

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 August 2011 (25.08.2011)

(10) International Publication Number
WO 2011/101442 A2

(51) International Patent Classification:
H04N 7/32 (2006.01)

SEN, Soren, Vang [DK/LU]; 38 Rue Wilson, L-2732 Luxembourg (LU).

(21) International Application Number:
PCT/EP2011/052446

(74) Agents: DRIVER, Virginia, Rozanne et al.; Page White & Farrer, Bedford House, John Street, London Greater London WC1N 2BF (GB).

(22) International Filing Date:
18 February 2011 (18.02.2011)

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/306,399 19 February 2010 (19.02.2010) US
12/838,197 16 July 2010 (16.07.2010) US

(71) Applicant (for all designated States except US): SKYPE LIMITED [IE/IE]; Arthur Cox Building, Earlsfort Centre, Earlsfort Terrace, Dublin, 2 (IE).

(72) Inventors; and

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU,

(75) Inventors/Applicants (for US only): BIVOLARSKY, Lazar [BG/US]; 11024 Holly Tree Lane, Cupertino, CA 95014 (US). NILSSON, Mattias [SE/SE]; Egnahemsvagen 5, S-17446 Sundbyberg (SE). VAFIN, Renat [EE/EE]; Raagu 33-7, EE-13417 Tallinn (EE). ANDER-

[Continued on next page]

(54) Title: DATA COMPRESSION FOR VIDEO

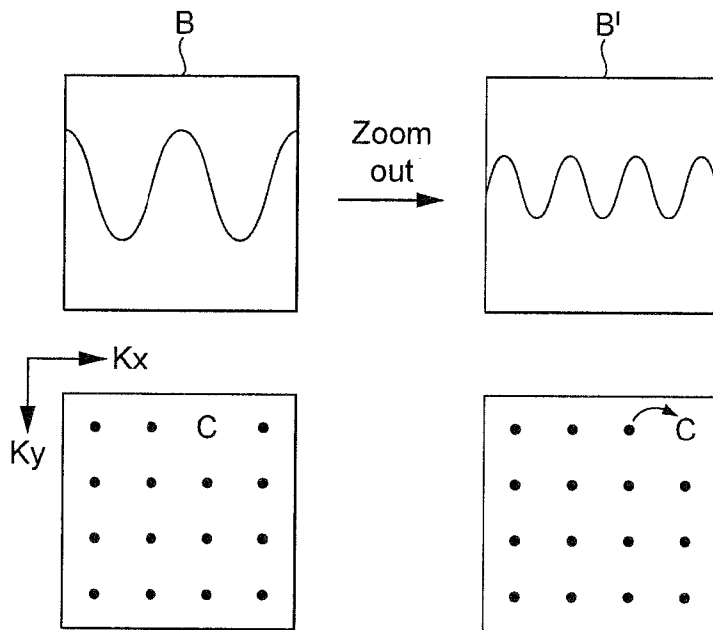


FIG. 13a

(57) Abstract: A method of compressing a video signal. The method comprises: for each of a plurality of image portions, transforming from a spatial domain representation into a transform domain representation; and for each of a plurality of target image portions to be encoded, determining a respective reference portion, determining a prediction of the target image portion based on the reference portion, and encoding the target image portion based on the prediction. The prediction comprises predicting a non-lateral motion of the target image portion relative to the reference portion based on processing of their transform domain coefficients. The encoding comprises generating an encoded bitstream comprising information about the non-lateral motion for use by a decoder.

WO 2011/101442 A2

LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, **Published:**
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, — *without international search report and to be republished*
GW, ML, MR, NE, SN, TD, TG). *upon receipt of that report (Rule 48.2(g))*

DATA COMPRESSION FOR VIDEO

FIELD OF THE INVENTION

5 The present invention relates to the encoding and transmission of video streams.

BACKGROUND

10 In the transmission of video streams, efforts are continually being made to reduce the amount of data that needs to be transmitted whilst still allowing the moving images to be adequately recreated at the receiving end of the transmission. A video encoder receives an input video stream comprising a sequence of "raw" video frames to be encoded, each representing an image at a respective moment in time. The encoder then encodes each input frame into one of two types of
15 encoded frame: either an intra frame (also known as a key frame), or an inter frame. The purpose of the encoding is to compress the video data so as to incur fewer bits when transmitted over a transmission medium or stored on a storage medium.

20 An intra frame is compressed using data only from the current video frame being encoded, typically using intra frame prediction coding whereby one image portion within the frame is encoded and signalled relative to another image portion within that same frame. This is similar to static image coding. An inter frame on the other hand is compressed using knowledge of a preceding frame (a reference frame)
25 and allows for transmission of only the differences between that reference frame and the current frame which follows it in time. This allows for much more efficient compression, particularly when the scene has relatively few changes. Inter frame prediction typically uses motion estimation to encode and signal the video in terms of motion vectors describing the movement of image portions between frames,
30 and then motion compensation to predict that motion at the receiver based on the signalled vectors. Various international standards for video communications such as MPEG 1, 2 & 4, and H.261, H.263 & H.264 employ motion estimation and compensation based on regular block based partitions of source frames.

Depending on the resolution, frame rate, bit rate and scene, an intra frame can be up to 20 to 100 times larger than an inter frame. On the other hand, an inter frame imposes a dependency relation to previous inter frames up to the most recent intra frame. If any of the frames are missing, decoding the current inter frame may result in errors and artefacts.

These techniques are used for example in the H.264/AVC standard (see T. Wiegand, G.J. Sullivan, G. Bjontegaard, A. Luthra: "Overview of the H.264/AVC video coding standard," in IEEE Transactions on Circuits and Systems for Video Technology, Volume: 13, Issue: 7, page(s): 560-576, July 2003).

Figure 7 illustrates a known video encoder for encoding a video stream into a stream of inter frames and interleaved intra frames, e.g. in accordance with the basic coding structure of H.264/AVC. The encoder receives an input video stream comprising a sequence of frames to be encoded (each divided into constituent macroblocks and subdivided into blocks), and outputs quantized transform coefficients and motion data which can then be transmitted to the decoder. The encoder comprises an input 70 for receiving an input macroblock of a video image, a subtraction stage 72, a forward transform stage 74, a forward quantization stage 76, an inverse quantization stage 78, an inverse transform stage 80, an intra frame prediction coding stage 82, a motion estimation & compensation stage 84, and an entropy encoder 86.

The subtraction stage 72 is arranged to receive the input signal comprising a series of input macroblocks, each corresponding to a portion of a frame. From each, the subtraction stage 72 subtracts a prediction of that macroblock so as to generate a residual signal (also sometimes referred to as the prediction error). In the case of intra prediction, the prediction of the block is supplied from the intra prediction stage 82 based on one or more neighbouring regions of the same frame (after feedback via the reverse quantization stage 78 and reverse transform stage 80). In the case of inter prediction, the prediction of the block is provided from the motion estimation & compensation stage 84 based on a selected region of a preceding frame (again after feedback via the reverse quantization stage 78

and reverse transform stage 80). For motion estimation the selected region is identified by means of a motion vector describing the offset between the position of the selected region in the preceding frame and the macroblock being encoded in the current frame.

5

The forward transform stage 74 then transforms the residuals of the blocks from a spatial domain representation into a transform domain representation, e.g. by means of a discrete cosine transform (DCT). That is to say, it transforms each residual block from a set of pixel values at different Cartesian x and y coordinates to a set of coefficients representing different spatial frequency terms with different wavenumbers k_x and k_y (having dimensions of 1/wavelength). The forward quantization stage 76 then quantizes the transform coefficients, and outputs quantised and transformed coefficients of the residual signal to be encoded into the video stream via the entropy encoder 86, to thus form part of the encoded video signal for transmission to one or more recipient terminals.

10
15

Furthermore, the output of the forward quantization stage 76 is also fed back via the inverse quantization stage 78 and inverse transform stage 80. The inverse transform stage 80 transforms the residual coefficients from the frequency domain back into spatial domain values where they are supplied to the intra prediction stage 82 (for intra frames) or the motion estimation & compensation stage 84 (for inter frames). These stages use the reverse transformed and reverse quantized residual signal along with knowledge of the input video stream in order to produce local predictions of the intra and inter frames (including the distorting effect of having been forward and reverse transformed and quantized as would be seen at the decoder). This local prediction is fed back to the subtraction stage 72 which produces the residual signal representing the difference between the input signal and the output of either the local intra frame prediction stage 82 or the local motion estimation & compensation stage 84. After transformation, the forward quantization stage 76 quantizes this residual signal, thus generating the quantized, transformed residual coefficients for output to the entropy encoder 86. The motion estimation stage 84 also outputs the motion vectors via the entropy encoder 86 for inclusion in the encoded bitstream.

20

25

30

When performing intra frame encoding, the idea is to only encode and transmit a measure of how a portion of image data within a frame differs from another portion within that same frame. That portion can then be predicted at the decoder (given some absolute data to begin with), and so it is only necessary to transmit the difference between the prediction and the actual data rather than the actual data itself. The difference signal is typically smaller in magnitude, so takes fewer bits to encode.

10 In the case of inter frame encoding, the motion compensation stage 84 is switched into the feedback path in place of the intra frame prediction stage 82, and a feedback loop is thus created between blocks of one frame and another in order to encode the inter frame relative to those of a preceding frame. This typically takes even fewer bits to encode than an intra frame.

15

Figure 8 illustrates a corresponding decoder which comprises an entropy decoder 90 for receiving the encoded video stream into a recipient terminal, an inverse quantization stage 92, an inverse transform stage 94, an intra prediction stage 96 and a motion compensation stage 98. The outputs of the intra prediction stage and the motion compensation stage are summed at a summing stage 100.

20

There are many known motion estimation techniques. Generally they rely on comparison of a block with one or more other image portions from a preceding frame (the reference frame). Each block is predicted from an area of the same size and shape as the block, but offset by any number of pixels in the horizontal or vertical direction or even a fractional number of pixels. The identity of the area used is signalled as overhead ("side information") in the form of a motion vector. A good motion estimation technique has to balance the requirements of low complexity with high quality video images. It is also desirable that it does not require too much overhead information.

30

In the standard system described above, it will be noted that the intra prediction coding and inter prediction coding (motion estimation) are performed in the unquantized spatial domain.

5 More recently, motion estimation techniques operating in the transform domain have attracted attention. However, none of the existing techniques are able to perform with low complexity (thus reducing computational overhead), while also delivering high quality. Hence no frequency domain techniques for motion estimation are currently in practical use.

10

The VC-1 video codec has an intra prediction mode which operates in the frequency domain, in which the first column and/or first row of AC coefficients in the DCT (Discrete Fourier Transform) domain are predicted from the first column (or first row) of the DCT blocks located immediately to the left or above the
15 processed block. That is to say, coefficients lying at the edge of one block are predicted from the direct spatial neighbours in an adjacent block. For reference see "The VC-1 and H.264 Video Compression Standards for Broadband Video Services", AvHari Kalva, Jae-Beom Lee, pp.251.

20

SUMMARY OF INVENTION

According to one aspect of the present invention, there is provided a method of
compressing a video signal comprising a plurality of video frames, each frame
25 being divided into a plurality of image portions, comprising: for each of a plurality of the image portions, transforming the image portion from a spatial domain representation into a transform domain representation so as to generate respective set of transform domain coefficients; and for each of a plurality of target ones of said image portions to be encoded, determining a respective reference
30 portion represented by a respective set of transform domain coefficients, determining a prediction of the target image portion based on the reference portion, and encoding the target image portion based on the prediction; wherein the prediction comprises predicting a non-lateral motion of the target image

portion relative to the reference portion based on processing of their transform domain coefficients; and wherein the encoding comprises generating an encoded bitstream comprising information about the non-lateral motion for use by a decoder.

5

In embodiments, the prediction of said non-lateral motion may comprise predicting a scaling of the target image portion relative to the reference portion based on processing of their transform domain coefficients. Said information may indicate a parameter of the scaling.

10

The prediction of said non-lateral motion may comprise predicting a rotation of the target image portion relative to the reference portion based on processing of their transform domain coefficients. Said information may indicate a degree of the rotation.

15

The determination of the reference portion may comprise selecting a selected reference portion from amongst a plurality of potential reference portions each being represented by a respective set of transform domain coefficients, and the prediction of the target image portion may be based on the selected reference portion.

20

The plurality of potential reference portions may comprises at least one potential reference portion for predicting a lateral motion of the target image portion according to one type of motion prediction, and at least one potential reference portion for predicting a non-lateral motion of the target image portion according to another type of motion prediction; and said information may indicate the type of motion prediction to be used by a decoder for decoding the target image portion.

25

The plurality of potential reference portions may comprise at least one potential reference portion for predicting a lateral motion of the target image portion according to a first type of motion prediction, at least one potential reference portion for predicting a scaling of the target image portion according to a second type of motion prediction, and at least one potential reference portion for

30

predicting a rotation of the target image portion according to a third type of motion prediction; and said information may indicate whether the first, second or third type of motion prediction is to be used by a decoder for decoding the target image portion.

5

The method may comprise determining a respective index value for the target image portion and each of a plurality of the reference portions based on the respective transform domain coefficients; wherein the selection of the selected reference portion may be based on the index values.

10

The selection of the selected reference portion may comprise selecting a subset of candidate portions having an index value within a predetermined range of the target image portion, and selecting the respective reference portion from amongst the candidate portions.

15

The index value may represent an energy of the respective portion. The index value may be based on one of: a number of non-zero transform domain coefficients in the respective portion, and a number of transform domain coefficients having a value of zero in the respective portion.

20

The index value may be based on an aggregate combination of the transform domain coefficients of the respective portion. The index value may be based on an average or total of a modulus of the transform domain coefficients of the respective portion.

25

The selection of the selected reference portion may be based on a determination of a number of bits that would be required in the encoded bitstream to encode the target image portion.

30

The selection of the selected reference portion from amongst the candidate portions may be based on a determination of a number of bits that would be required in the encoded bitstream to encode the target image portion relative to each of the candidate portions.

The method may comprise: during ongoing encoding of the video signal, updating a table indicating a subset of commonly usable reference portions; and including the table in the encoded bitstream along with side information identifying the selected reference portions by reference to an entry in said table.

One or more of said subset of reference portions may be artificial portions rather than actual portions from a frame.

10 The method may comprise determining the one or more artificial reference portions by averaging, interpolating or extrapolating from one or more actual image portions from one or more of said frames.

One or more of said subset of reference portions may be actual image portions from one or more of said frames.

The transform domain coefficients may represent chrominance and luminance in the transform domain.

20 Each image portion may comprise a block or macroblock corresponding to a two-dimensional array of pixels.

According to another aspect of the present invention, there may be provided an encoder for compressing video data, the encoder comprising: an input for receiving a video signal comprising a plurality of video frames, each frame being divided into a plurality of image portions; and processing apparatus configured, for each of a plurality of the image portions, to transform the image portion from a spatial domain representation into a transform domain representation so as to generate respective set of transform domain coefficients; and for each of a plurality of target ones of said image portions to be encoded, to determine a respective reference portion represented by a respective set of transform domain coefficients, determine a prediction of the target image portion based on the reference portion, and encode the target image portion based on the prediction;

wherein the processing apparatus is operable to perform said prediction by predicting a non-lateral motion of the target image portion relative to the reference portion based on processing of their transform domain coefficients; such that the encoding comprises generating an encoded bitstream comprising information about the non-lateral motion for use by a decoder; and wherein the encoder further comprises an output module arranged to output the encoded bitstream.

In embodiments, the processing apparatus may be further configured to generate an encoded bitstream in accordance with any of the above method features.

10

According to another aspect of the present invention, there is provided an encoder program product for compressing video data, the encoder program product comprising software embodied on a computer-readable medium and configured so as when executed to: for each of a plurality of image portions from said stream, transform the image portion from a spatial domain representation into a transform domain representation so as to generate respective set of transform domain coefficients; and for each of a plurality of target ones of said image portions to be encoded, determine a respective reference portion represented by a respective set of transform domain coefficients, determine a prediction of the target image portion based on the reference portion, and encode the target image portion based on the prediction; wherein the prediction comprises predicting a non-lateral motion of the target image portion relative to the reference portion based on processing of their transform domain coefficients; and wherein the encoding comprises generating an encoded bitstream comprising information about the non-lateral motion for use by a decoder.

In embodiments, the software may be further configured to generate an encoded bitstream in accordance with any of the above features.

30 According to another aspect of the present invention, there may be provided an encoded bit stream of compressed video data, representing a plurality of successive video frames with each frame being divided into a plurality of image portions and each portion corresponding to a plurality of pixels, the encoded bit

stream comprising: a set of quantized transform domain coefficients representing a first one of said portions; and information relating to a non-lateral motion of a second one of said portions which is to be predicted from the first portion, allowing the second portion to be predicted based on the non-lateral motion and on the transform domain coefficients of the first portion.

In embodiments the encoded bitstream may be encoded according to any of the above features.

10 According to another aspect of the present invention, there is provided a network equipment comprising a transmission medium carrying the bitstream.

According to another aspect of the present invention, there is provided a method of decoding the encoded bitstream, comprising receiving the encoded bitstream, and predicting the second portion based on the information relating to the non-lateral motion and on the transform domain coefficients of the first portion.

According to another aspect of the present invention, there is provided a decoder for decoding the encoded bitstream, comprising an input for receiving the encoded bitstream, and processing apparatus configured to predict the second portion based on the information relating to the non-lateral motion and on the transform domain coefficients of the first portion.

According to another aspect of the present invention, there is provided a decoder program product for decoding the encoded bitstream, the decoder program product comprising software embodied on a computer-readable medium and configured so as when executed to: receive the encoded bitstream, and predict the second portion based on the information relating to the non-lateral motion and on the transform domain coefficients of the first portion.

30

BRIEF DESCRIPTION OF THE DRAWINGS

For a better understanding of the present invention and to show how the same may be carried into effect, reference will now be made by way of example, to the accompanying drawings, in which:

- Figure 1 is a schematic illustration of two frames of a video stream,
Figure 1a shows the pixel values of blocks represented in the spatial domain,
Figure 1b shows coefficients of blocks represented in the frequency domain,
- 10 Figure 2 is a flow diagram showing an encoding method,
Figure 3 is a graph showing a sorted block list,
Figure 3a illustrates an example of block-sorting,
Figure 3b illustrates an example of block matching prediction
Figure 4 is a schematic block diagram of an encoder,
- 15 Figure 5A is a schematic example of an intra frame prediction,
Figure 5B is a schematic example of an inter frame prediction,
Figure 6 is a schematic diagram of a decoder,
Figure 7 is a schematic block diagram of an encoder,
Figure 8 is a schematic block diagram of a decoder,
- 20 Figure 9 is a schematic illustration of selection of candidates for block matching,
Figure 10 is a flow chart of an encoding method,
Figure 11 is a flow chart of a decoding method,
Figure 12 is a schematic representation of a transmitted bitstream,
Figure 13a is a schematic illustration of block scaling,
- 25 Figure 13b is a schematic illustration of block rotation,
Figure 13c is another schematic illustration of block rotation,
Figure 13d is another schematic illustration of block rotation,
Figure 13e is a diagram showing a calculation for use in block rotation,
Figure 14a schematically illustrates a motion shift between two frames,
- 30 Figure 14b is another schematic illustration of a motion shift, and
Figure 14c schematically shows using a motion shift to reduce data transmission.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

In the following there is described a method and system for data compression in a video transmission system. First is described an exemplary technique of block matching performed in the frequency domain for selecting a reference block to use in prediction coding of a target block. Next are described some exemplary techniques for signalling the identity of reference blocks for use in video prediction coding. Following that are described some exemplary image processing techniques that can be advantageously performed in the frequency domain, and a scheme for further reducing the bitrate of a transmitted video stream. In particularly preferred embodiments these techniques may be combined, but alternatively they may be used independently.

Block Matching

Figure 1 schematically illustrates two successive frames f_t and f_{t+1} of a video image at two respective moments in time t and $t+1$. For the purpose of inter frame prediction the first frame f_t may be considered a reference frame, i.e. a frame which has just been encoded from a moving sequence at the encoder, or a frame which has just been decoded at the decoder. The second frame f_{t+1} may be considered a target frame, i.e. the current frame whose motion is sought to be estimated for the purpose of encoding or decoding. An example with two moving objects is shown for the sake of illustration.

Motion estimation is itself known in the art and so is described herein only to the extent necessary to provide suitable background for the present invention. According to International Standards for Video Communications such as MPEG 1, 2 & 4 and H.261, H.263 & H.264, motion estimation is based on block-based partitions of source frames. For example, each block may comprise an array of 4x4 pixels, or 4x8, 8x4, 8x8, 16x8, 8x16 or 16x16 in various other standards. An exemplary block is denoted by B_i in Figure 1. The number of pixels per block can be selected in accordance with the required accuracy and decode rates. Each pixel can be represented in a number of different ways depending on the protocol

adopted in accordance with the standards. In the example herein, each pixel is represented by chrominance (U and V) and luminance (Y) values (though other possible colour-space representations are also known in the art). In this particular example chrominance values are shared by four pixels in a block. A macroblock
5 MB_i typically comprises four blocks, e.g. an array of 8x8 pixels for 4x4 blocks or an array of 16x16 pixels for 8x8 blocks. Each pixel has an associated bit rate which is the amount of data needed to transmit information about that pixel.

Figure 2 is a schematic flow chart of a data compression method in accordance
10 with a preferred embodiment of the present invention. The method preferably uses block matching based on objective metrics. That is, one or more metrics of a current target block to be encoded are compared to the corresponding metrics of a plurality of other blocks, and a reference block is selected based on a measure of similarity of those metrics. The reference block then forms the basis for
15 encoding the current block by means of prediction coding, either intra-frame coding in the case where the reference block is from the same frame f_{t+1} or inter-frame coding where the reference block is from a preceding frame f_t (or indeed f_{t-1} , or f_{t-2} , etc.). The idea behind the block matching is to choose a reference block which will result in a small residual signal when the current block is encoded
20 relative to that reference block (i.e. so that the difference between the actual current block and the prediction will be small when predicted from the selected reference block), hence requiring only a small number of bits to encode.

It is a particularly preferred aspect of the technique that block matching is carried
25 out in the frequency domain, i.e. based on comparison of one or more metrics of a transformed representation of the blocks.

Hence at step S1, a frequency domain transform is performed on each portion of the image of each of a plurality of frames, e.g. on each block. Each block is
30 initially expressed as a spatial domain representation whereby the chrominance and luminance of the block are represented as functions of spatial x and y coordinates, $U(x,y)$, $V(x,y)$ and $Y(x,y)$ (or other suitable colour-space representation). That is, each block is represented by a set of pixel values at

different spatial x and y coordinates. A mathematical transform is then applied to each block to transform into a transform domain representation whereby the chrominance and luminance of the block (or such like) are represented as a function of variables such as wavenumbers k_x and k_y having dimensions of 1/wavelength, i.e. $U(k_x, k_y)$, $V(k_x, k_y)$ and $Y(k_x, k_y)$. That is, the block is transformed to a set of coefficients representing the amplitudes of different spatial frequency terms which can be considered to make up the block. Possibilities for such transforms include the Discrete Cosine transform (DCT), Karhunen-Loeve Transform (KLT), or others. E.g. for a block of $N \times M$ pixels at discrete x and y coordinates within the block, a DCT would transform the luminance $Y(x,y)$ to a set of frequency domain coefficients $Y(k_x, k_y)$:

$$Y(k_x, k_y) = \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} Y(x, y) \cos \left[\frac{\pi k_x}{2N} (2x + 1) \right] \cos \left[\frac{\pi k_y}{2M} (2y + 1) \right]$$

And inversely, the x and y representation $Y(x,y)$ can be determined from a sum of the frequency domain terms summed over k_x and k_y . Hence each block can be represented as a sum of one or more different spatial frequency terms having respective amplitude coefficients $Y(k_x, k_y)$ (and similarly for U and V). The transform domain may be referred to as the frequency domain (in this case referring to spatial frequency).

In some embodiments of the invention, the transform could be applied in three dimensions. A short sequence of frames effectively form a three dimensional cube or cuboid $U(x,y,t)$, $V(x,y,t)$ and $Y(x,y,t)$. In the case of a three dimensional transform, these would transform to $U(k_x, k_y, f)$, $V(k_x, k_y, f)$ and $Y(k_x, k_y, f)$. The term "frequency domain" may be used herein may be used to refer to any transform domain representation in terms of spatial frequency (1/wavelength domain) transformed from a spatial domain and/or temporal frequency (1/time period domain) transformed from a temporal domain.

30

Once the blocks are transformed into the frequency domain, block matching is performed by comparing the transformed frequency domain coefficients of the

current block to those of a plurality of other blocks. A reference block for prediction coding of the current block (either intra or inter) can then be selected based on a measure of block similarity determined from the frequency domain coefficients.

5

An advantage of block-matching in the frequency domain is that the transform tends to compact the energy of a block into only a few non-zero (or non-negligible) coefficients, and thus that comparison can now be based only on only a few frequency coefficients instead of all the coefficients on the block. That is, since the frequency transform concentrates the energy into only a few significant coefficients, then efficient block matching (or indeed other processing) can be performed by only considering those few significant coefficients. This technique thus provides a unique approach to the problem of data compression in video transmission. Although not every pixel need be directly compared when comparing patterns, nevertheless, a complete search can be achieved.

10

15

For example consider an illustrative case as shown in Figures 1a and 1b. Here, the representation of a block in the frequency domain is achieved through a transform which converts the spatial domain pixel values to spatial frequencies.

20

Figure 1a shows some example pixel values of four 8x8 blocks in the spatial domain, e.g. which may comprise the luminance values $Y(x, y)$ of individual pixels at the different pixel locations x and y within the block. Figure 1b is the equivalent in the frequency domain after transform and quantization. E.g. in Figure 1b such coefficients may represent the amplitudes $Y(k_x, k_y)$ of the different possible frequency domain terms that may appear in the sum. The size of the block in spatial and frequency domain is the same, i.e. in this case 8x8 values or coefficients. However, due to the properties of these transforms then the energy of the block is compacted into only few coefficients in the frequency domain, so the entire block can be considered by processing only these few coefficients.

25

30

As can be seen from this example, only four values need to be processed to find a match for these four blocks in the frequency domain, whereas in the spatial domain there are 256 values that would need to be processed. Thus unlike prior

techniques, the present invention may allow a full true search to be performed but without the need to “touch” every pixel in the block, i.e. without the need to process each individual pixel.

5 It will be appreciated that while blocks and macroblocks are referred to herein, the techniques can similarly be used on other portions definable in the image. Frequency domain separation in blocks and/or portions may be dependent on the choice of transform. In the case of block transforms, for example, like the Discrete Cosine transform (DCT) and Karhunen-Loeve Transform (KLT) and
10 others, the target block or portions becomes an array of fixed or variable dimensions. Each array comprises a set of transformed quantized coefficients. E.g. in the more detailed example of Figure 5A, each macroblock MB of 16x16 pixels may be represented in the frequency domain by 16 luminance blocks and 8 chrominance blocks; each block $b_0 \dots b_{23}$ having a 4x4 array of quantized
15 coefficients.

According to another preferred aspect of the present invention, block matching may be performed within a sorted list based on an index value reflecting the relative importance of the block. In this case the selection of matching blocks may
20 be performed based on an aggregate of values used for the importance indexing. A preferred example will now be described with reference to steps S2 to S6 of Figure 2 and the example blocks of Figure 5A.

At Step S2, each block $b_0 \dots b_{23}$ in the frequency domain is assigned an index
25 value derived from one or more of its frequency domain coefficients. For example, the index value may represent the energy of the block. E.g. this may comprise an aggregate over the coefficients of the block, such as a number of zero coefficients, number of non-zero coefficients, or an average or total value of the moduli of the coefficients in each block.

30

At Step S3, the blocks from at least one frame are then sorted based on the index value. This may involve generating a sorted list in which the entries represent blocks ordered according to their index values, e.g. their block energies.

At Step S4, a subset of candidate blocks is identified from the sorted array by determining a search range or threshold Δ based on the index values. The candidate blocks will be potential matches as reference blocks for use in prediction coding of a current block to be encoded. This is illustrated in Figure 3. For example this may be achieved by determining an energy range $\pm\Delta$ from the current block to be encoded, and determining that all blocks within that range of the current block are candidates for potential selection as a reference block (i.e. candidates for a "match" to the current block for the purpose of prediction coding).

10

At Step S5, the candidate blocks are then evaluated for similarity. For example, block similarity is preferably determined based on bit rate, where the bit rate is a measure of the number of bits that would need to be transmitted in order to define the residuals for the current block if predicted from each candidate block. An example of this will be discussed in more detail shortly.

15

At Step S6, the best matching candidate is determined based on its similarity, and the current target block is encoded relative to that matching candidate. The encoding comprises subtracting the frequency domain coefficients of the reference block from those of the current block in order to generate a residual signal, and then encoding the residual of the current block into the encoded bitstream along with the identity of the respective selected reference block (instead of encoding the target block's actual absolute coefficients). The reference block is thus used as a prediction of the current block. The residual is the difference between the frequency domain coefficients of the current block and the frequency domain coefficients of the reference block, which requires fewer bits to encode and so the encoding results in a compressed video signal. The best candidate for use as the reference block is preferably selected by calculating the bit rate that would be required to transmit the residuals for the current block based on the candidate plus overhead information identifying the candidate block, in comparison with the bit rate that would be required for other such candidates. It will be readily appreciated that a match does not imply identical blocks, but blocks that are sufficiently similar that residuals can be transmitted at a lower bit rate.

20

25

30

Figure 3 is a graph illustrating the arrangement of a sorted array. The list of sorted blocks is shown on the horizontal axis, with block energy index value on the vertical axis. The block energy index value is an example of an objective metric derived from the block's coefficients.

As described above, a best matching reference block is selected having an index within a certain search range or threshold Δ . Thus according to one preferred aspect, the invention provides a method of searching amongst the blocks for matches based on similarity of their indices. By searching for matches by their energy index or such like, this advantageously expands the potential for matches to anywhere within the frame or another frame. Hence the matching need not be restricted to adjacent regions of the target block. For instance, blocks having similar energies may achieve a good even if located on opposite sides of a frame, e.g. blocks of a similar background area appearing at different locations in the frame.

According to another preferred aspect of the invention, block matching is performed by first selecting a subset of candidate blocks based on a first metric (e.g. the index value), and then selecting a matching candidate block from within the subset based on a second metric (e.g. bitrate cost). The matching block is then used as a reference block in prediction coding of a current block to be encoded. One advantage of narrowing the possible matches down to a preliminary subset of candidates based on a first metric, particularly based on an aggregate metric such as block energy, is that unlikely candidates can be eliminated early on without incurring significant processing burden. That is, the sort may be used to discard unlikely candidates. Thus the more processor-intensive comparison based on the second metric, such as the bit rate comparison, need only be performed for a relatively small number of pre-vetted candidates, thus reducing the processing burden incurred by the block matching algorithm. E.g. blocks with very different block energies are unlikely to be good matches and therefore it is unlikely to be worth the processing cost of comparing their potential bitrate contributions. To minimize processing, the selection of a

matching block in Step S6 is preferably performed within a small neighbourhood within the list (search range $\pm\Delta$).

Note though that the sort only gives a certain probability of a match and may be
5 chosen depending on performance considerations. A smaller choice of Δ results
in a lower processing cost but fewer candidates, risking not find the best possible
match. A larger choice of Δ on the other hand incurs a higher processing cost but
will include more candidates and so have a better chance of finding the best
10 match. In embodiments, Δ could even be adapted dynamically based on one or
more performance factors such as available up or downlink bandwidth or available
processing resources. Note also that the same value of Δ need not necessarily be
use in the + Δ direction as in the - Δ direction.

It will be appreciated that at Step S3, the sorted array can be generated for a
15 macroblock (as shown in the example of Figure 5A), for a single frame (for intra
frame data compression) or for a current target frame and one or more reference
frames (for inter frame motion estimation).

In one particularly advantageous embodiment, the same sorted list is used to
20 match multiple target blocks (by determining respective subsets of candidates
within the same list). Further, if the list contains blocks from both the current frame
and one or more preceding frames, then the same list can even be used for both
inter and intra matching within the same sorted list. E.g. when processing a
particular target frame it may be that a good match may is not found within that
25 frame, in which case the method may look to other frames since the complexity is
low and the matching method is the same. According to preferred embodiments of
the present invention, there is no need to use a different method for finding inter
frame matches between frames than is used for intra matching within a frame.

30 By replacing an exhaustive, repetitive search performed for every block with a
single sort that is performed once for an entire frame or even multiple frames, the
selection of a matching block can be performed in a small neighbourhood using
the sorted list. Preferably the sort is performed once for multiple frames, so that

both inter and intra matches can be processed at the same stage over the same sorted list. E.g. this may involve looking for a match within the sorted list of the current frame and, if no satisfactory match is found, looking into the sorted lists of one or more other frames to find a better match.

5

The above-described aspect of the present invention thus provides a method of compressing video data which can be applicable both to intra frame compression and to inter frame motion estimation. In the past, algorithms have adopted different approaches to inter versus intra data compression. The invention on the other hand can advantageously provide a unified technique used for both intra and inter frame prediction.

10

Another benefit of the method is that due to its low complexity, the number of used reference frames can be substantially higher in comparison with existing algorithms.

15

Furthermore, note that conventional motion estimation predicts each block from an area offset by any arbitrary number of pixels or even fractional number of pixels in the horizontal or vertical direction, whereas the approach used in the present invention differs by restricting the prediction to performing only block-to-block matching. That is, matching on a block-by-block basis whereby a block is matched to another whole block (rather than any arbitrarily offset area requiring a motion vector specifying any number of pixels or fractional number of pixels). In a particularly advantageous combination of features, the block-to-block matching may be performed in the frequency domain where efficiency can be derived by predicting only a subset of frequency domain coefficients between two or more blocks.

20

25

Once a matching block has been selected at step S6 and the current target block has been encoded relative to that matching block, the residual of the frequency domain coefficients is output via an entropy encoder for inclusion in the encoded bitstream. In addition, side information is included in the bitstream in order to identify the reference block from which each encoded block is to be predicted at

30

the decoder. Each block may be identified by its location, i.e. by its address or position within a particular frame. Each frame may be identified by a frame number. Because of the above distinction, note that the side information identifying the selected reference block may be signaled in the bitstream in the form of a block address identifying the location of the reference block in terms of a whole number of blocks. This may take the form of an absolute block address, i.e. a position relative to a fixed point in the frame. Alternatively it may take the form of a relative address. The side information may also identify the frame of the selected reference block if candidates may be selected from a plurality of different potential frames.

This is different from a conventional motion vector, which is signaled in the form of a small vector relative to the current block, the vector being any number of pixels or fractional pixels.

15

As mentioned, the VC-1 video codec has an intra prediction mode in which the first column and/or first row of AC coefficients in the DCT domain are predicted from the first column (or first row) of the DCT blocks located immediately to the left or on the top of the processed block. However, this differs from the approach used in aspects of the present invention in that it is restricted to using only predetermined spatially-adjacent coefficients for intra prediction. VC-1 does not allow intra matching to a *selected* reference block, e.g. selected based on block energy and/or bitrate contribution (and therefore VC-1 also does involve signaling the identity of a selected reference block to the decoder).

20

Figure 4 is a schematic block diagram showing the architecture of an encoding technique in accordance with one embodiment of the invention. The raw input video stream is received by a forward transform stage 2. The output of this stage is supplied to a forward quantization stage 4. The forward transform stage 2 applies spatial or spatial-temporal transform into the frequency domain as a first coding step. The forward quantization stage 2 applies quantization and generates for each block a set of quantized coefficients in the frequency domain. The transform coefficients from the forward quantization stage 2 of each intra frame in

25

30

the temporal domain of the input video stream are supplied to an intra prediction stage 6.

The intra prediction stage 6 operates to locate candidate blocks for prediction within each frame, using the method described above. The transform coefficients of inter frames are supplied from the forward quantization stage 4 to an inter-prediction stage 8, which separates the candidate blocks for prediction of target frames as described above. The outputs of the intra prediction stage and the inter-prediction stage 8 are supplied to an entropy encoder 10 which encodes the data to provide an encoded stream for transmission. The encoded stream contains a sequence of information comprising, for each block, a set of coefficients (actual or residual), data defining whether the block is to be predicted and, if it is, an indication of the reference block from which it is to be predicted. The identity of the reference block may be represented in the encoded bitstream as an absolute block location within a frame, i.e. by reference to a fixed point, and not relative to the current block. Alternatively the location may be represented in the encoded bitstream as a difference between the location of the current block and the block from which it is predicted. Either way, the block location is expressed in terms of a number of intervals of whole blocks, i.e. as a block address, and so a benefit is achieved because this requires far less overhead to encode than a conventional motion vector expressing an offset in pixels or even fractions of pixels.

Note that the arrangement does not involve a loop back into the spatial domain as in the standard encoder of Figure 7. Hence block matching is performed in the transformed frequency domain based on frequency domain coefficients.

Note also that in preferred embodiments, the selection of the reference block is performed in the quantized domain, i.e. a non-distorting, lossless environment. Therefore no additional distortion is applied to the candidate blocks or current blocks before performing the selection.

Figure 5A illustrates schematically a prediction example. The case illustrated in Figure 5A is where the technique is used for intra prediction between different blocks of the same macroblock in one frame. Figure 5A illustrates on the left hand side luminance and chrominance data transformed into the frequency domain for a macroblock (16x16 pixels). The frequency transformed coefficients are organised into blocks b0, b1, etc, each block comprising a 4x4 array of coefficients. Blocks b0 to b15 represent luminance data (y) for the macroblock, and blocks b16 to b23 represent chrominance data (u,v) for the macroblock.

10 There are different schemes for treating the luma and chroma channels. A common way is the 4:2:0 format which implies that the chroma channels are being downsampled by a factor two in both the horizontal and in the vertical direction.

In the example shown, block b0 contains 16 coefficients: one DC (the first one at coordinate 0,0) and 15 AC coefficients (the rest of the block). The DC represents the so-called "constant" value of luminance (for 'Y' blocks) and of the chrominance (for 'U' and 'V' blocks), and the ACs form the variable part meaning their contribution for each pixel is different. The combination of the DC and all ACs are used to represent the value of each pixel after decoding based on the used transform. The 16x16 luma frequency domain coefficients 'Y' are fully utilized to represent 16x16 spatial domain pixels. In the explanation above, the chrominance 'Us' are sub-sampled. This format is known as YUV 4:2:0, which means that four luminance pixels for each 2x2 square of the Y pixels share one 'U' and one 'V' pixel respectively.

25

There also exist other formats known as YUV 4:2:2 or YUV 4:4:4, in which the chrominance is not sub-sampled at all in YUV 4:4:4, or where the chrominance has twice as much data in 4:2:2. The present invention can work for any of these formats.

30

In the described example, the blocks b0 to b23 for the macroblock are sorted based on a measure (index value) of block energy or activity. Figure 3a illustrates an example of block-sorting. The block energy used to order the sort can be

measured in a number of different ways. According to one technique, the sort is based on the number of zero value coefficients in a block. In another technique, the sort is carried out using the average value of the modulus of non zero coefficients. Using a measure Δ of block energy, a search range is established within the sorted list as illustrated in Figure 3 to identify candidate blocks (Step S4 of Figure 2). The best candidate for prediction is then established as described above based on bit rate evaluation (Step S6 of Figure 2).

The right hand diagram in Figure 5A illustrates the effect of these predictions. Block b12 is labelled P1 to denote it as the first predicted block. Instead of transmitting the actual coefficients in block b12, coefficients (residuals) representing the differential between block b12 and b10 are transmitted, together with the information that in the transmitted data block 12 has been predicted from reference block 10. An indication of the reference block 10 is also transmitted, e.g. identified by its frame number and position in the frame. This is shown schematically in the list on the right hand side of Figure 5A where P1 denotes prediction 1, block 12 minus block b10 in the luma block. The next candidate to be selected is block 20 labelled P2 which is predicted from block b21. The process continues and in this case results in 7 predicted blocks. This results in a reduction in the number of coefficients to be transmitted by 9 (from 132 to 123). In a specific example, when the video data is encoded for transmission in bins, this has the effect that bins 122 and 1008 are removed, while the content of bins 224 and 288 are increased. In Figure 5A, the arrows denote the relationship between a predicted block and the block from which it is being predicted.

Figure 5B shows a prediction example for motion prediction between different blocks of different macroblocks of two frames.

Figure 6 is a schematic block diagram of a decoder for decoding a video stream which has been subject to the block prediction technique described above. In addition to the encoded coefficients, the video stream includes data defining the predicted blocks, the identity of the blocks from which they have been predicted and the order in which they have been predicted. The encoded stream is supplied

to an entropy decoder 12 which determines for the incoming data whether the blocks to be decoded are for reconstruction of an intra frame or reconstruction of an inter frame. Blocks for reconstruction of an intra frame are passed to intra reconstruction stage 14, while blocks intended for reconstruction of an inter frame are passed to inter reconstruction stage 16. A predicted block is reconstructed by adding the residuals to the correspondingly located coefficients in the block it is predicted from. The output of the reconstruction stages 14 and 16 are supplied to an inverse quantization stage 18 and then to an inverse transform stage 20 where the quantization coefficients are transformed from the frequency domain into the time domain as a decoded stream.

Details of a preferred technique for matching blocks based on bitrate contribution are now discussed in more detail. This technique decreases the bitrate in video compression by means of block prediction in the quantized domain. The input to the method is e.g. a slice or a set of slices of blocks of transformed and quantized coefficients (e.g. residuals from the H.264). A slice means a group of macroblocks, so one slice per frame means all macroblocks in the frame belong to the slice. For each transformed and quantized block in the current slice a block from previous encoded slices or a block in the current slice (care has then to be taken to ensure a decodable stream) is a potential candidate to be used for prediction in order to reduce the bitrate (compared to direct entropy coding of the block itself). An example embodiment of a predictor and the "optimal" selection of the block to be used for prediction and required side-information to identify that block (needed description for reconstruction in the decoder) is described below. The side information is entropy encoded into the encoded bitstream along with the residual, by entropy encoder 10.

In the preferred embodiments, the present invention performs block matching using two classes of metrics: one based on an aggregate or pattern of the block (e.g. energy, structure etc.) and a second based on bit rate. These two metrics are used in two separate stages: the first stage to sort and the second stage for the RD loop. In particularly preferred embodiments, the RD loop rate target is not only to find two blocks that can predict each other closely in terms of rate, but also

to solve this problem for groups of blocks at the same time. One simple example could be the following patterns – (a) 1,2,1,2,1,2,1,2 and (b) 46,47, 46,47, 46,47, 46,47, that will result in (a) 1,2,1,2,1,2,1,2 and (b) 45, 45, 45, 45, 45, 45, 45, 45. That is to say, multiple blocks can be matched from the same sorted list of candidate blocks, including potentially both interframe and intraframe prediction being performed based on the same sorted list.

The advantages are improved entropy coding due to improved redundancy removal prior to an arithmetic or Huffman coder in the entropy encoder 10. Compared to VC-1 [1, pp.251] there are a number of potential advantages: (i) all coefficients in the block are used in the prediction and not just the first row and/or column; (ii) all blocks in the frame/slice as candidate blocks for prediction and not just block to the left or on the top; (iii) generalized prediction structures, e.g. weighted prediction or prediction from multiple blocks; and (iv) explicit rate estimation for finding the best block for prediction (taking the cost for side information into account).

Let $X(m, n)$ denote a block $m \in M$ (a frame/slice consists of M blocks in total) of quantized coefficients (e.g. quantized DCT coefficients) at time-instance n . The blocks are conventionally fed to an entropy coder 10 (in H.264 more specifically the context adaptive variable length coder or the context adaptive binary arithmetic coder). That is, from the point where we have $X(m, n)$, lossless compression is performed, i.e., the distortion is fixed. The method seeks to remove remaining redundancies (and thereby reduce the rate) prior to the arithmetic coder by means of a predictor. In one embodiment the prediction is formed as a subtraction between a current block and a reference block. The optimal indices (o_{opt}, p_{opt}) for prediction of the current block $X(m, n)$ is selected based on rate calculation/estimation, i.e.,

$$(o_{opt}, p_{opt}) = \underset{o, p}{\operatorname{argmin}} (R(X(m, n) - X(o, p)) + R(o, p)) \quad [1]$$

where $R(X(m, n) - X(o, p))$ denotes the bitrate of the prediction residual and $R(o, p)$ the bitrate of side-information (i.e., the bitrate for transmission of the prediction

block index o of frame p). The rate estimation can for instance be provided from parts of the arithmetic coding routine where the sum of \log_2 of the symbol probabilities can be used to estimate the rate. It could also be beneficial, from e.g. a computational aspect, to approximate the criterion in equation [1] by using
 5 another measure that correlates well with the rate. Generally, any metric can be used that relates in some way to a number of bits that would be required in the encoded bitstream to encode both the residual block and the side information identifying the respective reference block (i.e. would be required for each candidate if that candidate was chosen as the reference block), whether the
 10 metric is a direct measure the number or rate of bits or a metric that correlates with the number/rate.

The search for the optimal predictor can be made computationally more efficient by pre-ordering the candidates such that potential good candidates are located in
 15 the proximity of a specific position in an ordered array. Let $Y(k, n)$ now denote element k in an M dimensional ordered array of the block indices $m \in M$ of frame n according to some measure. For instance, the blocks $X(m, n)_{m \in M}$ can be ordered according to their ascending energy (or some other signal dependent properties).

20

To find the predictors we start e.g. with the first block in the current frame/slice in the raster-scan order (or some other order beneficial from either a rate or computational complexity point of view) and find its position in the ordered array $Y(n)$ of the current frame and the (re-)ordered arrays of the previously processed
 25 frames $Y(n-NumRef), \dots, Y(n-1)$. NumRef is the number of reference frames, i.e. here the number of previous quantized frames that have been processed and can be used for inter prediction. As prediction candidates from the current frame/slice we select the candidates that are located within the range $\pm W$ around the current index in the ordered array, i.e., the "intra" prediction candidates plugged
 30 into expression (1) are the blocks corresponding to the sorted indices $Y(q(n)-W, n), \dots, Y(q(n)-1, n)$; and $Y(q(n)+1, n), \dots, Y(q(n)+W, n)$; where q denotes the position of the current block in the ordered array. Note that special caution has to be taken to avoid cyclic predictions, i.e., avoid prediction of block m from block n if

block n has already been predicted from block m , making decoding infeasible. It should also be mentioned that direct encoding (i.e., no prediction) of the residual is also included as a candidate for the rate estimation.

- 5 Similar to the selection of candidates for intra prediction, the inter prediction candidates are selected as $Y(q(n-i)-W, n-i), \dots, Y(q(n-i)+W, n-i)_{i=1: NumRef}$.

All intra and inter candidates are evaluated according to equation (1) and the optimal index pair is selected. This procedure is repeated for all blocks in the frame/slice. The resulting prediction residuals (variable/index differences) together
10 with required side- information for decoding is e.g. arithmetically encoded and sent to decoder.

Referring to Figure 10 one embodiment of the method performed by the encoder
15 is as follows.

Step T1: order all the blocks in the frame according to some measure.

Step T2: set block index to $m=0$.

20

Step T3: find the equivalent position q of the block index m in the ordered lists (both current and previous quantized frames, i.e., find $q(n), \dots, q(n-NumRef)$).

Step T4: select the intra and inter prediction candidates as

25

$Y(q(n)-W, n), \dots, Y(q(n)-1, n);$
 $Y(q(n)+1, n), \dots, Y(q(n)+W, n);$ and
 $Y(q(n-i)-W, n-i), \dots, Y(q(n-i)+W, n-i)_{i=1: NumRef},$ respectively.

The size of the search range W is a trade-off between performance and
30 computational complexity.

Step T5: find the best candidate according to expression [1] or some approximation of it.

Step T6: send optimal prediction residual together with side-information (e.g. the position of the residual block within the frame and the position (e.g. space and time) of the block that was used for prediction) to the arithmetic coder.

5

Step T7: increment block index $m = m + 1$ and go to step T3, until $m = M$ when the method moves to the next frame $n=n+1$.

Referring to Figure 11, one embodiment of the method performed by the decoder is as follows.

10

Step U1: decode all prediction residuals and side information (this gives a frame of prediction residuals together with the description for each block how to undo the prediction).

15

Step U2: reconstruct all blocks that do not depend on unreconstructed blocks (i.e. (undo prediction)).

Step U3: repeat step U2 until all blocks have been reconstructed.

20

The above example embodiment can be extended in several ways. For instance it could be beneficial to use weighted prediction or prediction from multiple blocks. Additional side information would be needed to be transmitted which for weighted prediction and prediction using multiple blocks would be prediction weights and block positions/addresses.

25

An illustration of the prediction in the encoder is shown in Figure 9. This gives a high-level illustration of the block prediction in the encoder. The prediction residual together with side information is sent to the entropy coder. In the decoder the reverse procedure is performed, i.e. first reconstruct the residual frame and then reconstruct the frame given side information.

30

The above described embodiments of the present invention may provide several advantages. Matching blocks are located by examining the difference between blocks to be certain that the bit rate of the ultimately transmitted video data will be reduced with respect to the bit rate for sending coefficients of those blocks.

5 Moreover, the pre sort has identified candidate blocks within which this comparison takes place. The blocks do not have to be physical neighbours in the image frame – instead, they are sorted on the basis of an index value associated with the blocks, for example, representing energy. This allows a best matching block to be selected from any part of a frame (or indeed a different frame). When

10 selecting best candidates, the comparison of bit rates can take into account the overhead information that needs to be transmitted to identify that the block is a predicted block, and to identify the block from which it is predicted. The identity of the block from which it is predicted can be provided to the decoder in the form of an location within the frame expressed as a number of intervals of whole blocks,

15 i.e. a block address, rather than by a motion vector expressed as an offset in terms of a number of pixels or even fractions of pixels.

The method described removes redundancy in the temporal and frequency domain before and/or after quantization in a compressed digital video stream by

20 means of block prediction. The input to the method is a set of transformed and/or quantized transform coefficients of a set of frames in the temporal domain of the input video stream. The input video stream frame can be separated into blocks and groups of blocks. The groups of blocks are not limited by the location of the individual blocks participating in the group. The prediction is performed between

25 the blocks of the current frame (intra) and is not limited by location of the blocks but by the factor of the block similarity. The same technique can be used for inter frame predictions. Inter frame block matching is not restricted by location either. The block similarity is determined from the point of view of reduction of bit rate.

30 Furthermore, as explained, in a preferred embodiment processing is carried out in the frequency domain where the transform has already compacted the energy of the target object such that comparison can now be carried out using a few frequency domain coefficients instead of a whole image. In these embodiments,

both components of the method, i.e. processing in the frequency domain and the sort versus search, reduce the complexity while maintaining a very high quality. Another benefit of the method is that due to the low complexity of the calculations involved, the number of used reference frames for inter frame motion
5 compensation can be substantially higher in comparison with existing algorithms. Another major benefit is that, due to the low complexity, matches can be made on several level sub block divisions. That is, an image portion can be a macroblock, a block or even a smaller number of pixels than a block. This is because the described method achieves low complexity and therefore incurs fewer of clock
10 cycles, which if desired means that some of the saved complexity can then be spent searching for sub-blocks such as 4x4 or 2x2 sub-blocks instead of just blocks. Alternatively the search could be performed at a higher level of 16x16, 32x32 or 64x64 aggregate blocks for example, which would save on the side information necessary to signal them in the encoded stream.

15

A particular advantage arises from processing in the frequency domain. While there are frequency domain processing models existing, there is none that explores redundancy reduction as in the method described in the following embodiments; in particular which provides a unique set of benefits including
20 complexity reduction, preserving and improving video quality and at the same time lowering the bit rate of the encoded video stream.

The method need not require loop filter or loop back to the spatial domain for motion estimation due to the fact that all processing is now concentrated in the
25 frequency domain. This is a major advantage with respect to existing video coding methods and a point of significant reduction of complexity.

Another advantage is that processing of all the colour components can be done at the same time. That is, processing done in the luminance channel can affect
30 processing done in the chrominance channels.

Another advantage of processing in the frequency domain relates to blocks lying on the edge of a frame or slice of a sub frame. That is, the blocks that lie on the

edge of a frame (or if a sub frame separation in multiple slices is used, the blocks that are on the edge of the slice) can be efficiently predicted. As the blocks are grouped in accordance with similarity, the method allows grouping of blocks or slices in any order and hence there is no penalty in the prediction of blocks sitting
5 on the edge of a slice or frame. This is a significant improvement in comparison with the current FMO (Flexible Macroblock Ordering) in the current Standards like MPEG-4 AVC/H.264.

Another advantage of the described embodiments of the invention herein is that
10 deep sub-block sub-divisions can be utilised without excessive processor load.

Note that the different preferred techniques discussed above need not necessarily be used in conjunction with one another. For example, it is possible to perform block matching in the frequency domain without using the additional technique of
15 a sorted list based on block energy or other such index. Alternative block matching techniques could also be used, for either intra and/or inter frame block matching, e.g. by matching based on a measure of correlation or a measure of minimum error. Conversely, it is possible to use the sorting technique for block matching without a frequency domain transform, e.g. by determining a measure of
20 block energy based on the spatial domain coefficients (though this is less [preferred since it will tend to be more computationally intense).

Further, where sorting is discussed as a method of determining a subset of candidates within a search range Δ , note that it is not necessarily required to
25 rearrange list entries in memory. More generally, the search for candidates may be performed by any method of identifying blocks having an energy or other index within the desired range.

Further, the sort index need not necessarily be a measure of block energy.
30 Another possibility would be a metric relating to the structure of the block, such as the structural similarity index (SSIM). In other embodiments, multiple metrics could be combined in order to determine the index used for sorting. Furthermore, once the list is sorted, aspects of the invention need not necessarily be limited to

finding the best match from amongst the candidates based on bitrate contribution. Other second metrics could be used for this purpose, e.g. a more conventional motion based matching as used in H.264.

5

Signalling Blocks by Address

The above describes a particularly advantageous method of selecting reference blocks from a frame; but regardless of how blocks are selected from frames, the present invention provides an improved method of encoding the identity of reference blocks for transmission to the decoder. Exemplary details of a method of signalling the intra and inter prediction information for prediction in the frequency domain are now discussed in more detail.

15 As mentioned, according to one aspect of the present invention a block is matched only to another whole block rather than to a block-sized area offset by any number of pixels as in more conventional block matching techniques. Therefore the signalling algorithm of the present invention sends block addresses instead of motion vectors, i.e. represented in terms of a whole number of blocks rather than a pixel offset. Note however that whilst the term "block" may be used

20 herein, in its most general sense this is not intended to imply and particular size, shape or level subdivision. It will be appreciated that in different schemes then various different divisions and subdivisions may be referred to by terms such as macroblock, block and sub-block, etc., but that the term "block" as used most

25 generally herein may correspond to any of these or indeed any other constituent image portion being a division of a video frame corresponding to multiple pixels. Whatever manner of division is employed, according to the present invention the address of the reference portion for use in prediction is signalled as a whole number of multi-pixel portions instead of a pixel offset.

30

In embodiments, the bitstream may also contain one or more prediction method flags indicating a prediction method to be used by the decoder (corresponding to that used by the encoder).

Further, the bitstream may contain a frame number of the reference block, as the reference block for prediction can be chosen from any of multiple different frames.

- 5 In one particularly preferred embodiment, the side information signalled in the bitstream to the decoder comprises: frame number, an addition or subtraction flag, absolute value flag, a macroblock address, a block address within the macroblock, and a sub-block address within the block. The signalling structure of this side information is shown in the following table.

10

Field	No. Bits
Frame Index (Frameldx)	4
Add/Sub	1
Nat/Abs	1
Macroblock Address (MBAAddr)	9
Block Address (BlockAdr)	3
Sub-block Address (SubBAdr)	2

For each resolution the prediction indexes can be encoded as follows. This shows signalling structure size and encoding for a plurality of different resolutions.

Inter	SIF	WVGA	480p	4CIF	720p	1080p	4kx2k	8kx4k	Values	Max		4 bits
Frameldx	4	4	4	4	4	4	4	4	-8 .. 7		0	Intra
Sub/Add	1	1	1	1	1	1	1	1	0 .. 1		1	-1
Nat/Abs	1	1	1	1	1	1	1	1	0 .. 1		2	-2
MBAAddrX	5	6	6	6	7	7	8	9	0 .. Max	480	3	-3
MBAAddrY	4	5	5	6	6	7	8	9	0 .. Max	270	4	-4
BlockAdr	3	3	3	3	3	3	3	3	0 .. 5		5	-5
SubBAdr	2	2	2	2	2	2	2	2	0 .. 3		6	-6
Total/B	20	22	22	23	24	25	27	29			7	-7
Total/MB	120	132	132	138	144	150	162	174			-8	List

15

This improved prediction scheme is more effective than the current prediction scheme which use a higher bit rate to signal only part of the information that the improved scheme can transmit. The streamlined inter and intra prediction allows for simplified signalling method. Figure 3b shows a block matching prediction example achieving bit savings. The table below shows the effective side information and coding for multiple resolutions.

	Res X	Res Y	MB_x	MB_y	MBs	MBBits	UpToMBs	Bs	BBits	UpToBs	Bits_X	Bits_Y	Bits_XY
SIF	320	240	20	15	300	9	512	1800	11	2048	5	4	9
WVGA	640	400	40	25	1000	10	1024	6000	13	8192	6	5	11
480p	640	480	40	30	1200	11	2048	7200	13	8192	6	5	11
4CIF	704	576	44	36	1584	11	2048	9504	14	16384	6	6	12
720p	1280	720	80	45	3600	12	4096	21600	15	32768	7	6	13
1080p	1920	1080	120	68	8160	13	8192	48960	16	65536	7	7	14
4kx2k	3840	2160	240	135	32400	15	32768	194400	18	262144	8	8	16
8kx4k	7680	4320	480	270	129600	17	131072	777600	20	1048576	9	9	18

Signalling Blocks by Means of a “Global Block List”

The following describes a second improved method of encoding the identity of reference blocks for transmission to the decoder. Again, this may be used regardless of how blocks are selected from frames. In embodiments this method may furthermore extend the available candidates to include certain “notional” or “artificial” blocks (rather than just actual blocks appearing in actual frames).

As discussed with reference to Figure 12, according to this aspect of the present invention the encoder generates a table of prediction blocks (i.e. reference blocks for use in prediction) having the most regularly (often) used block characteristics. A block in the table can then be referenced in the encoded signal instead of a block in the frame. The table of most common blocks is determined during

encoding and will be updated dynamically for transmission to the decoder. Thus the encoder generates an ad-hoc codebook for signalling the reference block to the decoder.

- 5 For example, certain blocks such as those shown in Figures 1b, 5A or 5B may occur more regularly within a certain frame, sequence of frames or part of a frame to be encoded (i.e. a higher number of instances). If certain blocks (i.e. certain sets of block coefficients or approximate sets) occur often enough, then it may become more efficient to dynamically maintain and transmit to the decoder a look-
- 10 up table of these regularly encountered blocks and then signal the identity of reference blocks used in the prediction coding by reference to an entry in the look-up table, rather than identifying the block by some other means such as a motion vector or block location address.

15	Table entry #	Block definition
	0	$B_a(\dots)$
	1	$B_b(\dots)$
	2	$B_c(\dots)$
	3	$B_d(\dots)$
20	4	$B_e(\dots)$
	etc...	

Each block definition indicates a certain respective regularly-encountered set of block coefficients (or approximate set).

25

- Preferably the table is updated and transmitted periodically, and each updated table may be interleaved into the encoded data stream as shown in Figure 12 (though other methods of separate transmission are not excluded). Thus after a certain table L_n is transmitted in the bitstream to the decoder, then one or more
- 30 subsequent encoded blocks $(n,1)$, $(n,2)$, etc. that are each encoded by prediction coding based on another reference block are each transmitted in conjunction with side information $S_{n,1}$, $S_{n,2}$ indicating a respective entry in the look-up table. When the look-up table is updated and retransmitted (L_{n+1}), then subsequent side

information $S_{n+1,1}$ may then signal a reference block for use in prediction by indicating an entry in that updated table. The decoder stores a copy of the most recent look-up table L received in the bitstream and uses it in conjunction with the side information S to identify a reference block for use in predicting a current block
5 to be decoded (and combines the predicted block with the respective residual data).

This technique is particularly useful when block matching is performed in the frequency domain, because the energy of each block compacted into only a few
10 non-zero coefficients - for example see Figure 1b. In this case certain blocks are likely to be selected in the block matching process often enough for the maintenance and transmission of a dynamic look-up table to be an efficient choice of technique. Nonetheless, the described technique is not limited to frequency domain processing nor selection based on block sorting, and could also be used
15 to encode the results of other block matching techniques.

Note that in most practical embodiments the look-up table will not be exhaustive. I.e. some reference blocks will not be selected very regularly by the block matching process, and those blocks will not be included in the look-up table. Such
20 reference blocks may be signalled in the encoded bitstream by a different method, e.g. preferably in the form of a block address identifying an absolute block position in the frame or a relative block address between the current block and respective reference block. That is, a block location expressed in terms of a whole number of multi-pixel blocks (rather than any number of pixels or fractional number of pixels
25 as in a conventional motion vector).

In embodiments the actual coefficients of the most regular reference blocks may be signalled to the decoder in the look-up table L (though an alternative would be for the encoder and decoder to have all possible block definitions that could
30 potentially be signalled in the table pre-stored at either end).

Thus according to the above techniques, a subset of blocks frequently appearing in multiple frames of matched blocks will be separately transmitted to the decoder

in a group. In a particularly preferred embodiments, this group may not only include existing blocks but can also include "artificial" blocks with sets of coefficients that help the prediction process and are calculated from the blocks in the input stream. That is, one some or all of the blocks in the group need not be
5 actual blocks of the frame, but instead could be notional blocks comprising pre-defined "artificial" coefficients which may be set by the system designer or calculated by the encoder, e.g. by averaging, interpolating or extrapolating from other actual blocks regularly found in the video. In the preferred embodiments these artificial blocks would be included in the sorted list or as candidates in the
10 block matching process.

For example, say a number of blocks regularly occur having a particular frequency domain coefficient which is regularly within a certain range, e.g. 200 to 210. In this case the encoder may create an artificial block having a corresponding coefficient
15 with an average or interpolated magnitude within that range, e.g. 205. That way, the regular similar blocks encoded relative to that artificial block will result in only a small residual, e.g. typically no more than about 5 in size, thus reducing the bitrate required to encode that residual. This example of a particular coefficient is considered for the sake of illustration, but note that the greatest benefit of this
20 scheme will be achieved when blocks regularly occur with similar patterns of multiple coefficients, in which case an artificial reference block can be generated having an approximate pattern close to that of a number of those blocks (the single coefficient case is likely to handled well in the entropy encoder anyway). It is these multi-coefficient blocks which tend to incur most of the bitrate in the
25 encoded signal.

The group of blocks which populate the look-up table may be referred to herein as the "global group", which may comprise artificial blocks and/or actual blocks which are extracted from one or more actual frames of the video. As described, the
30 global group is updated dynamically based regularity of use in the prediction coding.

The best candidate is preferably selected as the lowest bit rate contributor, selected from the set of blocks in the global group and the actual blocks in the frames in any combination.

- 5 One unique aspect of this algorithm is in the approach of application of the global group of blocks to aid the prediction process. The extraction of existing blocks to the global group will further reduce the bitrate.

10 The members of this group will be defined during the encoding process and will be based on rate calculation (i.e. how many bits would be required in the encoded bitstream), thus creating a set of coefficients that are the best predictors for a set of blocks in one or multiple slices of frames and not necessarily existing blocks.

15 Due to the rate prediction nature of the process then a perfect match is not necessarily sought, but rather rate reduction which allows for finding the closest match instead of the exact pattern.

20 All of the blocks in the frame or only the blocks submitted as the global list can be shifted in this way, creating another opportunity for bitrate reduction. Also, slices of the frame can be shifted as well as parts of the global group.

25 The look-up table will be periodically updated by allowing the decoder to drop blocks that will no longer be required from the table, and update them with new ones from the incoming stream.

30 In practice, the look-up table may not always be a useful way of signaling reference blocks, e.g. if the encoder determines that few or no blocks are selected significantly more than any others. In this case, the look-up table may not actually save on bits in the bitstream since it will be rarely referenced whilst itself incurring a number of bits in overhead. Therefore in particularly advantageous embodiments, the encoder may be configured to detect some bit-saving threshold, and if the saving is small or even negative then it will cease using the look-up table method (and stop sending updates), and instead will signal

reference blocks by another means such as identifying the location of the reference blocks by their address within the stream. In order to inform the decoder which encoding method is used, the encoder will also include a flag in the bitstream indicating whether or not the look-up table method is currently being used. The flag may also be sent periodically, or on an ad-hoc basis as and when the encoder decides to change mode.

The algorithm described here is lossless. The encoding and the decoding process can be made very low complexity algorithms since the process operates in the frequency domain on a very small set of coefficients (in comparison with the spatial domain process where the complexity is exponentially higher).

Scaling and Rotation

One flaw with conventional codecs which perform motion estimation in the spatial domain is that they require high computational complexity in order to handle prediction based on non lateral (non translational) motion, i.e. scaling or rotation. Scaling occurs when the camera zooms in or out, or an object moves closer or further away from the camera, or indeed if the object expands or shrinks. Rotation occurs when the camera and object or background rotate relative to one another.

An advantage of performing motion estimation in the frequency domain is that the complexity of handling scaling and rotation type prediction is greatly reduced.

For example, consider the illustrative example of Figure 13a. On the left hand side is represented a block B at some point in time, comprising a frequency domain coefficient C. That is, the block may be considered to comprise a frequency domain term such as a sinusoid with amplitude C which varies with some wavelength across the block, e.g. representing a variation in chrominance or luminance across the block. An example of a sinusoidal variation in chrominance or luminance is illustrated in the top left of Figure 13a and the corresponding array of coefficients used to represent the block is shown in the bottom left. Of course

other frequency domain terms may also be present, but for illustrative purposes only one is shown here.

Now consider a corresponding block B' at a later point in time when the camera
5 has zoomed out from the object or scene in question (or the object has moved further away or shrunk). As shown in the top right of Figure 13a, this means the wavelength of the frequency domain term decreases. Suppose for the sake of an illustrative example that the zoom causes the wavelength to halve: the effect in the frequency domain representation of the block is that the coefficient C moves
10 from one position to another, as shown in the bottom right of Figure 13b. That is, the energy of the block is redistributed from one frequency domain coefficient of the block to another. In reality there is unlikely to be a sudden zooming out to exactly half the wavelength, but the zoom may still result in a gradual "fading" transition of the block energy from some coefficients to others – i.e. coefficients
15 representing the amplitude of lower frequency terms will gradually decrease while coefficients representing higher frequency terms gradually increase. E.g. the coefficient at the second (right-hand) position in Figure 13a will gradually increase at the expense of the coefficient at the first (left-hand) position. A similar effect will occur in the opposite direction for a zooming in (or the object moving closer or
20 expanding).

Given such a scenario, the scaling motion of one block can be predicted from another block and may be encoded relative to that block using only very few bits. It is the transformed representation of the blocks that allows this prediction and
25 encoding to be achieved with a low bitrate and low computational burden, because in the frequency domain the scaling will typically involve only a very low-complexity transition of frequency domain terms.

In one scenario the scaling may just be encoded in terms of the difference
30 between the frequency domain coefficients of one block and another, on the basis that as the energy of the block gradually fades from one frequency domain coefficient to another then the residual between one block to the next will be small. According to a further aspect of the present invention however, the encoder

may estimate some parameter of the scaling and signal that parameter to the decoder as side information in the encoded bitstream. For example the parameter may comprise an indication of a scaling factor indicating that the target block can be predicted by scaling the selected reference block by some factor +/-S (e.g. a certain percentage zoom in or out). The scaling may be estimated by a suitable image analysis algorithm in the encoder or by an auto-focus feature of the camera.

In the case where an estimation of the scaling factor is used in the encoding, the residual will represent only the difference between the scaled prediction and the actual target block, and will be even smaller so require even fewer bits to encode. That is, the encoder estimates a local prediction of the scaling to be applied to the selected reference block, subtracts the frequency domain coefficients of the scaled reference block from those of the target block so as to generate a residual (in practice this may just involve comparing the coefficients from shifted positions within the block), and then encodes the target block in the form of the residual, an indication of the scaling parameter, and an indication of the reference block. The signalled scaling parameter enables the decoder to determine the shift or transition to apply to the frequency domain coefficients of the reference block in order to recreate the predicted version. Adding the frequency domain residual at the decoder then recreates the target block. Alternatively, in a non-lossless case the residual may be omitted from the encoding and decoding altogether.

Similar lox-complexity bitrate savings can be achieved by predicting rotation in the frequency domain. An illustrative example is shown in Figure 13b

In this example, the left hand side again represents a block B at some point in time, having a frequency domain coefficient C representing a frequency domain term such as a sinusoid with amplitude C which varies with some wavelength in the horizontal or vertical direction across the block (e.g. representing a variation in chrominance or luminance). An example of a sinusoidal variation in chrominance or luminance is illustrated in the top left of Figure 13b and the corresponding array of coefficients used to represent the block is shown in the bottom left. In this case,

the coefficient C in the first block B represents a certain variation in the horizontal direction across the block. Again note that other frequency domain terms may also be present, but for illustrative purposes only one is shown here.

5 Now consider a corresponding block B' at a later point in time when the camera or object has rotated by 90°. As shown in the top right of Figure 13b, this means the frequency domain term is flipped from a horizontal to a vertical orientation. The effect in the frequency domain representation of the block is that the coefficient C is flipped about the diagonal axis from one position to another, as shown in the
10 bottom right of Figure 13b. That is, the energy of the block is redistributed from one frequency domain coefficient of the block to another corresponding coefficient representing the same frequency term but in the transverse direction across the block.

15 In reality there is unlikely to be a sudden right-angled rotation. However, the effect can be generalised to other angles of rotation. This takes advantage of the fact that the block or macroblock tends to have approximately the same total or average energy when rotated – i.e. the object in question is not emitting more light, just changing orientation relative to the camera. As shown in Figure 13c, in a
20 generalised rotation, the target block (centre) of some later frame F' can be predicted based on a reference block from an earlier frame F, and the contribution to the block energy from the neighbours may be approximated to be small (the contributions being shown shaded black in Figure 13c) and/or similar to the energy lost to the neighbouring from other regions.

25

Hence as the image rotates, the energy of the block will gradually fade between coefficients. That is, the energy from the coefficient at one position is gradually redistributed to it's diagonal counterpart. E.g. in the following rotation the factors "a" and "b" are given by the calculation illustrated in Figure 13e.

30

$$\begin{matrix}
 \bullet & \bullet & C & \bullet & & \bullet & \bullet & b \cdot C & \bullet \\
 \bullet & \bullet & \bullet & \bullet & \rightarrow & \bullet & \bullet & \bullet & \bullet \\
 \bullet & \bullet & \bullet & \bullet & & a \cdot C & \bullet & \bullet & \bullet \\
 \bullet & \bullet & \bullet & \bullet & & \bullet & \bullet & \bullet & \bullet
 \end{matrix}$$

5

That is, $a = \sin(\alpha)$ and $b = \cos(\alpha)$ where α is the angle of rotation. Preferably, the encoder will estimate the rotation using one of a set of computationally relatively low complexity rotations such as 30°, 45°, 60° and 90° as a best approximation.

- 10 For 30°, $a = 1/2$, $b = (\sqrt{3})/2$
- For 45°, $a = 1/(\sqrt{2})$ and $b = 1/(\sqrt{2})$
- For 60°, $a = (\sqrt{3})/2$, $b = 1/2$

15 The residual may then encode the difference between the approximated predicted rotation and the actual block.

In one scenario the rotation may just be encoded in terms of the difference between the frequency domain coefficients of one block and another, on the basis that as the energy of the block gradually fades from one frequency domain

20 coefficient to another then the residual between one block to the next will be small. In a further aspect of the present invention however, the encoder may estimate some parameter of the rotation and signal that parameter to the decoder as side information in the encoded bitstream. For example the parameter may comprise a rotation angle indicating that the target block can be predicted by

25 rotating the selected reference block by the specified angle. The rotation may be determined by an image analysis algorithm in the encoder or by gyro sensors of a mobile terminal in which the camera is housed.

In the case where an estimation of the rotation angle is used in the encoding, the

30 residual will represent only the difference between the rotated prediction and the actual target block, and will be even smaller so require even fewer bits to encode. That is, the encoder generates a local prediction of the rotation to be applied to the selected reference block, subtracts the frequency domain coefficients of the

rotated reference block from those of the target block so as to generate a residual (in practice this may just involve comparing rows of coefficients with columns and vice versa), and then encodes the target block in the form of the residual, an indication of the rotation angle, and an indication of the reference block. The signalled rotation parameter enables the decoder to determine the flip or transition to apply to the frequency domain coefficients of the reference block in order to recreate the predicted version. Adding the frequency domain residual at the decoder then recreates the target block. Alternatively, in a non-lossless case the residual may be omitted from the encoding and decoding altogether.

10

In embodiments the encoder will have the option of using any of the lateral, scaling and rotational types of motion prediction for encoding any given target block. In that case, it is useful to provide a mechanism for selecting the type of prediction to use for each target block. One such mechanism is, for a group of potential reference blocks, for the encoder to try each type of prediction in turn according to a type-hierarchy. Preferably, the encoder first attempts a lateral (i.e. translational) type prediction based on each of the candidates (this typically being the least computationally complex type of motion prediction). If a match is found which will result bitrate contribution within a maximum threshold, then the lateral type of prediction used based on that match, and the matching process halts there so that the scaling and rotation type predictions are not even considered for that target block. That is, if the number of bits required to encode the residual plus side information for the target block based on the best matching reference block is found to be within a certain threshold using a lateral type prediction, then the lateral prediction is used and other types are ignored so as to try to avoid wasting unnecessary machine cycles. However, if no match is found which would provide a bitrate contribution within the threshold using lateral type prediction, then one of a scaling or rotation type prediction may be tried. E.g. the next in the hierarchy may be scaling. The encoder therefore attempts a scaling type prediction based on each of the candidates in the list and tests whether the best matching candidate falls within the maximum bitrate contribution threshold if scaling is used instead of lateral prediction. If so, it encodes the target block based on the best matching reference block using scaling type motion prediction and halts the block

30

matching process for that target block. If not however, the encoder then attempts a rotational type prediction for each candidate in the list and tests whether the best match falls within the bitrate contribution threshold using rotation type prediction. If so, it encodes the target block accordingly. If no matching candidate
5 is found within the maximum bitrate contribution threshold, the encoder may accept the best of a bad lot, or may extend the list of candidates, or may encode by conventional intra encoding or encoding of absolute values.

In other embodiments the hierarchy may be different, e.g. lateral, rotation, scaling.
10 An alternative would be to compare all types of prediction together without hierarchy. However, that would incur a high processing burden and is less likely to be desirable for a live video stream.

The encoder may signal an indication of the type of prediction in the encoded
15 bitstream, so the decoder knows what type of prediction to apply. As mentioned, the encoder may also signal a parameter of the motion. In the case of rotation, the signalled information may indicate a degree of the rotation. In the case of scaling, the signalled information may indicate a scaling factor. This information allows the decoder to reconstruct a prediction of the target block based on the signalled
20 rotation

Referring to Figure 13d, prediction of rotation in the frequency domain can be particularly advantageous when combined with the feature of selecting reference blocks from a sorted list, as discussed above. Figure 13d shows a screen or video
25 window 50. As shown, if a large area is rotated then a closely matching candidate B for rotation type prediction of a target block B' may in fact be found a large distance away within the screen or viewing window 50. A conventional codec which predicts blocks only based on spatially neighbouring regions of the image would miss this situation. However, using the sorted list according to certain
30 aspects of the present invention, blocks from any part of the screen may become candidates for prediction. Particularly when the list is sorted according to block energy, then blocks which closely resemble rotations of one another will become very close in the sorted list (regardless of distance from one another) since the

rotation typically involves little variation in total block energy. Therefore a sorted list in which candidates are identified based on similarity of block energy is particularly likely to find good candidates for rotation type prediction.

5 The scaling and rotation types of prediction can also be particularly advantageously implemented using the feature of signalling reference blocks according to a global block list of the kind discussed above. In such an implementation, the reference blocks in the global block list can include representations of scaled and rotated patterns. For example, an artificial reference
10 block may be generated which is suitable for prediction coding of multiple target blocks according to a number of different types of prediction.

E.g. consider an artificial reference block having energy condensed into the following two non-zero coefficients:

15

224	0	0	0
0	0	0	0
280	0	0	0
0	0	0	0

20

This can be used to encode a 4x4 target block according to any of the following predictions with a reduced bitrate.

90° rotation

25

~~~~~

|     |   |   |   |   |     |   |     |   |
|-----|---|---|---|---|-----|---|-----|---|
| 224 | 0 | 0 | 0 |   | 224 | 0 | 280 | 0 |
| 0   | 0 | 0 | 0 | → | 0   | 0 | 0   | 0 |
| 280 | 0 | 0 | 0 |   | 0   | 0 | 0   | 0 |
| 0   | 0 | 0 | 0 |   | 0   | 0 | 0   | 0 |

30

45 ° rotation

~~~~~

	224	0	0	0		224	0	140	0
	0	0	0	0	→	0	0	0	0
5	280	0	0	0		140	0	0	0
	0	0	0	0		0	0	0	0

Zoom out

~~~~~

|    |     |   |   |   |   |     |   |   |   |
|----|-----|---|---|---|---|-----|---|---|---|
| 10 | 224 | 0 | 0 | 0 |   | 224 | 0 | 0 | 0 |
|    | 0   | 0 | 0 | 0 | → | 0   | 0 | 0 | 0 |
|    | 280 | 0 | 0 | 0 |   | 0   | 0 | 0 | 0 |
|    | 0   | 0 | 0 | 0 |   | 280 | 0 | 0 | 0 |

15 An example of spatial correlation where scaling would be applicable is in the color representation of 4:2:0 and 4:2:2 formats where the scaling is defined by the color sampling.

20 This approach allows for prediction of scaled or rotated blocks due to the similarity of the pattern that each block is covering. The rotation or scaling will be expressed as a reordering of coefficients in the same block. Due to the rate prediction nature of the process then a perfect match is not necessarily sought, but rather rate reduction which allows for finding the closest match instead of the exact pattern.

25

Super Resolution

30 As shown in Figures 14a and 14b, it is possible when reconstructing an image at a receiver to overlay frames which are offset by a fraction of a pixel from one another in order to achieve a higher resolution. This idea may be referred to as "super-resolution".

Figure 14a illustrates a pixel grid (raster) having some particular resolution defined by the pixel size of the camera. When a frame is captured, the image has a resolution of one pixel value per unit of the grid, i.e. one value per pixel (per statistic required to define a single pixel, e.g. one value of Y, one value of U and one value of V per pixel in YUV colour-space). Say now that the pixel grid is offset right by approximately half a pixel and down by approximately half a pixel, either because the camera moves slightly or because the scene or object being captured moves relative to the camera. If this movement is known or can be estimated, then it is possible to reconstruct a higher resolution image by superimposing the values captured from the two offset grids. In the example of Figure 14a this results in four "super-resolution" pixels A, B, C and D for each actual physical pixel of the camera's sensor array. Each super-resolution pixel value may be determined for example by interpolating between the two overlapping real pixel values which contribute to it.

15

Figure 14b illustrates the idea in more detail. For the sake of illustration, suppose there is a camera with only a 2x2 sensor array of four pixels, and in a first frame at some moment in time an object is captured appearing only in the top two pixels and not appearing in the bottom two. Thus in the first frame the object contributes only to the top two pixel values and not the bottom two. Suppose then that by a second frame at a later moment in time the object has moved down by half a pixel, or the camera has moved up by half a pixel. When the object is now captured in the second frame, different areas of the object appear in all four pixels and so it contributes to all four pixel values.

20

The idea of superimposing fractional shifts in a pixel grid has been used in the past to increase the resolution of satellite images for example. As mentioned, this idea may be referred to as "super-resolution" or sometimes "remote sensing" in the context of satellite images. However, this technique has only been used in the past to increase the resolution beyond the intrinsic physical resolution of the camera or detector in question. For example, some satellite detectors only have one "pixel" with resolution of the order 1km, and rely on this technique to greatly improve on the resolution of the satellite detector.

25

30

However, no-one has previously considered the potential to deliberately transmit a video image with a lower resolution than the intrinsic resolution of the camera, then use a super-resolution scheme to reconstruct an image at the receiver  
5 having a resolution more closely approaching the camera's intrinsic resolution. It is this idea that is the subject of a further aspect of the present invention. The advantage is that the transmitted video stream requires fewer bits per unit time. That is, instead of using fractional shifts of real size pixels to boost resolution beyond the camera's natural resolution, one aspect of the present invention  
10 instead uses the super-resolution technique to transmit averaged values for larger image units each corresponding to multiple real pixels (the averaged units thus having lower resolution than the real camera resolution) and to then reconstruct the real camera resolution at the receiver (or at least a higher resolution than that of the averaged units).

15

An example is discussed in relation to Figure 14c. Here, a region of an image is captured having higher-resolution values A to P:

|    |    |    |    |   |
|----|----|----|----|---|
|    | A, | B, | C, | D |
| 20 | E, | F, | G, | H |
|    | I, | J, | K, | L |
|    | M, | N, | O, | P |

In some embodiments, these higher-resolution values may correspond to the  
25 values captured from individual pixels of the camera's sensor array. In other embodiments however, these higher-resolution values need not necessarily correspond to the actual physical size of the camera's pixels, but rather may represent the smallest size unit that would be used by the encoder in question in some particular mode of operation. The point is that the following encoder will  
30 encode a frame with an even lower resolution, i.e. by averaging or otherwise combining groups of higher-resolution values to create larger, lower-resolution units represented by respective lower-resolution values. In the example of Figure

14c the lower-resolution units are 2x2 groups of higher-resolution values, but it will be appreciated that other schemes could equally well be used.

At an initial frame in a sequence, frame 0, the encoder averages the higher-resolution values F, G, J and K (or otherwise combines them, e.g. by totalling). This average provides a single overall lower-resolution value for a single, larger, lower-resolution unit covering the area of the respective group of four smaller, higher-resolution units. A similar averaging is performed for adjacent groups, thus generating a lower-resolution grid of larger size units represented by respective lower-resolution values. The encoder then encodes and transmits the frame based only on the lower-resolution grid of the averaged values.

Note that in embodiments, the image may still be divided into blocks and/or macroblocks, with each block or macroblock comprising a plurality of lower-resolution units (though fewer than if represented at the higher resolution). In this case, the blocks of multiple lower-resolution units may still be transformed into the frequency domain as part of the encoding process, though the transform may be considered optional according to this particular aspect of the invention. Either way, the super-resolution algorithm operates in the spatial domain (i.e. if there is a frequency domain transform, the super-resolution algorithm occurs before the transform at the encoder and after the reverse transform at the decoder)

At a first subsequent frame in the sequence, frame 1, the encoder shifts the lower-resolution grid up and left by one higher-resolution unit. The encoder then averages the higher-resolution values A, B, E and F to create a single overall lower-resolution value for a single, larger, lower-resolution unit covering the area of the respective group of four smaller, higher-resolution units – so now offset in each of the horizontal and vertical direction by one higher-resolution unit, which means a fractional offset of the lower-resolution grid. Again a similar averaging is performed for adjacent groups, thus generating a lower-resolution grid of larger size units each represented by respective lower-resolution values – but this time including the described offset. The encoder then encodes and transmits the frame using only the offset lower-resolution grid of the averaged values (again with a

transformation of blocks of multiple such units into the frequency domain if appropriate to the embodiment in question).

5 Note that the receiver has now been provided with two lower-resolution units covering the higher resolution unit F by means of the fractional overlap between the lower-resolution units of frames 0 and 1, thus allowing the receiver to generate an individual higher-resolution value for F. However, further subsequent frames 2-5 will be required in order to recreate the full higher-resolution grid.

10 At a second subsequent frame, frame 2, the encoder shifts the lower-resolution grid up and right by one higher-resolution unit relative to that of the initial frame 0. The encoder then averages the group of higher-resolution values C, D, G and H to obtain a respective lower-resolution value for a respective lower-resolution unit, and similarly for surrounding units, thus generating another offset grid of lower-  
15 resolution values which is encoded and transmitted to the receiver. The receiver now has enough information to recreate higher-resolution unit G by means of the fractional overlap between the lower-resolution units of frames 0 and 2.

20 At a third subsequent frame in the sequence, frame 3, the encoder shifts the lower-resolution grid down and right by one higher-resolution unit relative to that of the initial frame 0, and then averages the group of higher-resolution values K, L, O and P to obtain a respective lower-resolution value for another lower-resolution unit. This is encoded and transmitted to the receiver as part of a grid of similarly offset lower-resolution units, now allowing the receiver to recreate higher-  
25 resolution unit K by means of the fractional overlap between the lower-resolution units of frames 0 and 3.

The sequence then continues to a fourth subsequent frame, frame 4, where higher-resolution units I, J, M and N are averaged, and encoded and transmitted  
30 in a lower-resolution grid, thus allowing the receiver to recreate higher-resolution unit J by means of the fractional overlap between the lower-resolution units of frames 0 and 4.

Once the pattern of fractional shifts applied over the sequence of frames 0 to 5 has been completed, the full higher-resolution grid can be reconstructed at the receiver.

5 It will be appreciated however that the above is only one possible example. In other implementations, different ratios of higher to lower resolution unit sizes may be used, and/or other shift patterns may be used. For example, another possible shift pattern requiring only a four-frame cycle would transmit:

10 Av (B,C,F,G)  
Av (E,F,I,J)  
Av (J,K,N,O)  
Av (G,H,K,L)

15 In one embodiment, it is not necessary for there to be actual movement of the camera or object. Instead, the encoder may generate an indication of an artificial shift or pattern of shifts to be applied at the decoder to recreate the higher-resolution. That is to say, the "movement" may only be artificially generated for the sole purpose of reducing transmitted bitrate.

20

Alternatively, the shift may be based on actual movement. In this case, the movement may be detected by using gyro sensors of a mobile terminal in which the camera is housed so as to detect movement of the camera, or by using motion estimation techniques to detect movement of the scene or object being  
25 captured.

Along with the encoded lower-resolution units of frames 0 to 5, the encoder also transmits some side information in the encoded bitstream indicative of the shift or pattern of shifts to be applied as part of this scheme. This indication could take the  
30 form of a separate shift indication for each frame in the sequence; or more preferably in the case of an artificially generated shift the indication may take the form of a single indicator for the whole sequence, referring to a predetermined pattern of shifts to use for that sequence. In the latter case, the predetermined

patterns may be pre-stored at both the encoder at both the transmitter side and the receiver side. For example, the codec may be operable in one or more different super-resolution modes defined by different respective shift patterns and/or resolution ratios (and preferably the codec will also have a more conventional mode not using super-resolution). The nature of the different modes will be understood by both the encoder at the transmitter side and the decoder at the receiver side, and the side information signalled from the transmitter may indicate the mode that has been used at the transmit side by the encoder.

10 The present invention thus uniquely uses a super-resolution scheme to deliberately down-grade the resolution being transmitted in a video stream in order to reduce the bit-rate, and then reconstruct the higher-resolution image again at the receiver. Of course it is not possible to get "free data" – but the idea is to trade bitrate for reconstruction time, since the scheme will require multiple frames in order to reconstruct the higher resolution image at the receiver, thus taking a longer time to obtain the higher resolution than if the data was simply transmitted at the higher resolution in each frame.

For this reason, the above-described feature may not be suited to very fast motion, though it may be useful for encoding motion which is slower but more detailed. In a particularly advantageous embodiment, because the blocks are encoded on a block-by-block basis, it is possible to encode different regions of the same video differently. E.g. based on motion estimation analysis of the video image, a slow moving background may be encoded using lower-resolution units, whilst a faster moving foreground in the same image may be encoded using a higher-resolution; or even vice versa. In this case the shift information may be signalled on a block-by-block basis, where each block comprises multiple lower-resolution units (or on a macroblock-by-macroblock basis, etc.).

30 It is particularly preferred to use this idea in conjunction with the global block list described above. That is, some frames or some blocks or areas within a frame may be encoded using the global block list feature described above, whilst other frames or even other blocks or areas within the same frame may be encoded



using the super-resolution feature described in this section. For example, the global block list could be used to encode blocks in areas that are relatively static whilst the super-resolution could be used to encode blocks in other areas where more detailed motion is occurring (so as to reduce a peak in bitrate that such motion might otherwise cause); or the global block list could be used to encode the faster motion of large objects whilst the super-resolution feature could be used to encode areas where less motion is occurring (because it is less suited to fast motion due to the time required to reconstruct the higher-resolution image).

10 In other embodiments the global block list could alternatively be used to signal reference blocks for encoding and decoding video using a more conventional super-resolution approach for increasing resolution beyond the intrinsic resolution of the camera.

15 In addition to the scaling and rotation, the above allows an implementation of a super resolution approach in compilation of the frequency domain and spatial domain algorithm. In this approach every other frame is moved by  $\frac{1}{2}$  or  $\frac{1}{4}$  of pixel in a specific pattern that can be communicated to the decoder. While encoding in such a way, it is possible to derive a benefit of following minor motion shifts of  $\frac{1}{2}$  or  $\frac{1}{4}$  of a pixel or unit by simply finding matches in the direction of the motion which would otherwise have been missed. Additionally the reconstruction can be done in the spatial domain via pixel re-sampling. Figure 14a shows an example in which the  $\frac{1}{2}$  pixel shift between the two frames allows for four new pixels (A,B,C and D) to be created out of one pixel in the original frame.

25

The shift direction can come from the acquisition system as an encoding input or created to reduce the bitrate as a reverse scalability, e.g. sending CIF resolution instead of VGA.

30

## Implementation

The encoder elements 2, 4, 6, 8 and 10; and the decoder elements 12, 14, 16, 18 and 20 are each preferably implemented in software modules stored on a storage medium such as a hard drive or flash memory and arranged for execution on a processor. However, in other embodiments some or all of these elements could be at least partially implemented in dedicated hardwired circuitry.

In one possible embodiment, the invention may be implemented as an add-on to an existing encoder such as ISO standard H.264. That is, the input to the quantizer 4 in Figure 4 will be an output from a standard encoder such as an H.264 encoder.

It will be appreciated that the above embodiments have been described only by way of example.

For instance, note that whilst the term "block" is used herein, in its most general sense this is not intended to imply and particular size, shape or level subdivision. It will be appreciated that in different schemes then various different divisions and subdivisions may be referred to by terms such as macroblock, block and sub-block, etc., but that the term "block" as used most generally herein may correspond to any of these or indeed any other constituent image portion being a division of a video frame.

Further, whilst the above has been described with reference to the example of a Discrete Cosine Transform into the frequency domain, it will be appreciated that other transforms such as the KLT or others can be used (some of which may not represent the transform domain in terms of spatial frequency coefficients but in terms of some other transform domain coefficients).

30

Further, whilst the above has been described in terms of a residual representing the subtracted difference between the coefficients of the target block and the coefficients of the reference block, this is not the only possibility for encoding the

coefficients or values of the target block relative to those of the reference block. In other possible embodiments for example, the difference may be represented and signalled in terms of parameters of a correlation between the target block and the reference block such that the target can be predicted from the correlation, or in  
5 terms of coefficients of a filter that may be applied to the reference block to predict the target block. In these cases the prediction may not necessarily be lossless as in the case of a subtractive difference, but may instead be lossy such that the difference does not represent the exact difference. The term "difference" as used herein is not limited to subtractive difference nor to an exact difference.

10

Further, the present invention is not limited to implementation in any particular standard nor as an add-on to any particular standard, and may be implemented either as a new stand-alone codec, an add-on to an existing codec, or as a modification to an existing codec.

15

Other variants may be apparent to a person skilled in the art given the disclosure herein. The invention is not limited by the described embodiments, but only by the appendant claims.

CLAIMS:

1. A method of compressing a video signal comprising a plurality of video frames, each frame being divided into a plurality of image portions, comprising:
  - for each of a plurality of the image portions, transforming the image portion from a spatial domain representation into a transform domain representation so as to generate respective set of transform domain coefficients; and
  - for each of a plurality of target ones of said image portions to be encoded, determining a respective reference portion represented by a respective set of transform domain coefficients, determining a prediction of the target image portion based on the reference portion, and encoding the target image portion based on the prediction;
    - wherein the prediction comprises predicting a non-lateral motion of the target image portion relative to the reference portion based on processing of their transform domain coefficients; and
    - wherein the encoding comprises generating an encoded bitstream comprising information about the non-lateral motion for use by a decoder.
2. The method of claim 1, wherein the prediction of said non-lateral motion comprises predicting a scaling of the target image portion relative to the reference portion based on processing of their transform domain coefficients.
3. The method of claim 2, wherein said information indicates a parameter of the scaling.
4. The method of claim 1, 2 or 3, wherein the prediction of said non-lateral motion comprises predicting a rotation of the target image portion relative to the reference portion based on processing of their transform domain coefficients.
5. The method of claim 4, wherein said information indicates a degree of the rotation.
6. The method of any preceding claim, wherein the determination of the reference portion comprises selecting a selected reference portion from amongst a

plurality of potential reference portions each being represented by a respective set of transform domain coefficients, and the prediction of the target image portion is based on the selected reference portion.

7. The method of claim 6, wherein the plurality of potential reference portions comprises at least one potential reference portion for predicting a lateral motion of the target image portion according to one type of motion prediction, and at least one potential reference portion for predicting a non-lateral motion of the target image portion according to another type of motion prediction; and said information indicates the type of motion prediction to be used by a decoder for decoding the target image portion.

8. The method of claim 7, wherein the plurality of potential reference portions comprises at least one potential reference portion for predicting a lateral motion of the target image portion according to a first type of motion prediction, at least one potential reference portion for predicting a scaling of the target image portion according to a second type of motion prediction, and at least one potential reference portion for predicting a rotation of the target image portion according to a third type of motion prediction; and said information indicates whether the first, second or third type of motion prediction is to be used by a decoder for decoding the target image portion.

9. The method of any of claims 6, 7 or 8, comprising determining a respective index value for the target image portion and each of a plurality of the reference portions based on the respective transform domain coefficients; wherein the selection of the selected reference portion is based on the index values.

10. The method of claim 9, wherein the selection of the selected reference portion comprises selecting a subset of candidate portions having an index value within a predetermined range of the target image portion, and selecting the respective reference portion from amongst the candidate portions.

11. The method of claim 9 or 10, wherein the index value represents an energy of the respective portion.

12. The method of claim 11, wherein the index value is based on one of: a number of non-zero transform domain coefficients in the respective portion, and a number of transform domain coefficients having a value of zero in the respective portion.
13. The method of claim 11, wherein the index value is based on an aggregate combination of the transform domain coefficients of the respective portion.
14. The method of claim 13, wherein the index value is based on an average or total of a modulus of the transform domain coefficients of the respective portion.
15. The method of any preceding claim, wherein the selection of the selected reference portion is based on a determination of a number of bits that would be required in the encoded bitstream to encode the target image portion.
16. The method of claim 15, wherein the selection of the selected reference portion from amongst the candidate portions is based on a determination of a number of bits that would be required in the encoded bitstream to encode the target image portion relative to each of the candidate portions.
17. The method of any preceding claim, comprising:
  - during ongoing encoding of the video signal, updating a table indicating a subset of commonly usable reference portions; and
  - including the table in the encoded bitstream along with side information identifying the selected reference portions by reference to an entry in said table.
18. The method of claim 17, wherein one or more of said subset of reference portions are artificial portions rather than actual portions from a frame.
19. The method of claim 18, comprising determining the one or more artificial reference portions by averaging, interpolating or extrapolating from one or more actual image portions from one or more of said frames.

20. The method of claim 17, 18 or 19, wherein one or more of said subset of reference portions are actual image portions from one or more of said frames.
21. The method of any preceding claim, wherein the transform domain coefficients represent chrominance and luminance in the transform domain.
22. The method of any preceding claim, wherein each image portion comprises a block or macroblock corresponding to a two-dimensional array of pixels.
23. An encoder for compressing video data, the encoder comprising:  
an input for receiving a video signal comprising a plurality of video frames, each frame being divided into a plurality of image portions; and  
processing apparatus configured, for each of a plurality of the image portions, to transform the image portion from a spatial domain representation into a transform domain representation so as to generate respective set of transform domain coefficients; and for each of a plurality of target ones of said image portions to be encoded, to determine a respective reference portion represented by a respective set of transform domain coefficients, determine a prediction of the target image portion based on the reference portion, and encode the target image portion based on the prediction;  
wherein the processing apparatus is operable to perform said prediction by predicting a non-lateral motion of the target image portion relative to the reference portion based on processing of their transform domain coefficients; such that the encoding comprises generating an encoded bitstream comprising information about the non-lateral motion for use by a decoder; and  
wherein the encoder further comprises an output module arranged to output the encoded bitstream.
24. The encoder of claim 23, wherein the processing apparatus is further configured to generate an encoded bitstream in accordance with any of claims 2 to 22.

25. An encoder program product for compressing video data, the encoder program product comprising software embodied on a computer-readable medium and configured so as when executed to:

for each of a plurality of image portions from said stream, transform the image portion from a spatial domain representation into a transform domain representation so as to generate respective set of transform domain coefficients; and

for each of a plurality of target ones of said image portions to be encoded, determine a respective reference portion represented by a respective set of transform domain coefficients, determine a prediction of the target image portion based on the reference portion, and encode the target image portion based on the prediction;

wherein the prediction comprises predicting a non-lateral motion of the target image portion relative to the reference portion based on processing of their transform domain coefficients; and

wherein the encoding comprises generating an encoded bitstream comprising information about the non-lateral motion for use by a decoder.

26. The encoder program product of claim 25, wherein the software is further configured to generate an encoded bitstream in accordance with any of claims 2 to 22.

27. An encoded bit stream of compressed video data, representing a plurality of successive video frames with each frame being divided into a plurality of image portions and each portion corresponding to a plurality of pixels, the encoded bit stream comprising:

a set of quantized transform domain coefficients representing a first one of said portions; and

information relating to a non-lateral motion of a second one of said portions which is to be predicted from the first portion, allowing the second portion to be predicted based on the non-lateral motion and on the transform domain coefficients of the first portion.



28. The encoded bitstream of claim 27, encoded according to any of claims 1 to 22.
29. A network equipment comprising a transmission medium carrying a bitstream according to claim 27 or 28.
30. A method of decoding the encoded bitstream of claim 27 or 28, comprising receiving the encoded bitstream, and predicting the second portion based on the information relating to the non-lateral motion and on the transform domain coefficients of the first portion.
31. A decoder for decoding the encoded bitstream of claim 27 or 28, comprising an input for receiving the encoded bitstream, and processing apparatus configured to predict the second portion based on the information relating to the non-lateral motion and on the transform domain coefficients of the first portion.
32. A decoder program product for decoding the encoded bitstream of claim 27 or 28, the decoder program product comprising software embodied on a computer-readable medium and configured so as when executed to: receive the encoded bitstream, and predict the second portion based on the information relating to the non-lateral motion and on the transform domain coefficients of the first portion.

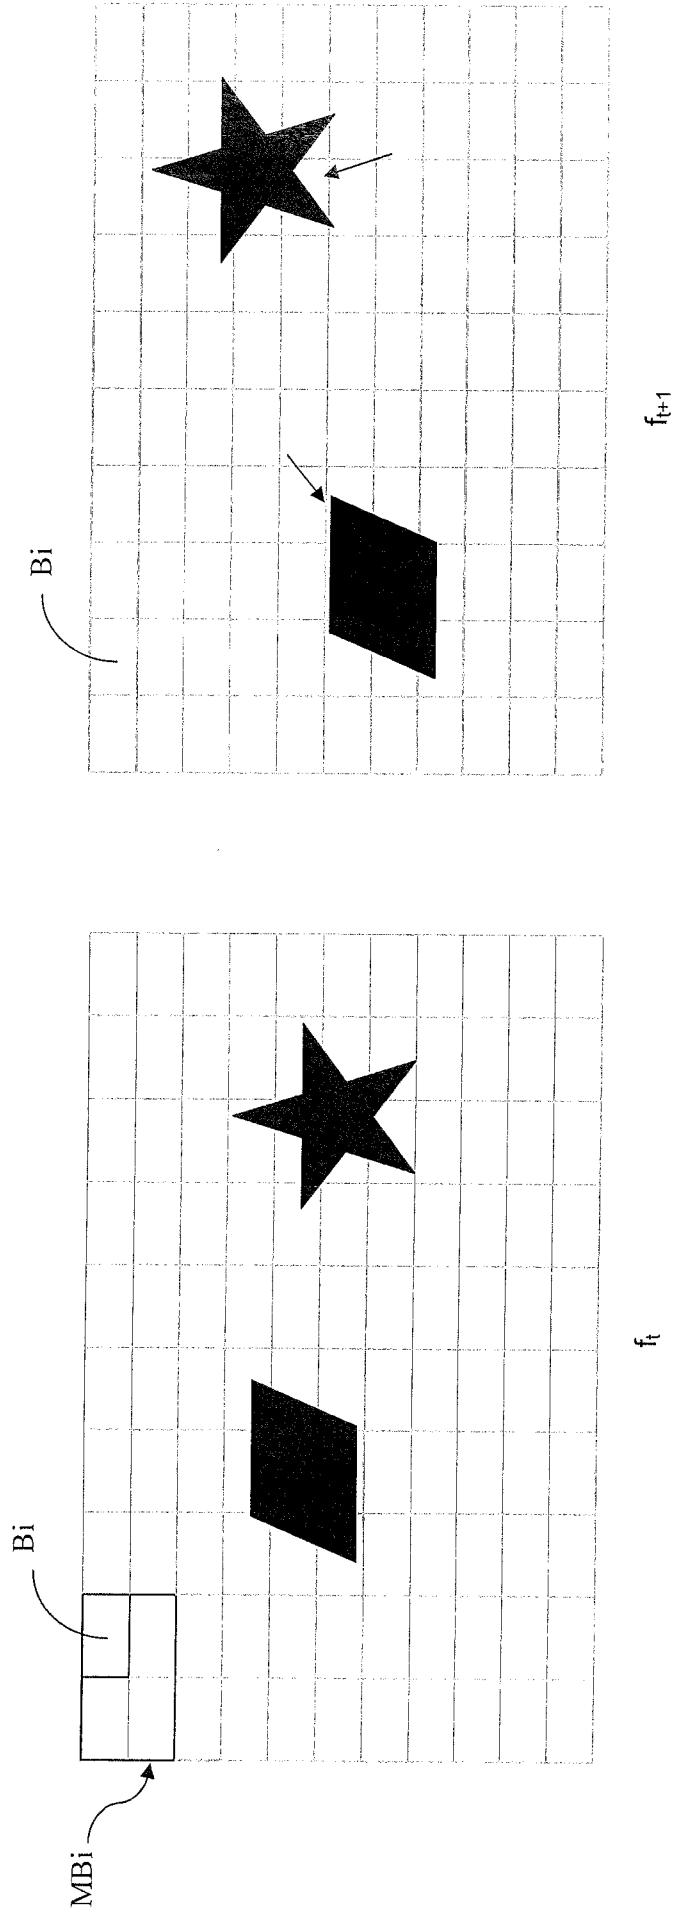


Fig. 1

+-----+  
| 130, 129, 129, 127, 125, 124, 124, 127, | 129, 126, 123, 117, 109, 104, 101, 99, |  
| 128, 128, 127, 126, 126, 124, 125, 126, | 128, 125, 122, 118, 110, 105, 102, 100, |  
| 126, 126, 126, 126, 125, 126, 126, | 127, 124, 122, 119, 114, 108, 105, 103, |  
| 125, 125, 125, 125, 125, 125, 127, | 126, 125, 123, 120, 115, 111, 107, 105, |  
| 124, 125, 125, 125, 125, 125, 126, | 125, 124, 123, 122, 119, 115, 112, 109, |  
| 124, 125, 125, 125, 125, 126, | 125, 125, 124, 123, 120, 117, 114, 111, |  
| 122, 124, 124, 124, 124, 124, 125, | 126, 126, 125, 124, 121, 119, 116, 114, |  
| 120, 122, 124, 124, 124, 124, 124, | 126, 126, 126, 125, 123, 121, 118, 115, |  
+-----+  
| 116, 121, 122, 123, 123, 123, 125, 127, | 127, 127, 126, 125, 123, 120, 118, 115, |  
| 114, 120, 122, 121, 121, 121, 123, 125, | 125, 126, 126, 124, 119, 118, 116, 115, |  
| 112, 118, 120, 119, 119, 120, 121, | 122, 121, 120, 117, 115, 114, 114, |  
| 111, 116, 118, 118, 118, 119, 120, | 119, 117, 115, 113, 112, 112, 112, |  
| 110, 114, 115, 116, 116, 116, 116, | 115, 113, 111, 109, 109, 109, 109, |  
| 112, 114, 114, 115, 115, 114, 114, | 113, 111, 109, 108, 108, 107, 108, |  
| 113, 113, 113, 113, 113, 112, 112, | 111, 110, 108, 106, 106, 106, 109, |  
| 114, 114, 114, 114, 114, 113, 113, | 113, 111, 110, 109, 109, 109, 101, |  
+-----+

FIG. 1a

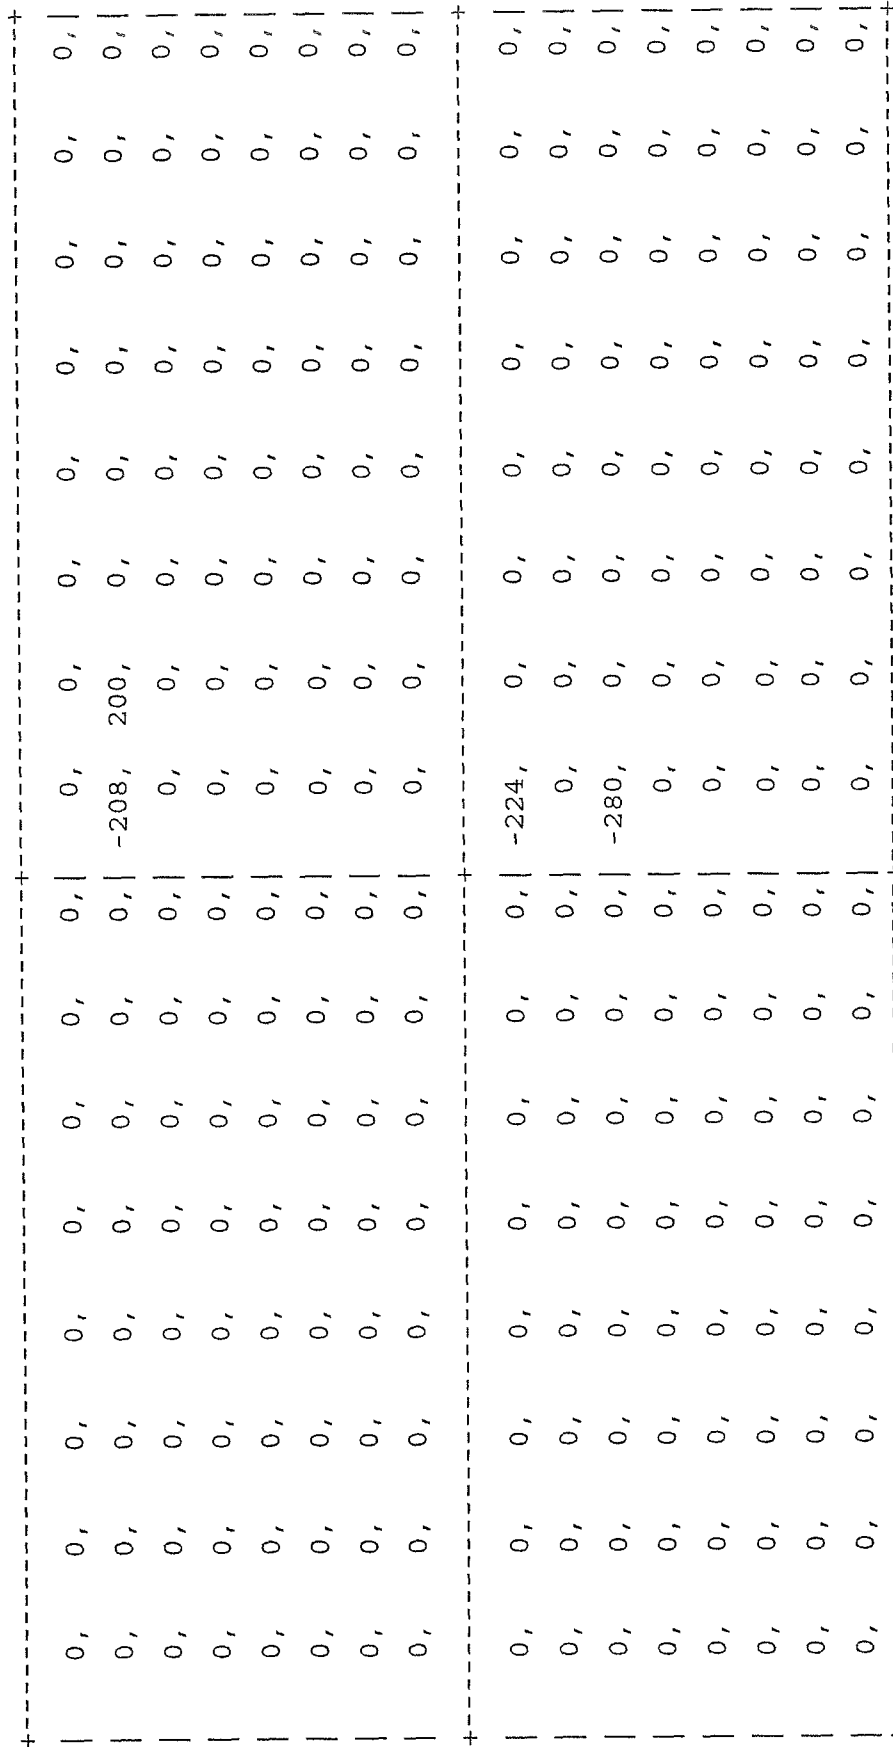


FIG. 1b

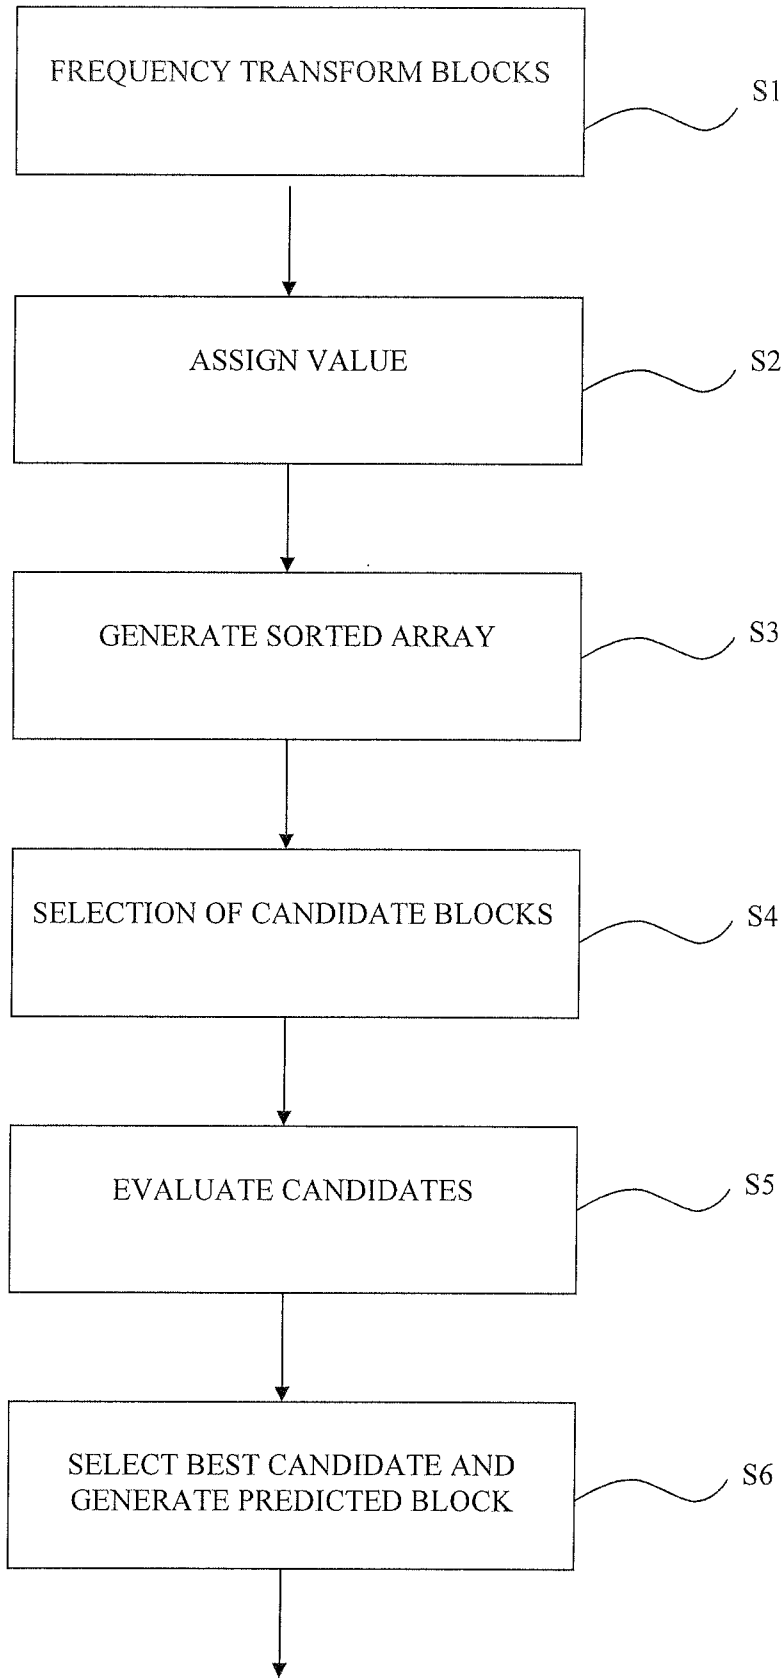


Fig. 2

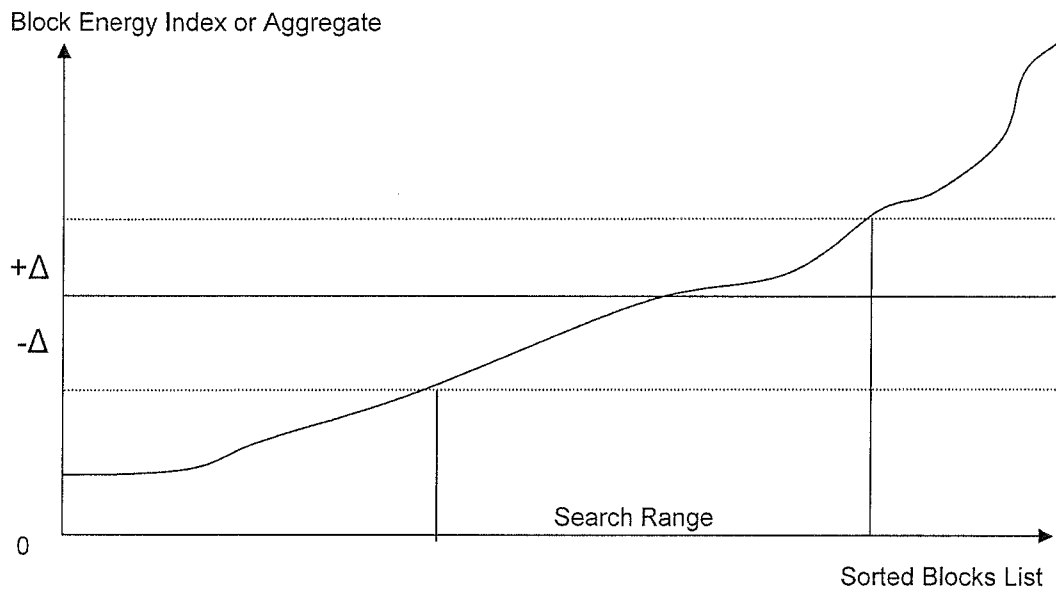


Fig. 3

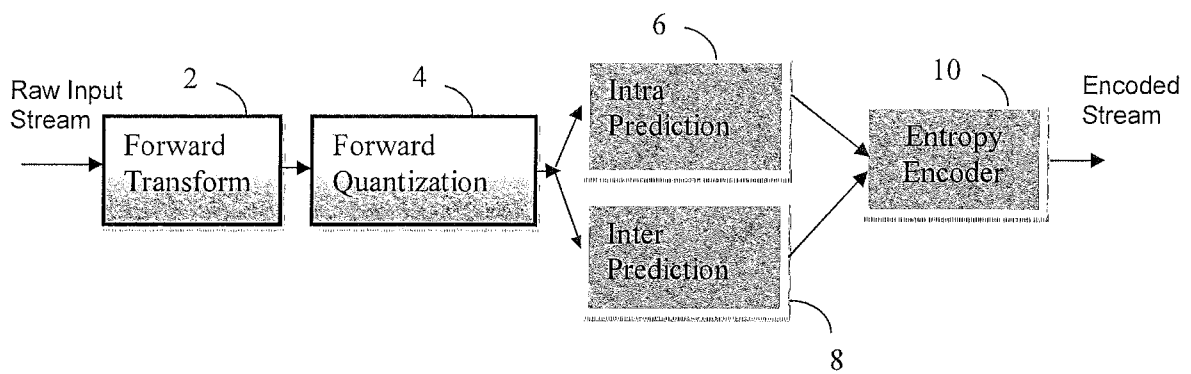


Fig. 4

Initial Blocks

|          |   |   |   |   |   |   |   |   |      |   |   |   |   |   |   |   |
|----------|---|---|---|---|---|---|---|---|------|---|---|---|---|---|---|---|
| -320 304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -320 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 400  | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 304  | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -640 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 400  | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

|          |   |   |   |   |   |   |   |   |      |   |   |   |   |   |   |   |
|----------|---|---|---|---|---|---|---|---|------|---|---|---|---|---|---|---|
| -320 304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -640 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -400 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 304  | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -320 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

|   |   |   |   |   |   |   |   |   |         |   |   |   |   |   |   |   |
|---|---|---|---|---|---|---|---|---|---------|---|---|---|---|---|---|---|
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 320-304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -304    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0       | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0       | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0       | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0       | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0       | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0       | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

|      |   |   |   |   |   |   |   |   |          |   |   |   |   |   |   |   |
|------|---|---|---|---|---|---|---|---|----------|---|---|---|---|---|---|---|
| 320  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -320-304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| -304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 400      | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0        | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

FIG. 3a

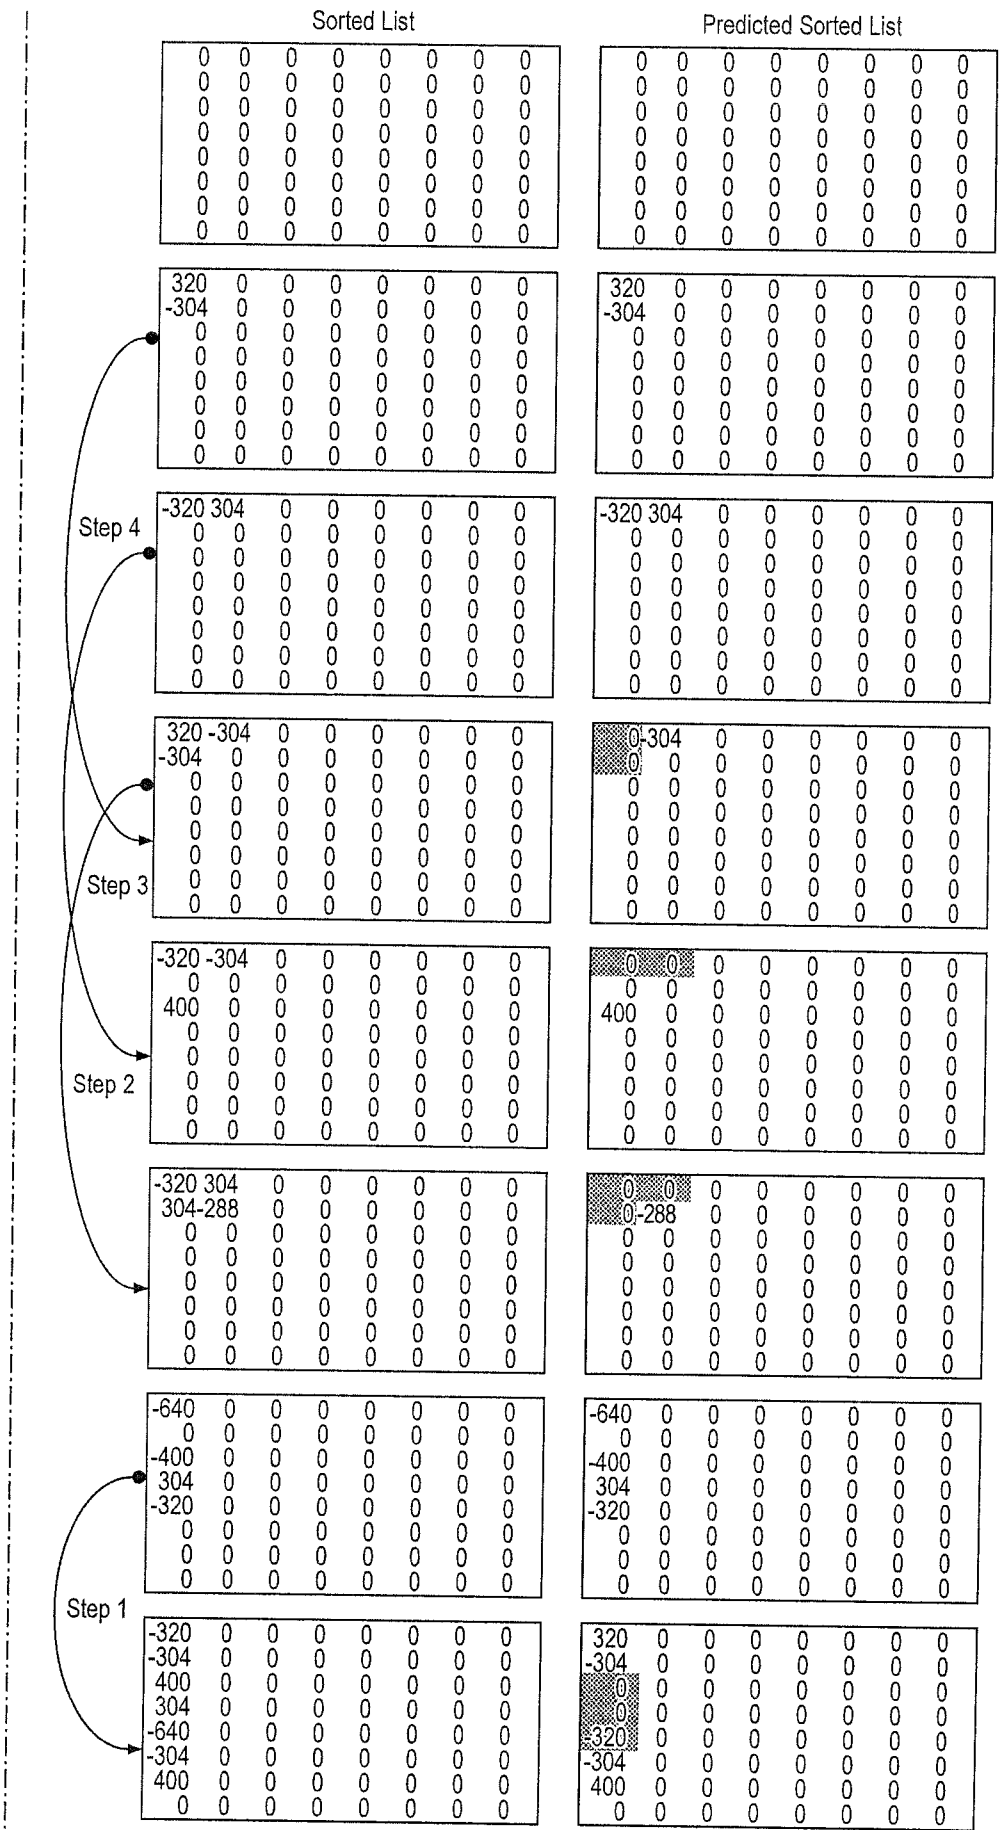


FIG. 3a Cont'd



Initial Blocks

|      |     |   |   |   |   |   |   |   |      |   |   |   |   |   |   |   |   |
|------|-----|---|---|---|---|---|---|---|------|---|---|---|---|---|---|---|---|
| -320 | 304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -320 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0   | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0   | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 400  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0   | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 304  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0   | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -640 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0   | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0   | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 400  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0   | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

|      |      |   |   |   |   |   |   |   |      |   |   |   |   |   |   |   |   |
|------|------|---|---|---|---|---|---|---|------|---|---|---|---|---|---|---|---|
| -320 | 304  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -640 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 304  | -288 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -400 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 304  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -320 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

|   |   |   |   |   |   |   |   |   |      |      |   |   |   |   |   |   |   |
|---|---|---|---|---|---|---|---|---|------|------|---|---|---|---|---|---|---|
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 320  | -304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -304 | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

|      |   |   |   |   |   |   |   |   |      |      |   |   |   |   |   |   |   |
|------|---|---|---|---|---|---|---|---|------|------|---|---|---|---|---|---|---|
| 320  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | -320 | -304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| -304 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 400  | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0    | 0    | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

|     |      |     |      |     |      |     |      |      |  |  |  |  |  |  |  |  |  |
|-----|------|-----|------|-----|------|-----|------|------|--|--|--|--|--|--|--|--|--|
| 1   | 2    | 3   | 4    | 5   | 6    | 7   | 8    | 9    |  |  |  |  |  |  |  |  |  |
| 0   | -288 | 304 | -304 | 320 | -320 | 400 | -400 | -640 |  |  |  |  |  |  |  |  |  |
| 487 | 1    | 5   | 6    | 2   | 5    | 3   | 1    | 2    |  |  |  |  |  |  |  |  |  |

Bins Before Prediction  
Values in each Bin  
Frequency

|     |      |     |      |     |      |     |      |      |  |  |  |  |  |  |  |  |  |
|-----|------|-----|------|-----|------|-----|------|------|--|--|--|--|--|--|--|--|--|
| 1   | 2    | 3   | 4    | 5   | 6    | 7   | 8    | 9    |  |  |  |  |  |  |  |  |  |
| 0   | -288 | 304 | -304 | 320 | -320 | 400 | -400 | -640 |  |  |  |  |  |  |  |  |  |
| 498 | 1    | 2   | 3    | 1   | 3    | 2   | 1    | 1    |  |  |  |  |  |  |  |  |  |

#Bins After Prediction  
Values in each Bin  
Frequency

|     |       |       |       |      |      |      |      |      |            |  |  |  |  |  |  |  |  |
|-----|-------|-------|-------|------|------|------|------|------|------------|--|--|--|--|--|--|--|--|
| 1   | 2     | 3     | 4     | 5    | 6    | 7    | 8    | 9    |            |  |  |  |  |  |  |  |  |
| 0   | -304  | 304   | -320  | 400  | 320  | -640 | -288 | -400 |            |  |  |  |  |  |  |  |  |
| 487 | 6     | 5     | 5     | 3    | 2    | 2    | 1    | 1    |            |  |  |  |  |  |  |  |  |
|     | 2.06  | 2.32  | 2.32  | 3.06 | 3.64 | 3.64 | 4.64 | 4.64 |            |  |  |  |  |  |  |  |  |
|     | 12.35 | 11.61 | 11.61 | 9.18 | 7.29 | 7.29 | 4.64 | 4.64 | 68.61 bits |  |  |  |  |  |  |  |  |

Values in each Bin  
Frequency  
Total Bits Before Prediction

|     |      |      |      |      |      |      |      |      |            |  |  |  |  |  |  |  |  |
|-----|------|------|------|------|------|------|------|------|------------|--|--|--|--|--|--|--|--|
| 1   | 2    | 3    | 4    | 5    | 6    | 7    | 8    | 9    |            |  |  |  |  |  |  |  |  |
| 0   | -320 | -304 | 400  | 304  | -288 | 320  | -400 | -640 |            |  |  |  |  |  |  |  |  |
| 498 | 3    | 3    | 2    | 2    | 1    | 1    | 1    | 1    |            |  |  |  |  |  |  |  |  |
|     | 2.22 | 2.22 | 2.81 | 2.81 | 3.81 | 3.81 | 3.81 | 3.81 |            |  |  |  |  |  |  |  |  |
|     | 6.67 | 6.67 | 5.61 | 5.61 | 3.81 | 3.81 | 3.81 | 3.81 | 39.79 bits |  |  |  |  |  |  |  |  |

Frequency  
Total Bits Before Prediction

##### Savings

FIG. 3b

28.82 Total Bits Saved

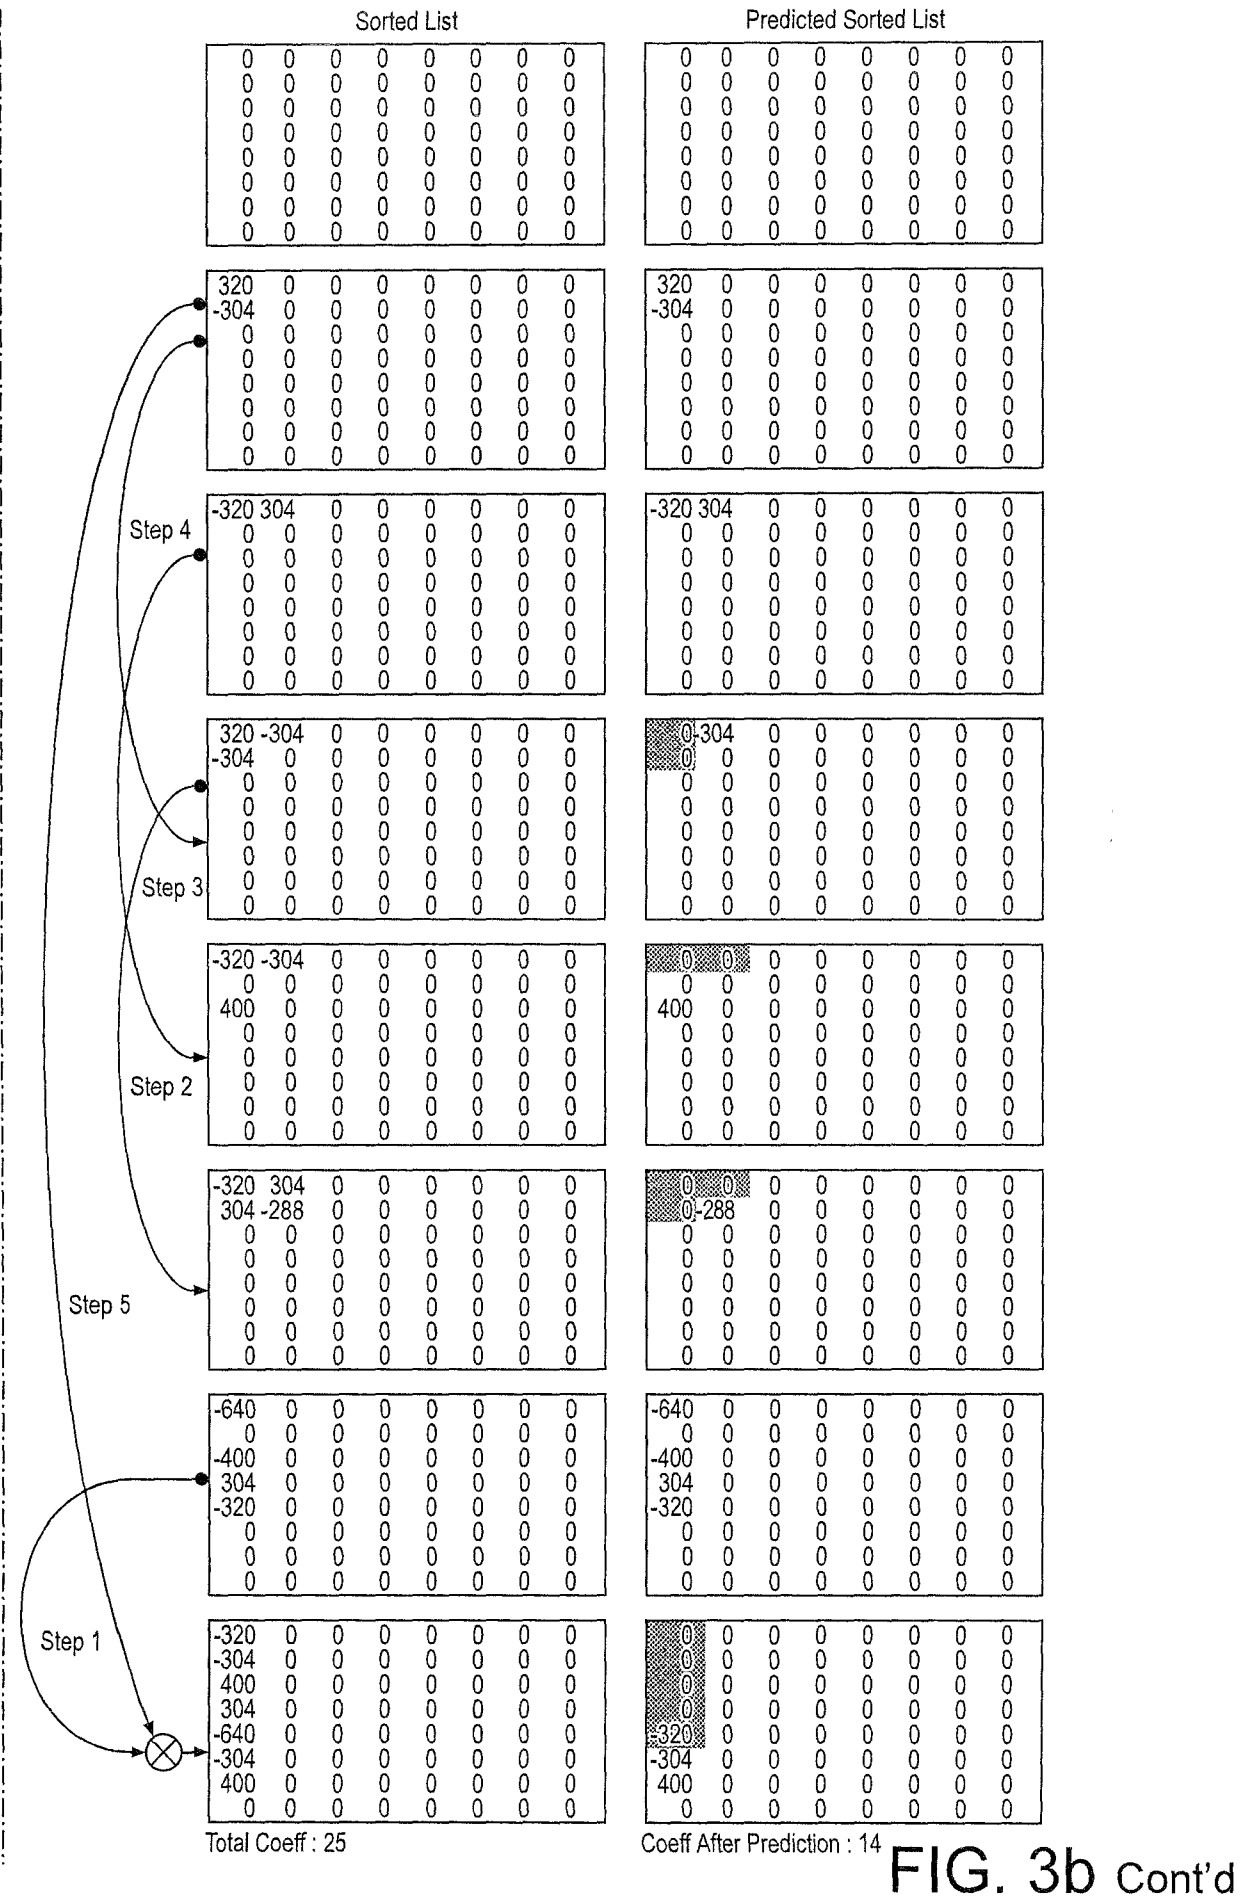


FIG. 3b Cont'd

|     |       |       |       |       |       |      |      |      |       |       |      |      |      |       |      |   |
|-----|-------|-------|-------|-------|-------|------|------|------|-------|-------|------|------|------|-------|------|---|
|     | b0    |       |       |       | b1    |      |      |      | b2    |       |      |      | b3   |       |      |   |
| b4  | 2464  | -4032 | -2688 | -1728 | 224   | 576  | -224 | 576  | 2688  | 0     | 224  | -576 | 896  | -576  | -672 | 0 |
|     | 864   | -1104 | 288   | 0     | 0     | -368 | 288  | -368 | 1728  | -1840 | 0    | 0    | 864  | -1104 | 288  | 0 |
|     | 672   | -576  | -448  | 288   | 224   | -576 | 0    | -288 | 448   | -288  | -448 | 288  | 672  | 0     | 224  | 0 |
|     | 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    | -288  | 368   | 0    | 0    | -288 | 0     | 0    | 0 |
|     | -1344 | -1440 | 1792  | 864   | 448   | 288  | -224 | 0    | 224   | 0     | 0    | 0    | -224 | 0     | 0    | 0 |
|     | 0     | 1104  | -576  | -368  | 576   | 736  | -288 | 0    | 288   | 0     | 0    | 0    | 0    | 0     | 0    | 0 |
|     | 448   | 0     | -448  | 0     | 224   | 0    | -224 | 0    | 224   | 0     | 0    | 0    | 0    | 0     | 0    | 0 |
|     | 288   | 0     | 0     | 0     | 0     | -368 | 0    | 0    | 0     | 0     | 0    | 0    | 0    | 0     | 0    | 0 |
|     | 224   | 0     | 0     | 0     | 0     | 0    | 0    | 0    | -672  | 576   | 0    | 0    | 448  | 288   | -224 | 0 |
| b8  | 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    | 864   | -736  | 0    | 0    | -576 | -368  | 288  | 0 |
|     | 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    | -448  | 576   | 0    | 0    | 448  | 288   | 0    | 0 |
|     | 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    | 288   | -368  | 0    | 0    | -288 | 0     | 0    | 0 |
| b12 | -672  | 864   | -224  | -228  | 0     | 0    | 0    | 0    | -2016 | 0     | -672 | 0    | 448  | -288  | -224 | 0 |
|     | 864   | -1104 | 288   | 0     | 1152  | -736 | 288  | 0    | 0     | 1104  | 0    | 736  | -288 | 368   | 0    | 0 |
|     | -448  | 288   | 0     | 0     | 0     | 0    | -224 | 0    | 448   | 0     | 0    | 0    | 672  | 576   | 0    | 0 |
|     | 288   | 0     | 0     | 0     | -288  | -368 | 0    | 0    | -288  | 736   | 0    | 0    | 0    | 0     | 0    | 0 |
| b16 | 336   | 0     | 0     | 0     | 336   | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | 288   | 0     | 0     | 0     | 576   | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | -224  | 0     | 0     | 0     | -224  | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | 0     | 0     | 0     | 0     | -288  | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| b18 | -336  | 288   | 0     | 0     | -784  | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | 288   | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| b20 | -1008 | 288   | 224   | 288   | -1232 | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | -864  | 368   | 0     | 0     | -1152 | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | 448   | -288  | 0     | 0     | 672   | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | 576   | 0     | 0     | 0     | 576   | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| b22 | 112   | 0     | 0     | 0     | 336   | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
|     | 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |

9 coeff less out of 132  
 bin 122 is removed  
 bin 1008 is removed  
 bin 224 increased  
 nib 288 increased

FIG. 5a

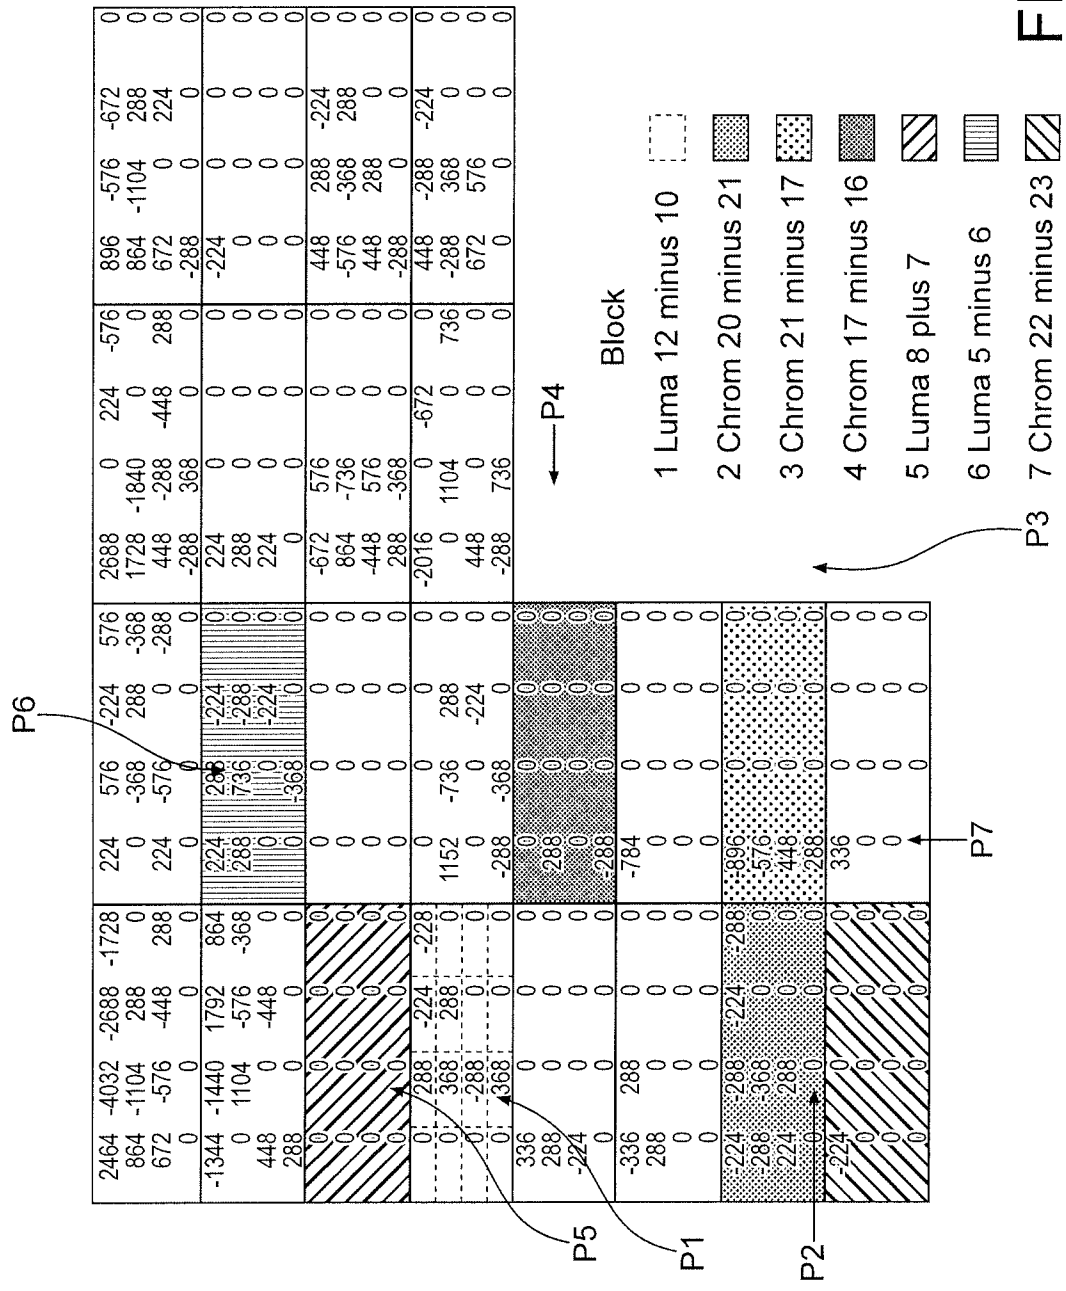


FIG. 5a Cont'd

Frame 1  
MB 1

|      |      |      |      |      |      |      |     |      |      |      |   |      |      |   |   |
|------|------|------|------|------|------|------|-----|------|------|------|---|------|------|---|---|
| -448 | 576  | 0    | 0    | -448 | -288 | 224  | 0   | 0    | 0    | 0    | 0 | 224  | -576 | 0 | 0 |
| 288  | 0    | 288  | 0    | 288  | 0    | -288 | 0   | 0    | 0    | 0    | 0 | 0    | 0    | 0 | 0 |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   | 0    | 0    | 0    | 0 | 0    | 0    | 0 | 0 |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   | 0    | 0    | 0    | 0 | 0    | 0    | 0 | 0 |
| 0    | -864 | 0    | 288  | -224 | 288  | 0    | 0   | 0    | 0    | 0    | 0 | 0    | 0    | 0 | 0 |
| -288 | 736  | -576 | 0    | 0    | 0    | 0    | 0   | 0    | 0    | 0    | 0 | 0    | 0    | 0 | 0 |
| 0    | 0    | 224  | -288 | 0    | 0    | 0    | 0   | 0    | 0    | 0    | 0 | 0    | 0    | 0 | 0 |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   | 0    | 0    | 0    | 0 | 0    | 0    | 0 | 0 |
| -672 | 288  | 0    | 0    | 672  | 0    | -448 | 0   | -448 | -288 | 224  | 0 | -224 | 0    | 0 | 0 |
| 0    | 0    | 0    | 0    | -288 | 736  | 576  | 0   | 288  | -368 | 0    | 0 | 0    | 0    | 0 | 0 |
| 0    | 0    | 0    | 0    | 0    | -288 | 0    | 288 | 0    | 0    | 0    | 0 | 0    | 0    | 0 | 0 |
| 0    | 0    | 0    | 0    | 0    | 0    | -288 | 0   | 0    | 0    | 0    | 0 | 0    | 0    | 0 | 0 |
| 224  | 0    | 0    | 0    | -224 | -576 | 0    | 0   | -224 | -288 | 224  | 0 | -224 | 0    | 0 | 0 |
| 0    | 0    | 0    | 0    | 576  | -368 | -288 | 0   | -288 | 736  | -288 | 0 | 288  | 0    | 0 | 0 |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   | -448 | -288 | 224  | 0 | 0    | 0    | 0 | 0 |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   | 288  | -368 | -288 | 0 | 0    | 0    | 0 | 0 |
| 224  | 288  | 224  | 0    | 0    | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| -228 | -368 | 0    | 0    | 0    | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| 672  | 288  | -224 | 0    | 0    | 288  | 0    | 0   |      |      |      |   |      |      |   |   |
| -228 | 368  | 0    | 0    | -288 | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| 336  | 576  | 224  | 0    | 112  | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| -288 | -368 | 0    | 0    | 0    | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| 2128 | 0    | -224 | 0    | 1008 | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| 0    | 736  | 0    | 0    | -288 | -368 | 0    | 0   |      |      |      |   |      |      |   |   |
| -224 | 0    | 0    | 0    | 0    | 0    | 0    | 0   |      |      |      |   |      |      |   |   |
| 0    | 0    | 0    | 0    | 0    | 0    | 0    | 0   |      |      |      |   |      |      |   |   |

Frame 2  
MB 0

|       |       |       |       |       |      |      |      |       |       |      |      |      |       |      |   |
|-------|-------|-------|-------|-------|------|------|------|-------|-------|------|------|------|-------|------|---|
| 2464  | -4032 | -2688 | -1728 | 224   | 576  | -224 | 576  | 2688  | 0     | 224  | -576 | 896  | -576  | -672 | 0 |
| 864   | -1104 | 288   | 0     | 0     | -368 | 288  | -368 | 1728  | -1840 | 0    | 0    | 864  | -1104 | 288  | 0 |
| 672   | -576  | -448  | 288   | 224   | -576 | 0    | -288 | 448   | -288  | 448  | 288  | 672  | 0     | 224  | 0 |
| 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    | -288  | 368   | 0    | 0    | -288 | 0     | 0    | 0 |
| -1344 | -1440 | 1792  | 864   | 448   | 288  | -224 | 0    | 224   | 0     | 0    | 0    | -224 | 0     | 0    | 0 |
| 0     | 1104  | -576  | -368  | 576   | 736  | -288 | 0    | 288   | 0     | 0    | 0    | 0    | 0     | 0    | 0 |
| 448   | 0     | -448  | 0     | 224   | 0    | -224 | 0    | 224   | 0     | 0    | 0    | 0    | 0     | 0    | 0 |
| 288   | 0     | 0     | 0     | 0     | -368 | 0    | 0    | 0     | 0     | 0    | 0    | 0    | 0     | 0    | 0 |
| 224   | 0     | 0     | 0     | 0     | 0    | 0    | 0    | -672  | 576   | 0    | 0    | 448  | 288   | -224 | 0 |
| 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    | 864   | -736  | 0    | 0    | -576 | -368  | 288  | 0 |
| 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    | -448  | 576   | 0    | 0    | 448  | 288   | 0    | 0 |
| 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    | 288   | -368  | 0    | 0    | -288 | 0     | 0    | 0 |
| -672  | 864   | -224  | -288  | 0     | 0    | 0    | 0    | -2016 | 0     | -672 | 0    | 448  | -288  | -224 | 0 |
| 864   | -1104 | 288   | 0     | 1152  | -736 | 288  | 0    | 0     | 1104  | 0    | 736  | -288 | 368   | 0    | 0 |
| -448  | 288   | 0     | 0     | 0     | 0    | -224 | 0    | 448   | 0     | 0    | 0    | 672  | 576   | 0    | 0 |
| 288   | 0     | 0     | 0     | -288  | -368 | 0    | 0    | -288  | 736   | 0    | 0    | 0    | 0     | 0    | 0 |
| 336   | 0     | 0     | 0     | 336   | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| 288   | 0     | 0     | 0     | 576   | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| -224  | 0     | 0     | 0     | -224  | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| 0     | 0     | 0     | 0     | -288  | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| -336  | 288   | 0     | 0     | -784  | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| 288   | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| -1008 | 288   | 224   | 288   | -1232 | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| -864  | 368   | 0     | 0     | -1152 | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| 448   | -288  | 0     | 0     | 672   | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| 576   | 0     | 0     | 0     | 576   | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| 112   | 0     | 0     | 0     | 336   | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |
| 0     | 0     | 0     | 0     | 0     | 0    | 0    | 0    |       |       |      |      |      |       |      |   |

FIG. 5b

Transmitted Coeff for Fr1 MB 1 Block 14

|      |      |      |   |              |
|------|------|------|---|--------------|
| 0    | -288 | 224  | 0 | 2 Coeff Less |
| 0    | 736  | -288 | 0 |              |
| -224 | -288 | 224  | 0 |              |
| 288  | -368 | -288 | 0 |              |

Transmitted Coeff for Fr1 MB 1 Block 20

|   |      |     |   |              |
|---|------|-----|---|--------------|
| 0 | 864  | 224 | 0 | 2 Coeff Less |
| 0 | -368 | 0   | 0 |              |
| 0 | 0    | 0   | 0 |              |
| 0 | 0    | 0   | 0 |              |

Transmitted Coeff for Fr2 MB 0 Block 0

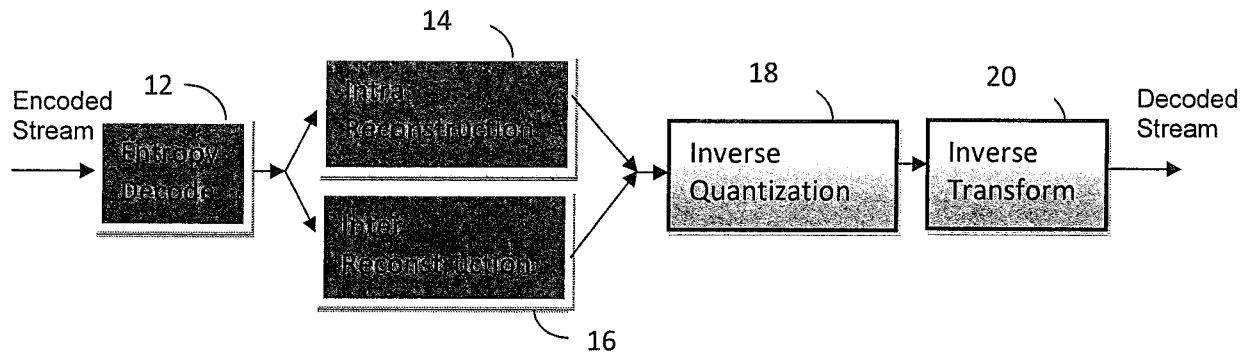
|      |       |       |       |              |
|------|-------|-------|-------|--------------|
| 2464 | -4896 | -2688 | -1440 | 1 Coeff Less |
| 576  | -368  | -288  | 0     |              |
| 672  | -576  | -224  | 0     |              |
| 0    | 0     | 0     | 0     |              |

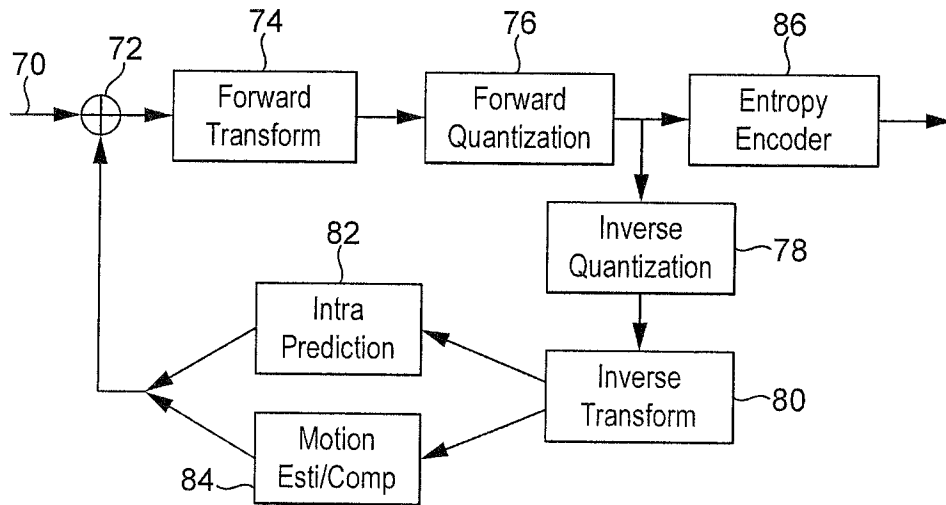
Transmitted Coeff for Fr2 MB 0 Block 11

|      |     |      |   |              |
|------|-----|------|---|--------------|
| 224  | 0   | -448 | 0 | 1 Coeff Less |
| -288 | 0   | 288  | 0 |              |
| 448  | 288 | 0    | 0 |              |
| -288 | 0   | 0    | 0 |              |

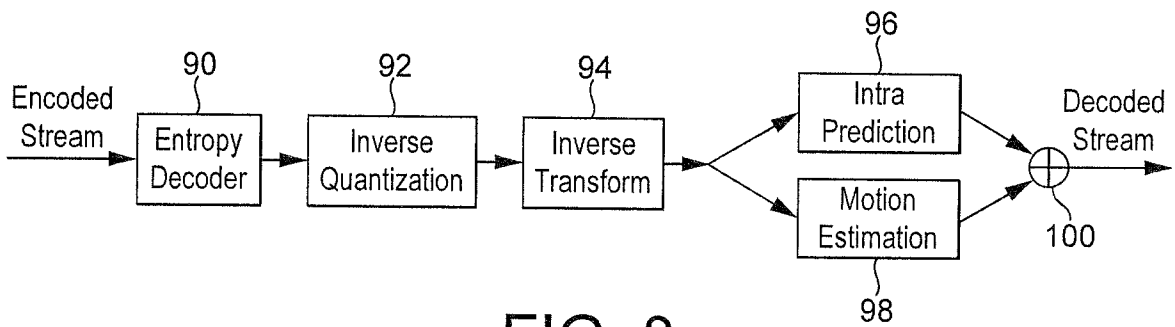
FIG. 5b cont'd

Fig. 6





**FIG. 7**  
(Prior Art)



**FIG. 8**  
(Prior Art)



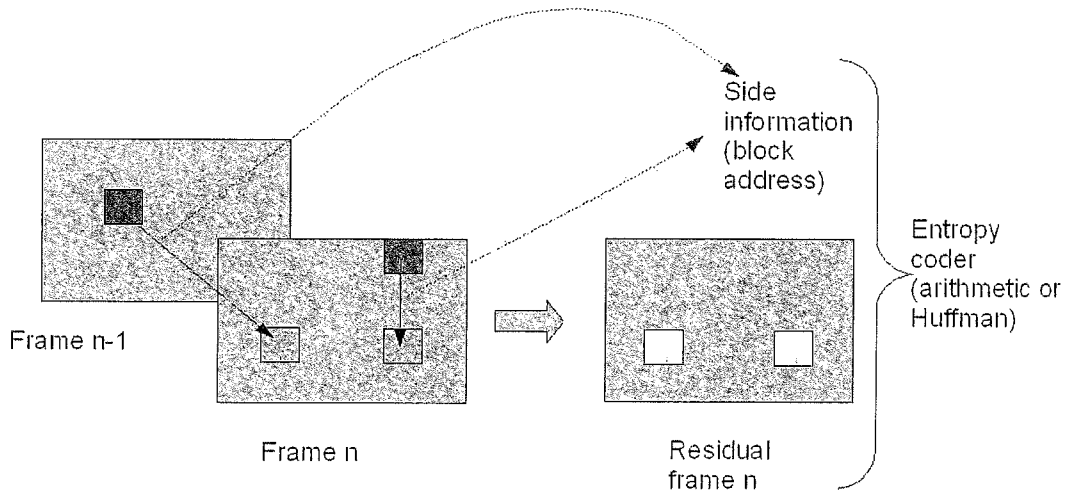


Fig 9

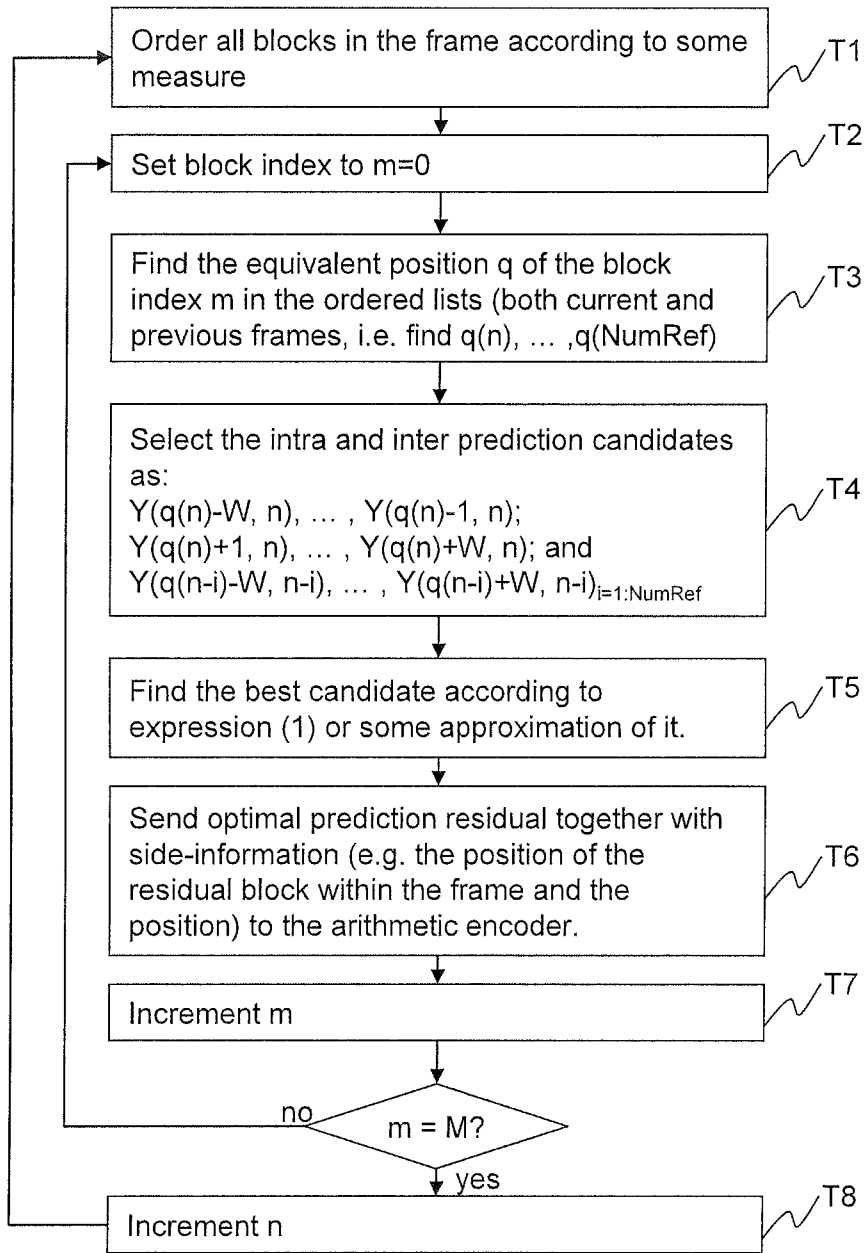


Fig 10

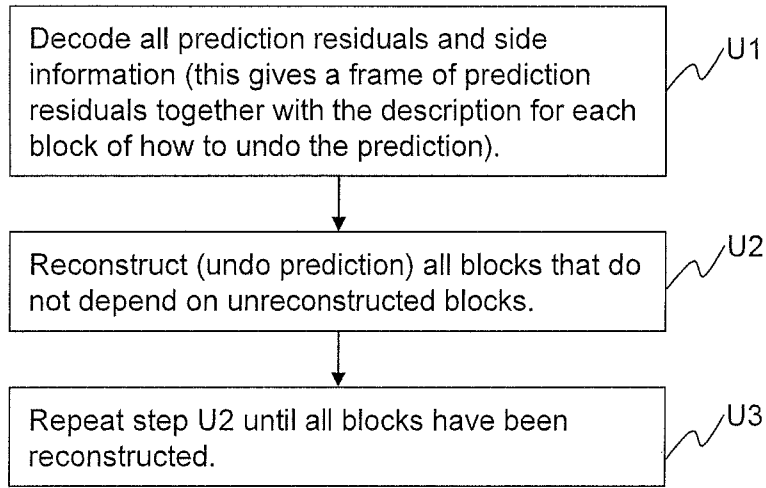


Fig 11

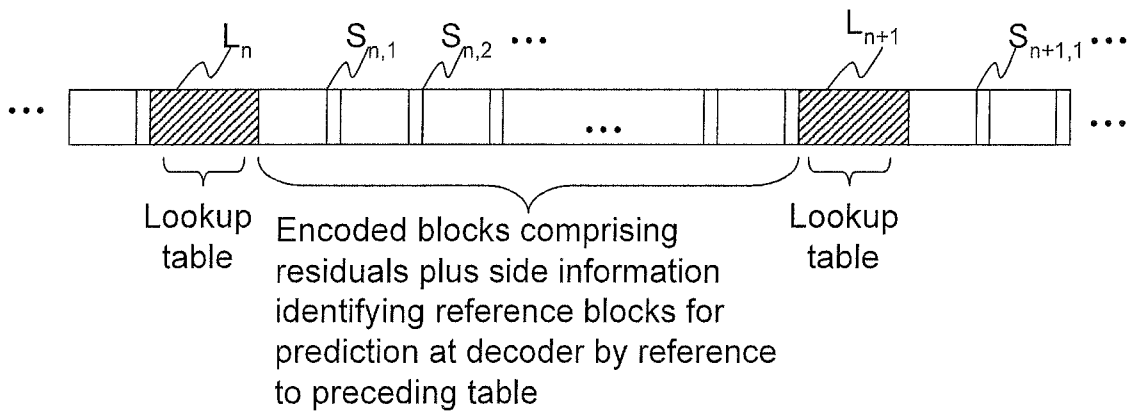


Fig 12

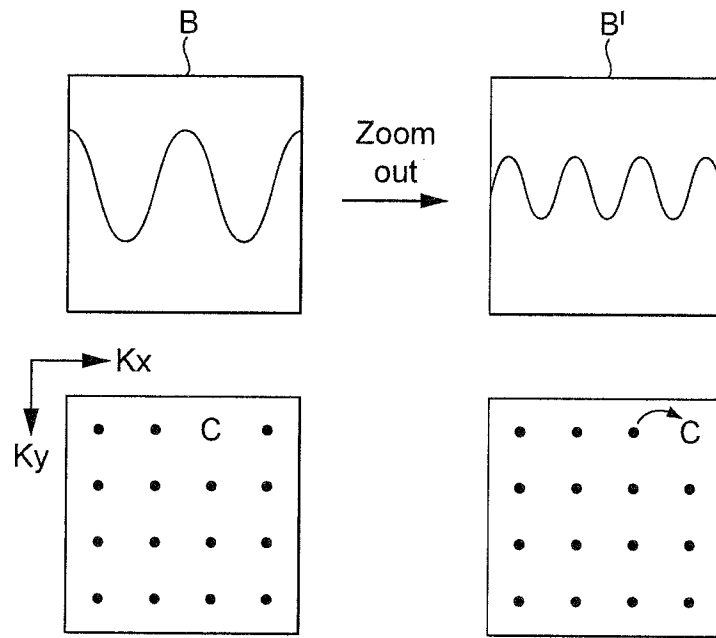


FIG. 13a

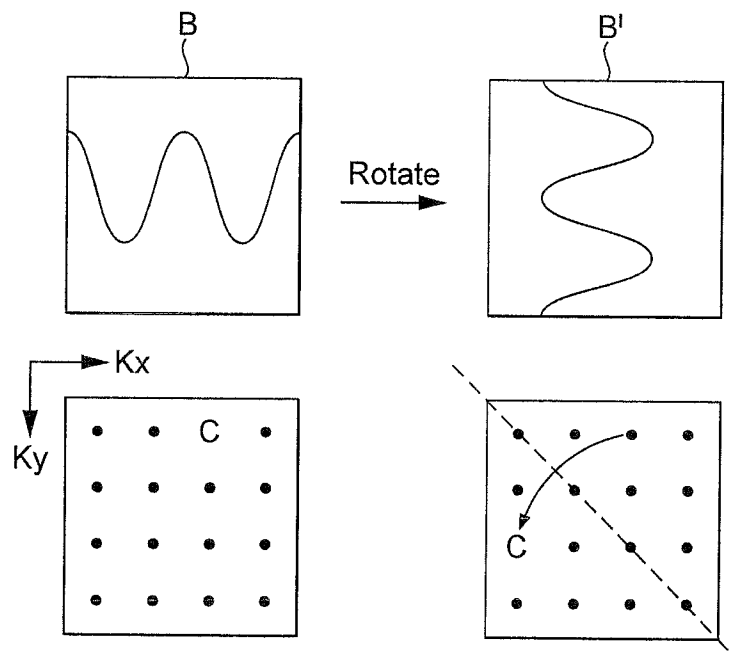


FIG. 13b

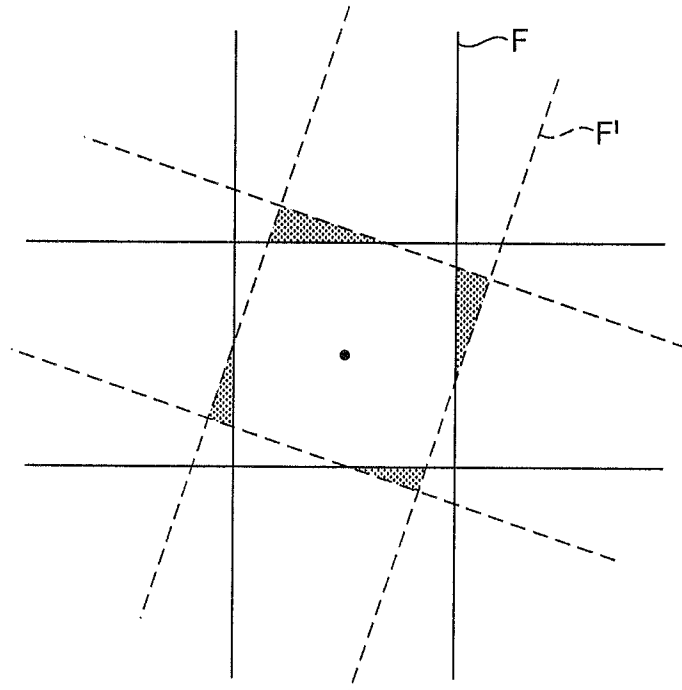


FIG. 13c

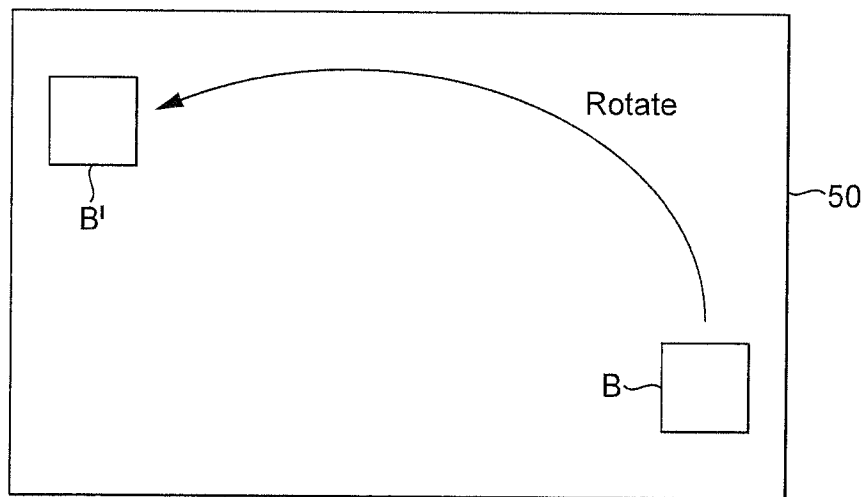


FIG. 13d

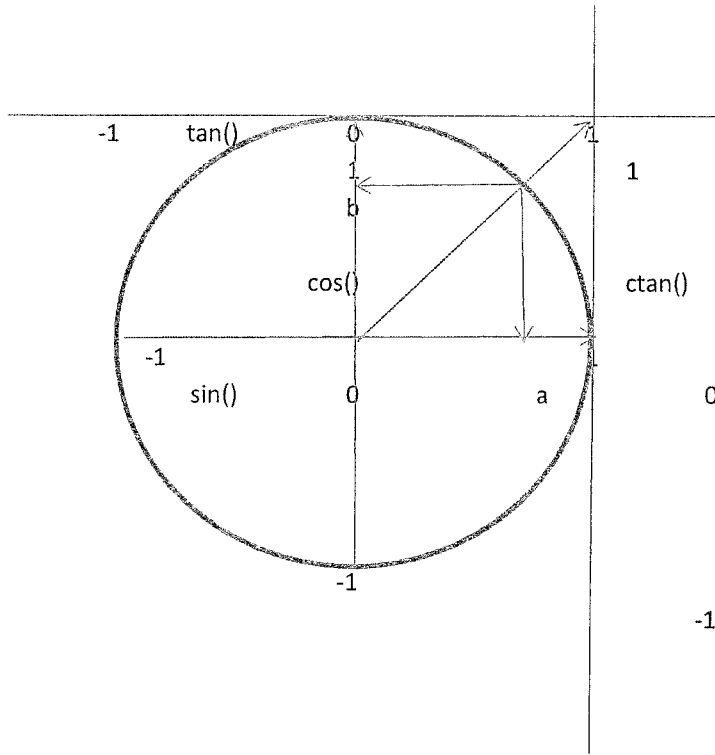


Figure 13e

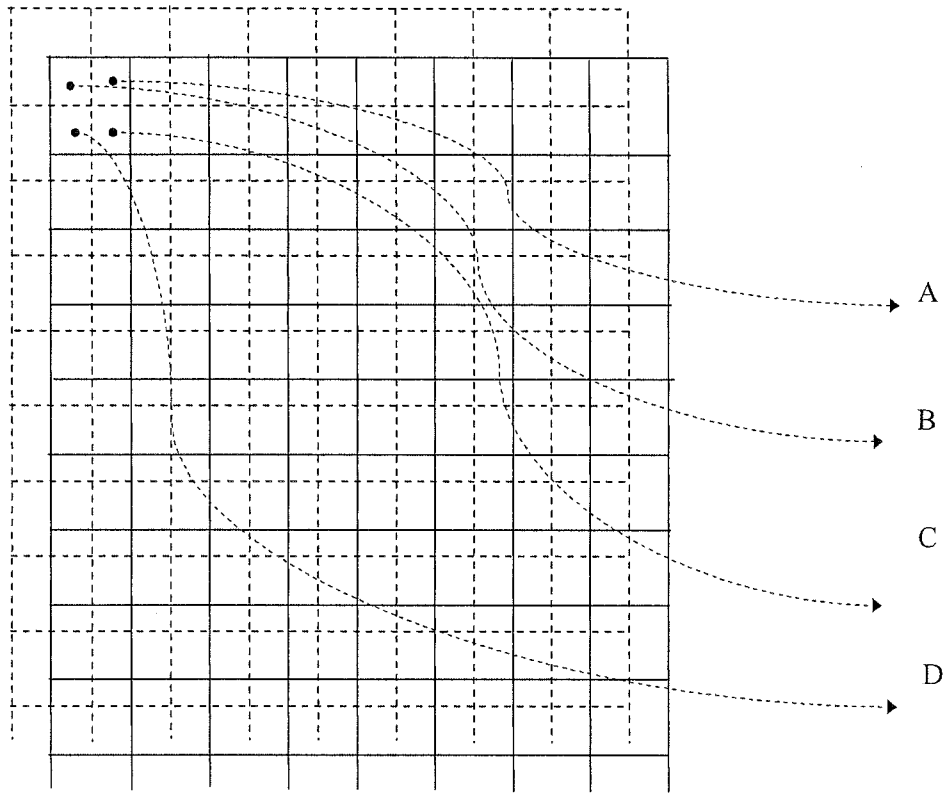


Figure 14a

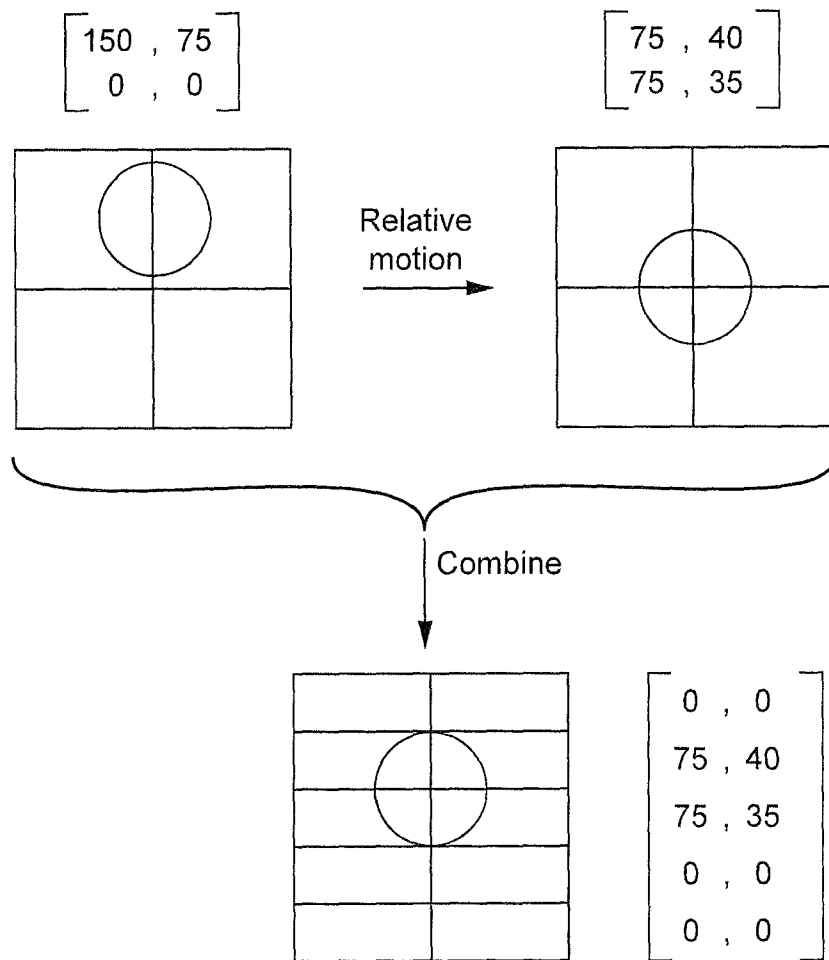


FIG. 14b



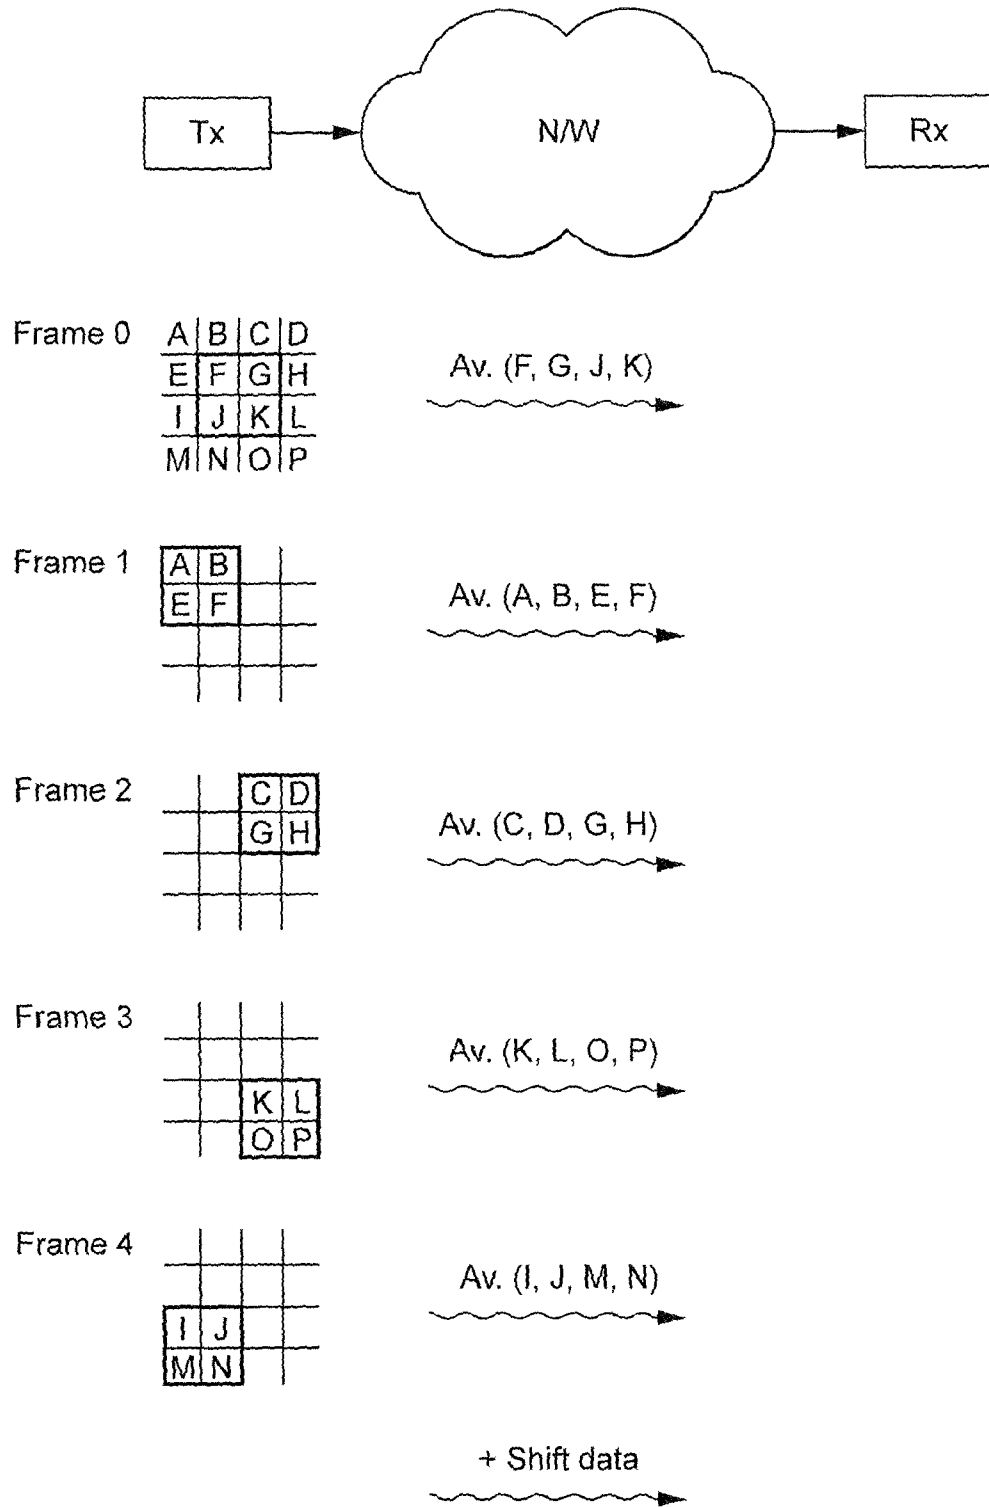


FIG. 14c