high frequency wavelet transform coefficients from a digital signal representative of a set a complete set of low frequency wavelet transform coefficients and an incomplete set of high frequency wavelet transform coefficients, both sets being derived from a biorthogonal wavelet transform consisting of a high pass filter and a low pass filter, said apparatus comprising:

means for selecting an operator to correspond to said high pass filter and to said low pass filter of said biorthogonal wavelet transform; and

means for applying said operator to said complete set of low frequency wavelet transform coefficients.

56. The apparatus of claim 55 wherein said complete set of low frequency wavelet transform coefficients and said incomplete set of high frequency wavelet transform coefficients are derived from application of a biorthogonal wavelet transform to an image obtained from a scanner.

This Sheet Has Been Replaced

57. A method for estimating the values of a first digital signal from the values from a second digital signal, said second digital signal obtained by

filtering said first digital signal with a low-pass filter to obtain its low frequency components,

filtering said first digital signal with a high-pass filter orthogonal to said low-pass filter to obtain its high frequency components, and

discarding all high frequency components having an amplitude smaller than a threshold,

said method comprising the steps of:

creating a first data vector, $\underline{\mathbf{x}}^{j+1}$, by filtering said first digital signal with a low-pass filter and discarding every other point,

creating a second data vector, $\underline{\mathbf{c}}^{j+1}$, by filtering said first digital signal with a high-pass filter and discarding every other point,

creating a unit vector having as many elements as there are values in said first digital signal, each element being unity,

creating a first matrix, $\mathbf{H}^j$, having

half as many rows as there are values in said first digital signal,

a first row corresponding to the circular convolution of said low-pass filter with said unit vector, and in which

each row corresponds to the previous row shifted two elements to the right,

creating a second matrix, $\mathbf{G}^j$, having

half as many rows as there are values in said first digital signal,

a first row corresponding to the circular convolution of said high-pass filter with said unit vector, and in which

each row corresponds to the previous row shifted two elements to the right,

creating a third matrix, $\tilde{\mathbf{G}}^j$, to reverse the operation of said second matrix $\mathbf{G}^j$,

creating a fourth matrix, $\tilde{\mathbf{H}}^j$, to reverse the operation of said first matrix $\mathbf{H}^j$,

determining a third data vector, $\hat{\underline{c}}^{j+1}$, and a positive scalar α to minimize

$$\left|\mathbf{G}^j\tilde{\mathbf{G}}^j\underline{\mathbf{x}}^{j+1}-\mathbf{H}^j\tilde{\mathbf{G}}^j\hat{\underline{c}}_x^{j+1}\right|^2+\alpha\left|\hat{\underline{c}}_x^{j+1}\right|^2,$$

whereby said third data vector corresponds to an estimate of the high frequency components of said first digital signal,

premultiplying said first data vector by said fourth matrix to create a first product,

premultiplying said third data vector by said third matrix to create a second product,

adding together said first and second product to create a vector $\mathbf{x}^j$,

whereby said vector $\mathbf{x}^j$ is an estimate of the values in said first digital signal.



This Sheet Has Been Replaced

