P28

NASA Case No. LEW-15,310-1

PRINT FIG. 1

NOTICE

The invention disclosed in this document resulted from research in aeronautical and space activities performed under programs of the National Aeronautics and Space Administration. The invention is owned by NASA and is, therefore, available for licensing in accordance with the NASA Patent Licensing Regulation (14 Code of Federal Regulations 1245.2).

To encourage commercial utilization of NASA-owned inventions, it is NASA policy to grant licenses to commercial concerns. Although NASA encourages nonexclusive licensing to promote competition and achieve the widest possible utilization, NASA will consider the granting of a limited exclusive license, pursuant to the NASA Patent Licensing Regulations, when such a license will provide the necessary incentive to the licensee to achieve early practical application of the invention.

Address inquiries and all applications for license for this invention to NASA Patent Counsel, Lewis Research Center, Mail Stop LE-LAW, 21000 Brookpark Road, Cleveland, OH 44135. Approved NASA forms for application for nonexclusive or exclusive license are available from the above address.

Serial No.: 08/167 107

Filed Date: 12/1

(NASA-Case-LEW-15310-1) NON-ORTHOGONAL SUBBAND/TRANSFORM CODER Patent Application (NASA. Lewis Research Center) 28 p

N94-29589

Unclas

G3/32 0003670

AWARDS ABSTRACT

NON-ORTHOGONAL SUBBAND/TRANSFORM CODER

The present invention is directed to a simplified digital subband coder/decoder. In the present invention a signal is fed into a coder. The coder uses a non-orthogonal algorithm that is simply implemented in the coder hardware. The simple non-orthogonal design is then used in the implementation of the decoder to decode the signal.

INVENTOR: Daniel R. Glover

EMPLOYER: NASA

EVALUATOR: Paul K. Wagner

Serial # 08/167,107
Filed: 12/15/93

NASA CASE NO. LEW-15,310-1 PATENT APPLICATION

NON-ORTHOGONAL SUBBAND/TRANSFORM CODER

ORIGIN OF THE INVENTION

The invention described herein was made by employees of the United States Government and may be manufactured and used by or for the Government for governmental purposes without the payment of any royalties thereon or therefore.

FIELD OF THE INVENTION

The present invention is directed to a simplified digital subband coder/decoder.

Various methods have been used to encode digital signals while maintaining the integrity of the original signal when decoding. Among these methods frequency subband coding has arisen as an excellent method when the encoded data needs to be transmitted. Subband coding enables the decomposition of a signal into different frequencies that are more easily compressed thereby enabling transmission at low data rates. Subband coding has been used for data compression since the late 1970's with Quadrature Mirror filters (QMF) making subbanding practical by eliminating aliasing effects. The first subband coding systems were built for audio systems. In 1985 subband coding was applied to image and video data.

In the prior art image data is coded into subbands by first applying it horizontally and then vertically to the two dimensional image data. It is also possible to use a tree structure to apply the same filter to data several times to break the signal into more subbands.

QMF filters have traditionally been used to filter image data, however, because of problems of stability, most of the applications of QMF's have been Finite Impulse Response (FIR) filters. Research has resulted in QMF's that allow perfect reconstruction of the original signal. The first QMF reconstruction filter recognized was a simple two-tap FIR filter bank with the lowpass filter given by $1+z^{-1}$ and the highpass filter by $1-z^{-1}$. Unfortunately, this filter has been dismissed as trivial because it lacks the frequency resolving power of higher order filters. The main disadvantage with higher order filters, using FIR and IIR filters, is that they require complex hardware to perform the filtering namely multipliers and accumulators. Many of these filters have been simulated and show excellent results, but the filters have not been reduced to hardware implementations for real time video because of the high throughput required.

It is therefore, an object of the present invention to provide hardware for real-time coding and decoding.

It is a further object of this invention to implement low complexity hardware while maintaining high speed filtering.

DESCRIPTION OF THE RELATED ART

Meyer-Ebrecht et al patent no. 4,463,377 discloses a system for coding and decoding picture element signals. A transform arrangement and a corresponding retransformation arrangement are formed by a plurality of auxiliary transformers/retransformers.

Ansari patent No. 4,918,524 discloses a subband coding technique wherein polyphase filter banks split a high sample rate signal into two sample streams at half the rate.

Reudink patent No. 4,891,840 discloses a method of subband signal transmission in which a wide band signal is separated into two or more narrower band channels which are subsequently modulated down to baseband.

SUMMARY OF THE INVENTION

The present invention is directed to a non-orthogonal subband/transform coder. The method and apparatus disclosed produces a digitally encoded signal which is a transformed version of an original signal. The digitally encoded signal maintains the integrity of the original signal but allows transmission at low data rates in progressive resolution. The transformed data is organized into frequency subbands with the lowest frequency subband being a low resolution version of the original signal. This low resolution version can then be used for browsing and progressive resolution retrieval of archival data.

In the present invention the original data is first converted to frequency subbands. The subband values are transformed coefficients and can be encoded for lossy or lossless compression with a method suited to the characteristics of the particular subband type. At the receiving end, the decoded signal is inverse transformed and recombined into a reconstruction of the original signal using methods of encoding such as coarse quantization of high bands, differential coding, predictive coding, transform coding, vector quantization, entropy coding and statistical coding.

BRIEF DESCRIPTION OF THE DRAWINGS

The advantages and novel features of the invention will be more fully apparent from the following detailed description when read in connection with the accompanying drawings in which;

Fig. 1 displays the data reduction and economics achieved using the non-orthogonal transform.

Fig. 2 displays a simplified three-channel subband analysis system.

Fig. 3 displays a simplified three-channel subband synthesis system.

Fig. 4 displays a nine band separable subband system.

Fig. 5 displays a two-dimensional four channel subband analysis filter bank using the non-orthogonal transform.

Fig. 6 displays a subband synthesis filter bank using the non-orthogonal transform.

Fig. 7 displays a video screen with the upper 9 pixels highlighted.

Fig. 8 displays the first nine pixels of the video screen depicted in Fig. 7.

DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention uses a perfect reconstruction filter set to perform the subband decimation and interpretation. The equations used to describe the subband decimation and interpolation filtering are the following:

Decimation	Interpolation
$H_o(z) = z^{-2}$	$F_o(z) = 1 + z^{-1} + z^{-2}$

$$H_1(z) = z^{-1} - z^{-2} \qquad F_1(z) = z^{-1}$$

$$H_2(z) = 1 - z^{-2} \qquad F_2(z) = z^{-2}$$

The subband decimation and reconstruction filters are a new simple set of filters for obtaining three subbands. The equivalent transform matrices for these filter banks are given below:

$$G = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & -1 \\ 1 & 0 & -1 \end{bmatrix}; \quad g = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

The components of the invention are: adders, buffers, registers, counters, memory, timing and address control logic.

Fig. 2 shows the basic subbanding (or transform) circuit block diagram. Three samples from the input data stream are processed at one time as a block. This is known as block transform coding. This block processing accomplishes the decimation that has been performed after filtering in prior subbanding art. The three samples are latched and passed on to the adders for processing and three new samples are read into the serial-to-parallel register.

The adders perform the filtering function. As a result of the arithmetic processing, the output of the adders will have twice the range of the input, requiring an additional bit per addition in the result. The additional bit must be accounted for in the design, in order to get NASA perfect reconstruction, but it can be dropped by a shift right, which is essentially dividing by 2, if a less than perfect reconstruction of the signal can be accepted (e.g., in lossy compression coding).

After the signal is processed through the circuit, the signal is divided into three subbands. Further divisions can be obtained by cascading the circuit in a tree structure or in separable subbanding for two-dimensional data, after which compression coding may be applied to the subbands.

The reconstruction circuit (see Fig. 3) consists of registers and adders. The corresponding low, mid, and high band samples are latched into registers. The reconstruction circuit can also be cascaded in a tree structure to provide the reconstruction of a tree structure subbanded signal.

In Fig. 2, input data is delivered from an image, to a first latch 10, a second latch 12 and a third latch 14. A sign compliment 16 converts the data from the first latch to the compliment of band 0. An adder 18 adds the output of the second latch 12 with the output of the sign compliment to produce the band 1 subband. The band 2 subband is created by using an adder 20 to combine the output from the third latch 14 with the output of the sign compliment 16.

Fig. 3, is a schematic of the synthesis circuit for the subbanding apparatus depicted in Fig. 1. In Fig. 3, a conduit 21 recovers the band 0 signal and transmits it to produce a first reconstructed signal. A first adder 22 adds the band 0 subband and a band 1 subband to produce the second reconstructed signal. Lastly, a second adder 24 adds the band 0 signal to the band 2 signal to produce a third reconstructed output. The first, second and third reconstructed outputs then serve as the basis for reconstructing the subbanded image.

The process for performing the three-band split on a one-dimensional data set is as follows:

1. Three samples from the input data stream are read into the coder at one time.

2. The one sample becomes the band 0 (B_o) subbanded value.

3. The B_o sample is subtracted from another sample to provide the band 1 value.

4. The B_o sample is subtracted from the remaining sample to produce the band 2 value.

5. Three new samples are read into the coder.

After the signal is split into bands, compression coding can be applied.

The process for performing the recombination of the signal from the subbands is follows:

1. The corresponding values from the three bands are read into the decoder in parallel.

2. The value from band 0 is the one reconstructed sample.

3. The values from band 0 and band 1 are added together to give the second reconstructed sample.

4. The values from band 0, band 1, and band 2 are added together to give the third reconstructed sample.

5. The next three values are read into the decoder.

The most useful implementation using the basic algorithm is in a two-dimensional, filter bank as would be used for image processing.

This can be done by separable filtering which applies the above filters to the data twice, once horizontally and once vertically as shown in Fig. 4. The two dimensional filter would require that samples be accessed in blocks, in this case 3x3 blocks. Cascading can be accomplished with FIFOs or fast Random Access Memory (RAM) and appropriate addressing circuitry.

A more direct two-dimensional implementation of the invention works directly on a 3X3 block of data. In this implementation, nine data values d(i, J) (for i = 0, 1, 2, and j = 0, 1, 2) are processed with the following analysis equations to get the nine bands:

1..$B_0 = d(0,0)$
2..$B_1 = d(0,1) - d(0,0)$
3..$B_2 = d(0,2) - d(0,0)$
4..$B_3 = d(1,0) - d(0,0)$
5..$B_4 = d(1,1) - d(1,0) - d(0,1) + d(0,0)$
6..$B_5 = d(1,2) - d(1,0) - d(0,2) + d(0,0)$
7..$B_6 = d(2,0) - d(0,0)$
8..$B_7 = d(2,1) - d(2,0) - d(0,1) + d(0,0)$
9..$B_8 = d(2,2) - d(2,0) - d(0,2) + d(0,0)$

This set of equations performs a two-dimensional block transform equivalent to separable subband coding. Equations 1 through 9 can be implemented in software or in hardware in the same manner as the 3X3 case presented, above. The two-dimensional transform is also equivalent to multiplying the data block by the basis pictures of the transform matrices.

The reconstruction of the signal from the subbands uses the following synthesis equations:

10..$d(0,0) = B_0$
11..$d(0,1) = B_0 + B_1$
12..$d(0,2) = B_0 + B_2$
13..$d(1,0) = B_0 + B_3$
14..$d(1,1) = B_0 + B_1 + B_3 + B_4$
15..$d(1,2) = B_0 + B_2 + B_3 + B_5$
16..$d(2,0) = B_0 + B_6$
17..$d(2,1) = B_7 + B_6 + B_1 + B_0$
18..$d(2,2) = B_8 + B_6 + B_2 + Bo$

Fig. 7 displays a video screen 280. An enlarged view of the upper nine pixels is shown by the 3x3 matrix in the upper left hand corner of the video screen. The enlarged view of section 276 is the 3x3 matrix of pixels displayed in Fig. 8. In Fig. 8 the digital domain representation of d(0,0), d(0,1), d(0,2), d(1,0), d(1,1), d(1,2), d(2,0), d(2,1), d(2,2) are represented by 300, 302, 304, 306, 308, 310, 312, 314, and 316, respectively. The digital domain representation is first transformed to the frequency domain representation where the process described in equations 1 through 9 given above, is performed. The subbands are then created as follows. The first 3x3 matrix of pixels is taken from the upper left hand corner of the screen. Shifting over to the right, by three pixels, the next 3x3 pixels are used. Once the first row of three pixels in one matrix is used, the second row of three pixels in the 3x3 matrix of pixels is used. This procedure continues until all the pixels comprising the screens have been separated into subbands. For example, in Fig. 8 band zero would be created by combining all of the pixels located in position 300. Band 1 would be created by taking all of the pixels in position 302 and subtracting them from the pixels in position 300. Band 2 through 8 would be derived using the method given above in equations 3 through 9 above.

Alternate Embodiments of the Invention:

Alternate embodiments of the invention include organizing the basic circuit in parallel processing configurations, using different tree structures to break down particular subbands into finer subbands, implementing buffers with other memory devices than RAM (e.g., CCD or FIFO), implementing the circuit on very large scale integration devices or semi-custom devices or gate arrays or application specific integrated circuits or programmable array logic or programmable logic devices or other methods of integrated circuit implementation. The use of various integrated circuit technologies, such as TTL, ECL, GaAs, CMOS, NMOS, etc., can be used to implement the invention. Discrete components on circuit boards or wire wrap can also be used to implement the invention.

Alternate embodiments could work on data sample sizes other than nine bits. For example, if the input data sample size was four or 12 bits, the circuit components could be changed to accommodate this size data word. The implementation of this invention shown is eight bits since this is a common size data sample and would likely be the size of luminance and chromance samples in a digital video signal.

An alternate embodiment of the invention would accept serial data by using a serial-to-parallel register. A simple compression coder can be obtained by coarsely quantizing values in the higher bands. This will somewhat degrade the reconstructed signal, but will provide data compression without loss of significant information. When decoding, the results of the quanitzed bands must be scaled back to the proper range, of course.

By dropping the high frequency band in a four band split of an image the data rate is reduced to 3/4 of the original because 1/4 of the information has been deleted. In a nine band split, dropping the highest band results in a 1/9 reduction in information giving a rate that is 8/9 of the original. With slow roll-off filters, simply dropping the highest band will result in some aliasing in the reconstructed image.

To minimize aliasing, the high band can be coarsely quantized, even with one bit. In this case, the high band pixels would be represented by a zero when the pixel value is below a threshold and by a one when it is above. Run-length coding can then be used to further compress the data. Any number of quantization levels between one (dropping the band) and 256 (no compression) could be used on any or all of the bands. Since most of the image information is in the lowest band, it makes sense to use finer quantization of that band.

As an example of compression due to quantizing subbands, consider the following nine band split of a 1920 X 1035 image with eight bits/pixel using coarse quantization:

Frequencies	Band	Original Size (bits)	Quantization Levels	Compressed Size
LL	0	1766400	256	1766400
LM	1	1766400	64	1324800
LH	2	1766400	32	1104000
ML	3	1766400	16	883200
MM	4	1766400	8	662400
MH	5	1766400	2	220800
HL	6	1766400	2	220800
HM	7	1766400	2	220800
HH	8	1766400	0	0
	Total	15897600		6403200

This example reduces the image data to approximately 40% of the original before run-length coding or Huffman coding. The next step is to apply run length coding and Huffman coding for additional compression. For coarse quantization, the best type of quantization is dead zone quantization. This is because deleting values with low energy will not result in noticeable aliasing until some threshold value is reached.

Color information can be handled by subbanding the chrominance signals corresponding to the luminance signal. Typically, the chrominance is subsampled by a factor of two in compression schemes. Subbanding and coarsely quantizing the high bands accomplishes a similar result. Two common chrominance/luminance formats that can be used are YIQ and YUV, but any format can be used.

An extension of the signal may be used around the border. This is typically done with higher order filters to minimized distortion. The extension may be a zero or a duplicate of the first or last value in a row or column. If the extension can be avoided if the number of samples processed at one time is an integer multiple of the block size, another at one time is an integer/approach is to drop any (one or two) extra samples in a row or column to get a length that is an integer multiple of the block length.

Data can be subbanded in several ways depending on how it is organized. For image data it is natural to process data in scan lines, fields, frames, etc. The most useful way is to process the images one frame at a time. Scan lines are concatenated with the end of one line followed by the beginning of another. If each line length is an integer multiple of the longest filter size, then no overlap should occur with respect to processing the end of one line and the beginning of the next. If the lowest band of an image is to be used for a low resolution (subsampled) version, it can be easily cropped by at least the length of the longest filter on all sides to remove any artifacts generated by line lengths that are not an integer multiple at the longest filter length. The number of circuits implementing the basic encoding or decoding circuit may be increased for parallel processing. Since the blocks are processed independently, massively parallel processing is possible.

Alternate embodiments can use filters other than the set shown above.

An alternate embodiment of the invention uses the non-orthogonal transform to produce four subbands, for example the following equations produce four subbands from 2x2 blocks of data:

$$B_0 = d_3$$
$$B_1 = d_1 - d_3$$

$$B_2 = d_2 - d_3$$
$$B_3 = d_o - d_1 - d_2 + d_3$$

Reconstruction:

$$\underline{d}_o = B_o + B_1 + B_2 + B_3$$
$$\underline{d}_1 = B_o + B_1$$
$$\underline{d}_2 = B_o + B_2$$
$$\underline{d}_3 = B_o$$

See Fig. 5 and 6 for a hardware implementation. The transform matrices are:

$$G = \begin{bmatrix} 0 & 1 \\ 1 & -1 \end{bmatrix} ; g = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$$

All of these transforms use filters with coefficients that are implementable by addition and subtraction and the filters have the same number of taps as the decimation factor. For example, a four band circuit will have equations with four or fewer terms defining the subbands; a nine band circuit will have equations with nine or fewer terms.

Cascading the algorithm produces results equivalent to using a larger matrix formed by the Kronecker product of the cascaded matrix. For example, cascading the 2x2 transform algorithm twice is equivalent to using the 4x4 transform matrix (formed by the Kronecker square) once. In cascading the algorithm, the algorithm is applied to each of the subbands produced by the first stage separately for a uniform tree structure or the algorithm can be applied only to the low band for an octave tree structure.

Fig 5. displays a two-dimensional, four-channel subband analysis bank. In Fig. 5 an input data stream 30 feeds the pixels from a video screen, into a store and convert unit 32. The store and convert unit 32 produces a first, a second, a third and a fourth output stream which travels along connections 34, 36, 38 and 40, respectively. Each of the output data streams are fed into a sample register where the data is held until the next timing change. The first, second, third and fourth sample registers are denoted by 44, 46, 48, and 50, respectively. The output 60 of the fourth sample register is summed to produce a band 0 subband. The output from the second sample register 56 is subtracted from band 0 denoted by 60 at a first summing means 62. As a result, a band 1 subband denoted by 68 is created. A second summing means 64 subtracts the band 0 output 60 from the output of the third sample register 58. As a result, a band 2 subband denoted by 70 is produced. Lastly, a third summing means 66 subtracts the band 2 subband, from the output from the first sample register 54. As a result, a band 3 subband is created. Therefore the two-dimensional four-channel subband analysis bank data displayed in Fig. 5, takes an input data stream and turns it into four subbands using the non-orthogonal transform.

Fig. 6 displays a schematic diagram of a circuit that reconstruct the various subband. In the synthesis bank band 0, band 1, band 2 and band 3 subbands are combined by summing means to produce d0, d1, d2 and d3 pixel outputs. In Fig. 6 band 0 denoted by 86 is converted directly into the d3 output. Band 2 denoted by 84 is combined with band 0 at a third summing means 92. As a result the output pixel d2 is produced at 98. Band 0 denoted by 86 is combined with band 1 denoted by 82, at summing means 90. As a result, pixels d1 is produced at 96. Lastly, pixel d0 is produced by combining band 1, band 2 and band 3 at a first summing means 88 to produce output 94.

Although specific embodiments of the method is detailed above, there is also a general approach to designing non-orthogonal transform matrices.

To design a non-orthogonal transform pair of matrices, it is easiest to start with diagonal matrices. One of the transform pairs is upper diagonal and the other is lower diagonal. For a 4x4 matrix, for example, this would give the following basic form to the matrices:

$$G(4) = \begin{bmatrix} 1 & a & b & d \\ 0 & 1 & c & e \\ 0 & 0 & 1 & f \\ 0 & 0 & 0 & 1 \end{bmatrix} \qquad g(4) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ g & 1 & 0 & 0 \\ h & j & 1 & 0 \\ i & k & l & 1 \end{bmatrix}$$

To get the identity matrix as the result of the matrix multiplication required from the equations above, a vector-by-vector dot product can be used between the basis vectors in the two matrices. The result of the dot product between each possible pair of vectors must be the corresponding element in the identity matrix. That is, for basis vectors $v_i \cdot u_j = I$ for $i = j$ and $v_i \cdot u_j = 0$ for $i \neq j$. In the 4x4 matrix case, this gives six equations with 12 unknowns allowing the designer to choose the forward transform to provide the processing of the data that is desired. For the matrices above the equations would be:

$$a + g = 0$$

$$c + j = 0$$

$$f + l = 0$$

$$b + cg + h = 0$$

$$e + fj + k = 0$$

$$d + eg + fh + i = 0$$

The designer can now choose a, b, c, d, e, and f to provide the desired filtering of the data. Typically some of the variables will be set to negative numbers to provide a difference function that will tend to reduce the dynamic range of the processed data.

As an example, select $a = c = d = 0$ and $b = e = f = -1$. Plugging into the equations above gives $g = j = 0$ and $h = i = k = l = 1$. The resulting matrices make a non-orthogonal transform pair.

$$G(4) = \begin{bmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 1 \end{bmatrix} \qquad g(4) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}$$

Notice that the 2x2 matrix case is a subset of the 4x4 case. The 2x2 matrices

$$G(2) = \begin{matrix} 1 & a \\ 0 & 1 \end{matrix} \qquad g(2) = \begin{matrix} 1 & 0 \\ g & 1 \end{matrix}$$

give the equation $a + g = 0$. The possible results are a=0, g=0; a=1, g= -1; and a=-1, g=1. The first case provides the trivial filter set, the last two cases are duals and provide the 2x2 non-orthogonal transform.

The non-orthogonal transform design can be done on any size square matrix: 2x2, 3x3, 4x4, 5x5, etc.

Alternate embodiments eliminate the circuitry for computing the higher frequency bands if those bands are just going to be deleted anyway, as in lossy compression coding.

While several embodiments of the method and apparatus of non-orthogonal subband/transform coding have been disclosed, it will be appreciated that various modifications may be made without departing from the spirit of the invention and the scope of the subjoined claims.

ABSTRACT OF THE DISCLOSURE

The present invention is directed to a simplified digital subband coder/decoder. In the present invention a signal is fed into a coder. The coder uses a non-orthogonal algorithm that is simply implemented in the coder hardware. The simple non-orthogonal design is then used in the implementation of the decoder to decode the signal.

Sample Index (n) Example Data x(n)	0 0	1 1	2 2	3 3	4 4	5 240	6 240	7 200	8 4	9 3	10 2	11 1	12 0	13 1	14 2
Subbanding															
Subband Result Index (m)			0			1			2			3			4
Band 0: B0(m)=x(3m)			0			3			240			3			0
Band 1: B1(m)=x(3m+1)-x(3m)			1			1			-40			-1			1
Band 2: B2(m)=x(3m+2)-x(3m) 2			2			237			-236			-2			
Reconstruction															
First Reconstructed Sample $\hat{x}(3m) = B0(m)$	0			3			240			3			0		
Second Reconstructed Sample $\hat{x}(3m+1) = B0(m)+B1(m)$		1			4			200			2			1	
Third Reconstructed Sample $\hat{x}(3m+2) = B0(m)+B2(m)$			2				240		4			1			2
Reconstructed Signal $\hat{x}(n)$	0	1	2	3	4	240	240	200	4	3	2	1	0	1	2

Fig.1

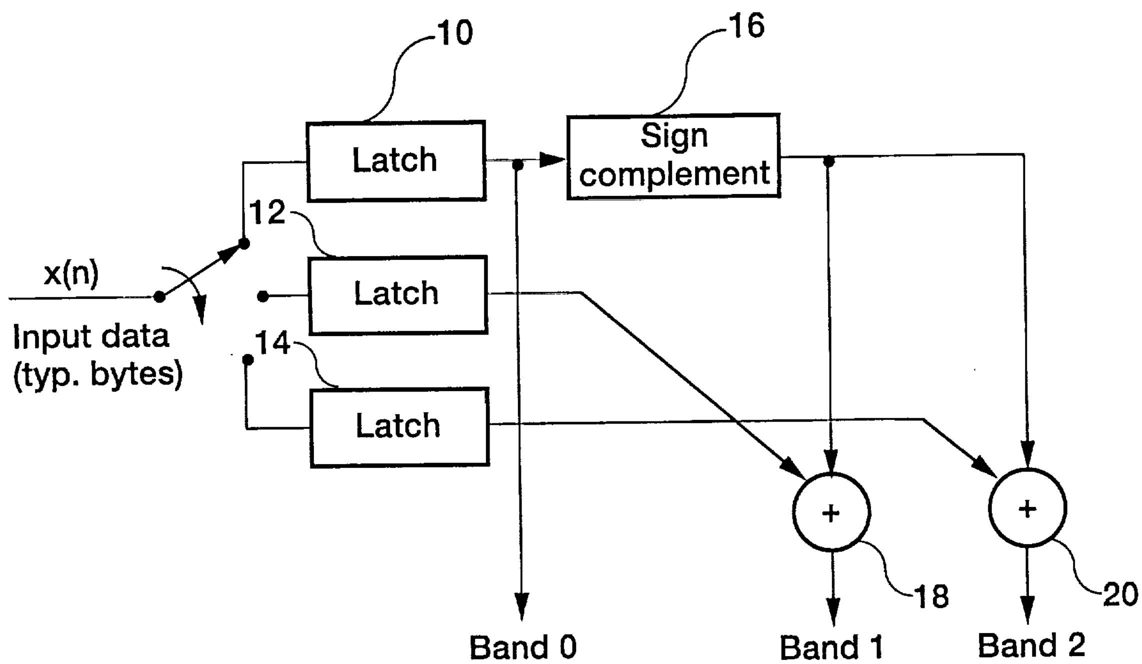
10
16
Latch
Sign
complement
12
x(n)
Latch
Input data
(typ. bytes)
14
Latch
+
+
18
20
Band 0
Band 1
Band 2

Fig. 2

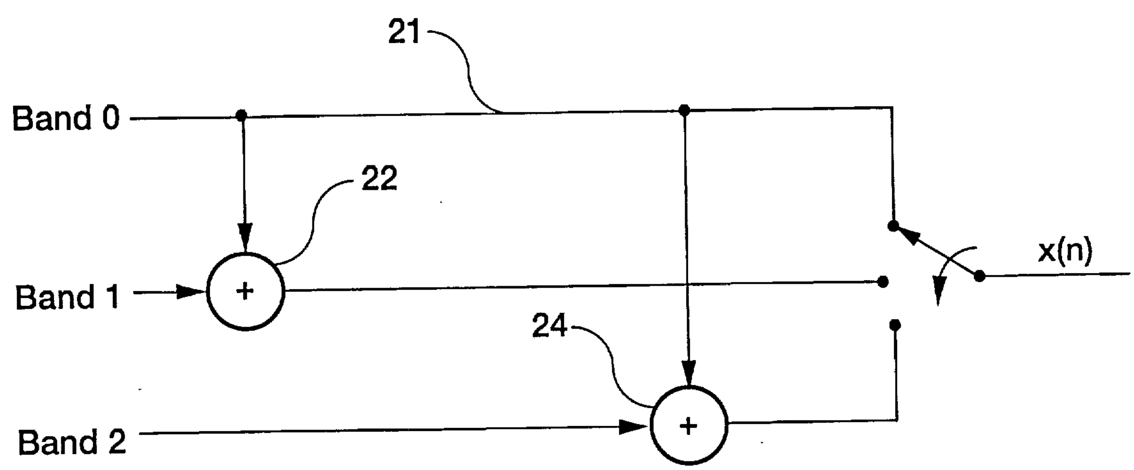

Fig. 3

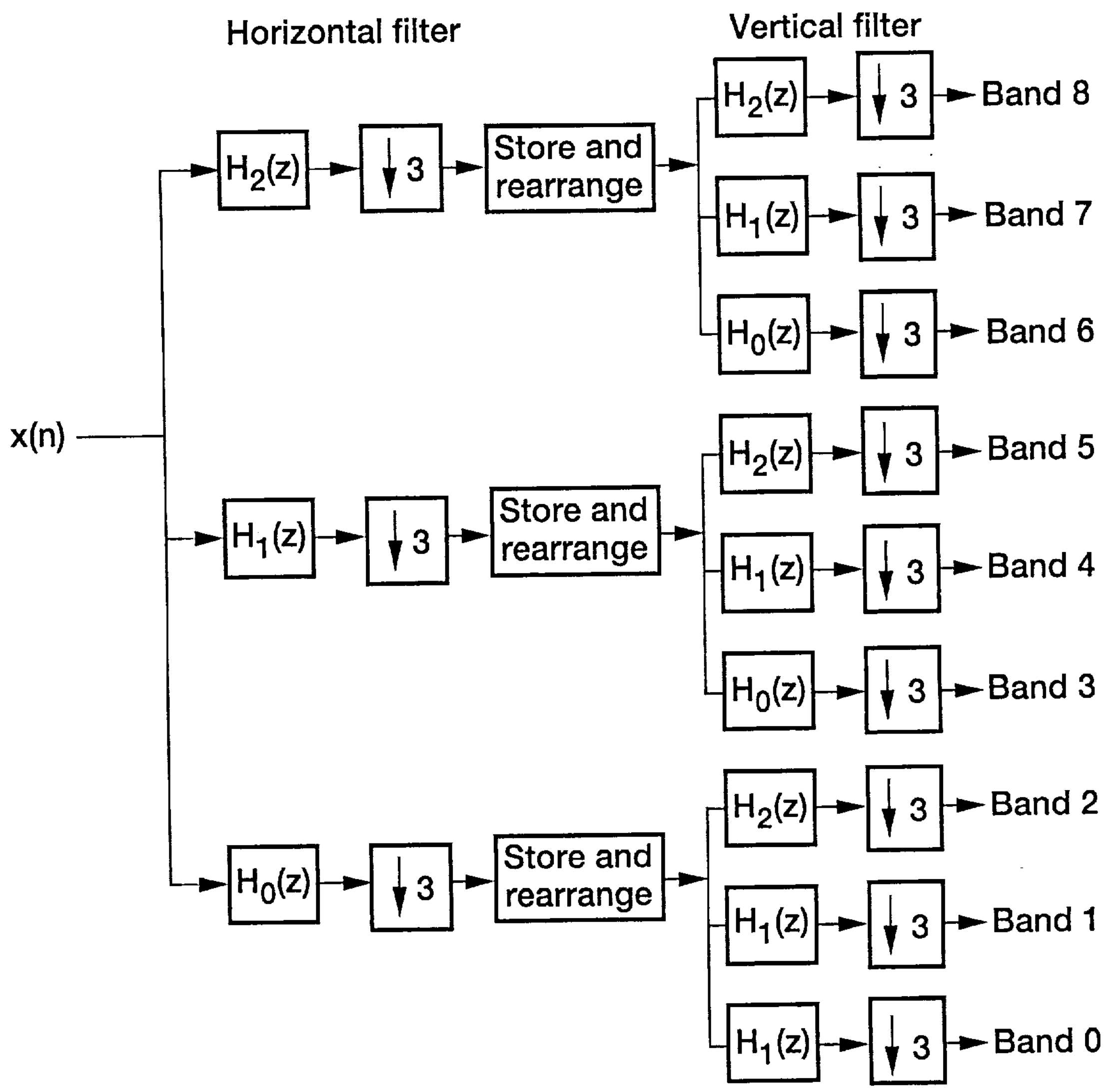
Horizontal filter
Vertical filter
x(n)
$H_2(z)$
$\downarrow 3$
Store and rearrange
$H_1(z)$
$H_0(z)$
Band 8
Band 7
Band 6
Band 5
Band 4
Band 3
Band 2
Band 1
Band 0

Figure 4.

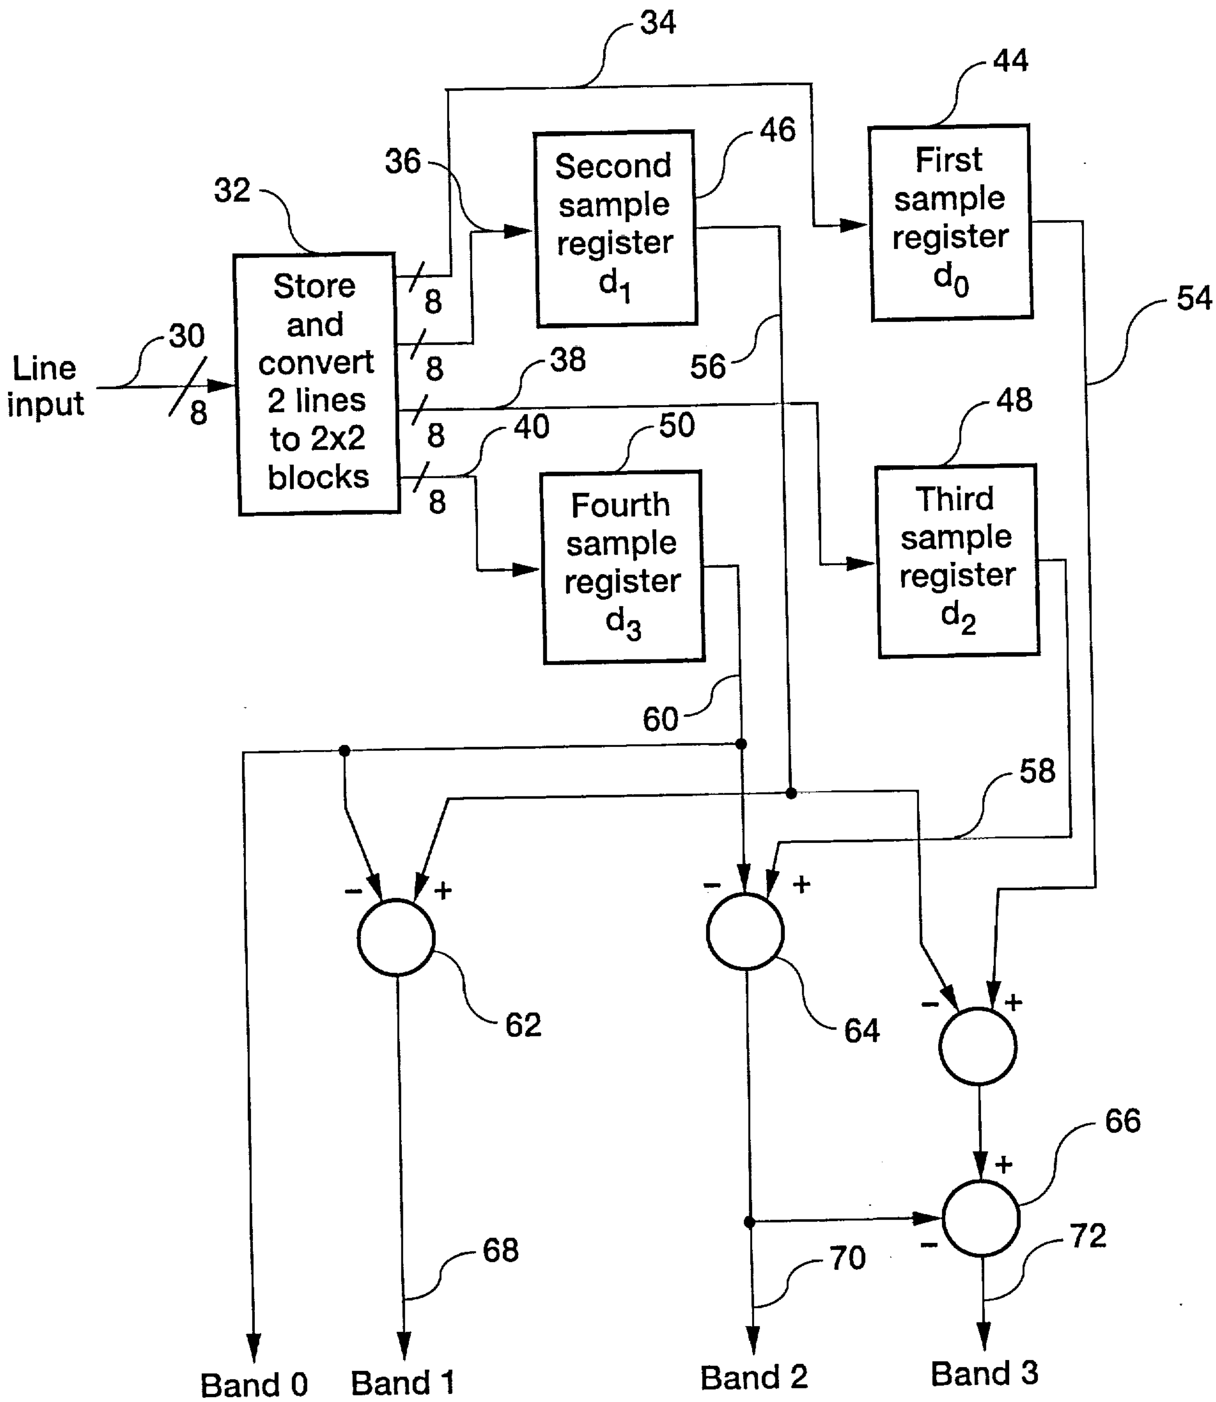
34
44
36
46
Second sample register d_1
First sample register d_0
32
Store and convert 2 lines to 2x2 blocks
8
8
8
8
8
30
Line input
38
56
54
40
50
48
Fourth sample register d_3
Third sample register d_2
60
58
62
64
66
68
70
72
Band 0
Band 1
Band 2
Band 3

Fig. 5

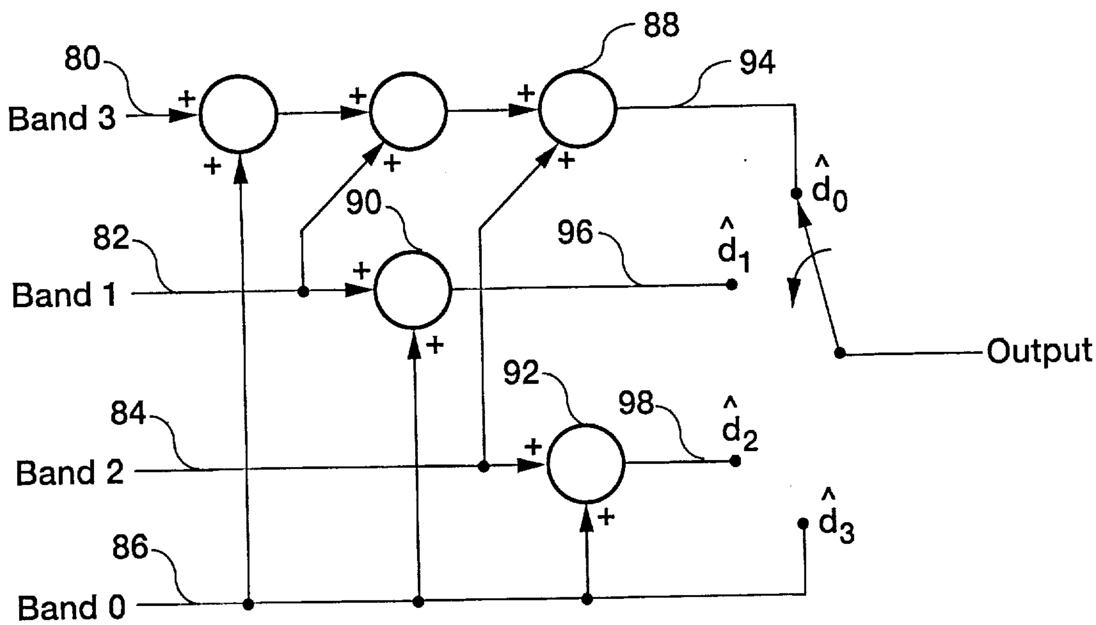
80
Band 3
82
Band 1
84
Band 2
86
Band 0
88
90
92
94
96
98
$\hat{d}_0$
$\hat{d}_1$
$\hat{d}_2$
$\hat{d}_3$
Output

Figure 6

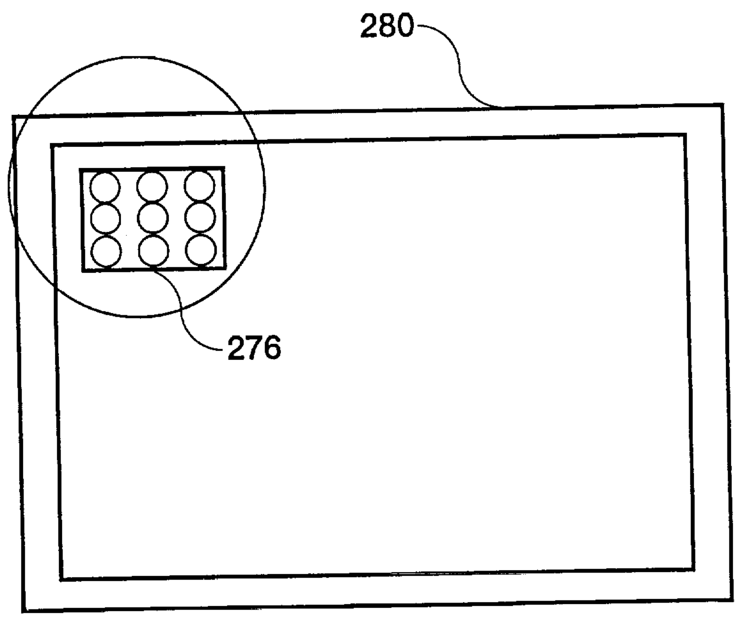

Fig.7

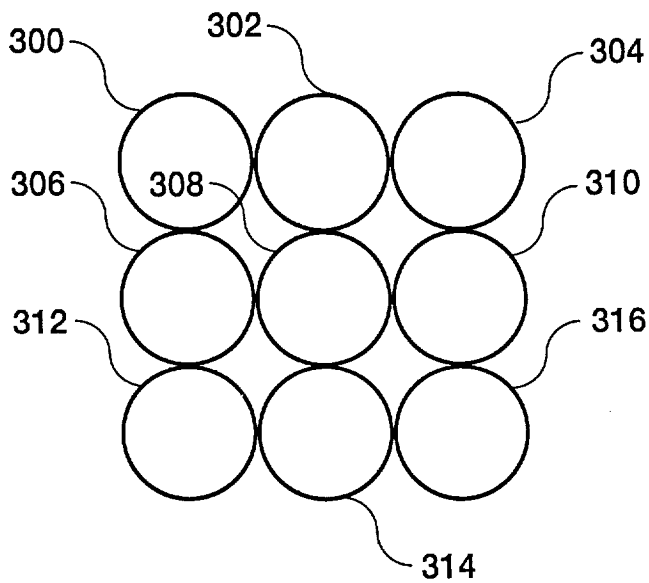

Fig.8