




(86) Date de dépôt PCT/PCT Filing Date: 1999/10/06
 (87) Date publication PCT/PCT Publication Date: 2000/04/20
 (45) Date de délivrance/Issue Date: 2009/12/22
 (85) Entrée phase nationale/National Entry: 2001/04/06
 (86) N° demande PCT/PCT Application No.: SE 1999/001785
 (87) N° publication PCT/PCT Publication No.: 2000/022832
 (30) Priorité/Priority: 1998/10/09 (SE9803454-9)

(51) Cl.Int./Int.Cl. *H04N 7/26* (2006.01),
G06T 9/00 (2006.01), *H04N 1/41* (2006.01)
 (72) Inventeurs/Inventors:
CHRISTOPOULOS, CHARILAOS, SE;
ASKELOF, JOEL, SE;
LARSSON, MATHIAS, SE
 (73) Propriétaire/Owner:
TELEFONAKTIEBOLAGET LM ERICSSON, SE
 (74) Agent: NICOLAESCU, SAWYER

(54) Titre : PROCÉDE ET SYSTÈME DE CODAGE DE RÉGIONS D'INTERET
 (54) Title: A METHOD AND A SYSTEM FOR CODING ROIs



(57) Abrégé/Abstract:

In a method and a system for encoding and transmission of still images having at least one region of interest (ROI), the ROI coefficients of an image transformed into the frequency domain, preferably using a wavelet transform are encoded so that they are

(57) **Abrégé(suite)/Abstract(continued):**

transmitted first and can be decoded by a receiver without transmission of the boundary of the ROI. In a preferred embodiment the coefficients belonging to the ROI are shifted so that the minimum ROI coefficient is larger than the largest background coefficient. A receiver can then perform an opposite procedure and thereby obtain the ROI. By specifying how much the coefficients need to be shifted in order to avoid sending shape information several advantages are achieved. Thus, it is possible to avoid sending shape information and to avoid shape encoding at encoder side. Furthermore, there is no need for a shape decoder at receiver side, and there is no need for the receiver to produce the ROI mask. Also, in another preferred embodiment, the shifting (or scaling operations) required at encoder and decoder are also avoided.

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁷ : H04N 7/26, 1/41</p>	A1	<p>(11) International Publication Number: WO 00/22832</p> <p>(43) International Publication Date: 20 April 2000 (20.04.00)</p>
<p>(21) International Application Number: PCT/SE99/01785</p> <p>(22) International Filing Date: 6 October 1999 (06.10.99)</p> <p>(30) Priority Data: 9803454-9 9 October 1998 (09.10.98) SE</p> <p>(71) Applicant (for all designated States except US): TELEFON-AKTIEBOLAGET LM ERICSSON (publ) [SE/SE]; S-126 25 Stockholm (SE).</p> <p>(72) Inventors; and (75) Inventors/Applicants (for US only): CHRISTOPOULOS, Charilaos [GR/SE]; Lomvägen 641, S-192 57 Sollentuna (SE). ASKELÖF, Joel [SE/SE]; S:t Johannesgatan 26, S-753 12 Uppsala (SE). LARSSON, Mathias [SE/SE]; Laxgatan 17, 5 tr, S-133 43 Saltsjöbaden (SE).</p> <p>(74) Agents: SANDSTRÖM, Staffan et al.; Bergensträhle & Lindvall AB, P.O. Box 17704, S-118 93 Stockholm (SE).</p>		<p>(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p>Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>
<p>(54) Title: A METHOD AND A SYSTEM FOR CODING ROIs</p>		
<p>(57) Abstract</p>		
<p>In a method and a system for encoding and transmission of still images having at least one <u>region of interest (ROI)</u>, the ROI coefficients of an image transformed into the frequency domain, preferably using a wavelet transform are encoded so that they are transmitted first and can be decoded by a receiver without transmission of the boundary of the ROI. In a preferred embodiment the coefficients belonging to the ROI are shifted so that the minimum ROI coefficient is larger than the largest background coefficient. A receiver can then perform an opposite procedure and thereby obtain the ROI. By specifying how much the coefficients need to be shifted in order to avoid sending shape information several advantages are achieved. Thus, it is possible to avoid sending shape information and to avoid shape encoding at encoder side. Furthermore, there is no need for a shape decoder at receiver side, and there is no need for the receiver to produce the ROI mask. Also, in another preferred embodiment, the shifting (or scaling operations) required at encoder and decoder are also avoided.</p>		<pre> graph TD 101[Receive Image] --> 103[Specify Bitrate] 103 --> 105[Transform] 105 --> 107[Calc. ROI-mask] 107 --> 109[Obtain Max_coeff.] 109 --> 111[Shift] 111 --> 113[Entropy Code] 113 --> 115[Add Shift Value] 115 --> 117[Add No bytes] 117 --> 119[Output] </pre>

Substitute sheet

1

A METHOD AND A SYSTEM FOR CODING ROIS

TECHNICAL FIELD

The present invention relates to a method and a system for coding of Region of interest (ROI) in still image coding schemes. The method and the system are particularly well suited for use in the JPEG 2000 standard and other wavelet based coders (as in MPEG 4) for still image compression.

BACKGROUND OF THE INVENTION AND PRIOR ART

In the JPEG 2000 standard there is support for the encoding of various parts of the image at various bitrates. A region encoded at a higher bit rate than the other parts of the image is considered a Region of Interest (ROI). Encoding of images with Regions of Interest has been a key issue in recent years. The JPEG 2000 standard under development has addressed the issue of efficient encoding of ROI's, see Charilaos Christopoulos (editor), ISO/IEC JTC1/SC29/WG1 N988 JPEG 2000 Verification Model Version 2.0/2.1., October 5, 1998. One of the modes for ROI coding in the JPEG 2000 verification model (VM) is called "scaling based method". In this method, the ROI coefficients are scaled up (basically shifted up), so that they are coded first during the encoding process. This gives the ability to see the important parts of the image at earlier stages of the transmission. The method increases slightly the bitrate for lossless coding of the image compared to not shifting the coefficients at all, but gives the ability of fast viewing of the important elements of the image, i.e. the ROI's.

In JPEG 2000 the transformed images are encoded bitplane wise. This means that the information about high transform coefficients will be placed earlier in the bit stream than the rest of the information. The current "scaling based coding method" for ROI coding is based on this fact. The coefficients corresponding to the ROI are upshifted prior to arithmetically encoding them. This means that information for these coefficients will be transmitted earlier in the bitstream than it would have without the shifting. At the early stages of the transmission, the ROI will be reconstructed with better quality than the BG. The whole operation is progressive by resolution or

Substitute sheet

2

by quality.

Furthermore, E. Atsumi and N. Farvardin, "Lossy/lossless region-of-interest coding based on set partitioning in hierarchical trees", Proceedings of IEEE International Conference on Image Processing (ICIP-98), Chicago, Illinois, USA, October 4-7, 1998, describes the general idea of the scaling based coding method. In addition, encoding of ROI's is disclosed in US 5,563,960, Oct. 8, 1996, although the ROI coding method described only performs scaling of the image data and not of the coefficients.

Using the methods as described above when encoding an image at various bitrates, information about what parts of the image should be encoded at what bit rate need be available to the encoder. Whereas the ROI might easily be described in the spatial domain, it will be more complicated in the transform domain. So far the information about the ROI shape must be available to the encoder and the decoder, thus it requires extra bits in addition to the bits representing the texture information. Moreover, a shape encoder is required (at the transmitter) and a shape decoder (at the receiver), making the whole system more complex and expensive to implement. The decoder has also to produce the ROI mask, i.e. it has to define which are the coefficients needed for the reconstruction of the ROI, see Charilaos Christopoulos (editor), ISO/IEC JTC1/SC29/WG1 N988 JPEG 2000 Verification Model Version 2.0/2.1., October 5, 1998, and this adds to the computational complexity and memory requirements of the receiver, which should be as simple as possible.

The currently used method to solve these problems is to include the description of the ROI in the spatial domain, in the bitstream. The necessary mask of ROI coefficients (ROI mask) for the transform domain is then created in both the encoder and the decoder, see for example Charilaos Christopoulos (editor), ISO/IEC JTC1/SC29/WG1 N988 JPEG 2000 Verification Model Version 2.0/2.1., October 5, 1998. The encoder encodes the shape information and the encoded bitstream with the shape information

Substitute sheet

3

is added to the total bitstream and transmitted to the receiver. The receiver from the shape information decodes the shape, makes the ROI mask and then decodes the texture information of the image.

In the case where the ROI shape is simple (for example rectangle or circle) the shape information is not requiring many bits. However, even in these simple cases, the receiver has to produce the ROI mask, which means that the receiver requires memory as large as the whole image (but of 1 bit/pixel) and has a certain computational complexity (since the creation of the mask is similar to doing a wavelet transform). For a complex ROI this means that a lot of information need be transmitted between encoder and decoder and computational complexity becomes an issue. The additional overhead for shape information is significant, particularly for low bitrates.

Also, the co-pending Swedish Patent number SE 512840 describes a method in which both encoder and decoder needs to use and to define the ROI mask, i.e. to find which coefficients belong to the ROI or are needed for the ROI.

SUMMARY

It is an object of the present invention to provide a method and a system whereby no shape information needs to be transmitted in an ROI coding scheme.

This object is obtained by a method and a system wherein the ROI coefficients are encoded so that they are transmitted first and can be decoded by a receiver without transmission of the boundary of the ROI.

In a preferred embodiment the coefficients belonging to the ROI are shifted so that the minimum ROI coefficient is larger than the largest background coefficient. A receiver can then perform an opposite procedure and thereby obtain the ROI.

By specifying how much the coefficients needs to be shifted in order to avoid sending shape information several advantages are

Substitute sheet

4

achieved. Thus, it is possible to avoid sending shape information and to avoid shape encoding at encoder side. Furthermore, there is no need for a shape decoder at receiver side, and there is no need for the receiver to produce the ROI mask.

Also, in another preferred embodiment the shifting (or scaling operations) required at encoder and decoder are also avoided.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will now be described in more detail and with reference to the accompanying drawings, in which:

- Fig. 1 is a flow chart illustrating the steps carried out at an encoder according to a first embodiment.
- Fig. 2 is a flow chart illustrating the steps carried out at an encoder according to a second embodiment.
- Fig. 3 is a flow chart illustrating the steps carried out at an encoder according to a third embodiment.
- Fig. 4 is a flow chart illustrating the steps carried out at a decoder according to the first and second embodiment.
- Fig. 5 is a flow chart illustrating the steps carried out at a decoder according to the third embodiment.
- Fig. 6 is a flow chart illustrating the steps carried out at an encoder according to a fourth embodiment.
- Fig. 7 is a flow chart illustrating the steps carried out at a decoder according to the fourth embodiment.
- Fig. 8 is an illustration of a bitstream syntax used.
- Fig. 9 is an illustration of an alternative bitstream syntax.

DETAILED DESCRIPTION

In Fig. 1, a flow chart illustrating the steps carried out at an encoder according to a first embodiment is shown. Thus, first in a step 101 an input image is received and its Region of Interest (ROI) is specified. Next, in a step 103, the required bitrate or quality for the ROI and the Background (BG) is received.

Thereupon, the image is transformed into wavelet domain, step 105. Next an ROI mask is calculated for example using the method described in Charilaos Christopoulos(editor), ISO/IEC JTC1/SC29/WG1 N988 JPEG 2000 Verification Model Version

Substitute sheet

5

2.0/2.1., October 5, 1998, step 107.

Thereupon, the maximum wavelet coefficient(s) (MAX_Coeff) in the BG or the whole image are obtained, step 109. All coefficients in the ROI mask are then shifted so much that the minimum coefficient in the ROI mask is larger than the MAX_coeff, step 111. The image is then entropy coded until the ROI quality or bitrate and BG quality or bitrate specified in step 103 is achieved, step 113.

Next, the shifting value is added in the bitstream so that the decoder can find and read it, step 115. This is required for the decoder, since the decoder needs to know how much will the coefficients be down-shifted. Next, the number of bytes which were needed for encoding the coefficients in the ROI mask are added, step 117. The result in step 117 is used as output from the encoder, step 119.

If the shifting value is selected so that the minimum coefficients in the ROI mask is larger than the maximum coefficient in the BG, then during encoding only ROI coefficients will be coded, until the BG coefficients become significant. At that stage, all ROI coefficients have been coded and have to downshifted at the receiver, while the remaining coefficients correspond to the BG and need not to be downshifted. The receiver needs to know the number of bytes (or bits) which correspond to full coding of the ROI coefficients (i.e. at the point where the first BG coefficients starts to be coded). This information is put in the bitstream header and extracted from the receiver.

As an example, assume that the ROI coefficients are shifted left 8 times (i.e. multiply each ROI coefficients by 2^8) and that all ROI coefficients then become larger than the largest BG coefficient. Then encoding starts and when all ROI coefficients are coded, the shifting value and the number of bytes needed for the ROI coefficients (Nbytes_ROI) is put in the bitstream header. The encoding continues as usual. The decoder gets the bitstream and starts decoding. When the decoder decodes certain

number of bytes which are less than the Nbytes_ROI, it shifts each coefficient down.

It should be noted that there are BG coefficients reconstructed which are zero since they were not coded at this stage and they will be shifted down since the decoder doesn't know anything about this. Up to when the number of bytes received is equal to Nbytes_ROI, then all coefficients are shifted down at the receiver. After this stage, no coefficients ^{are} is shifted down. Notice here that the ROI coefficients are no longer updated since they are zero coded.

Using this method, the decoder doesn't need any shape information. The decoder doesn't need to know which coefficients correspond to an ROI, since it will be shifting down all coefficients, i.e. BG coefficients will be zero till all ROI coefficients are coded. The decoder doesn't need to produce any ROI mask, making the coding scheme even simpler. The only thing that the decoder needs to do is the downshifting of the received ROI coefficients.

In Fig. 2, a flow chart illustrating the steps carried out in a second embodiment of an encoder is shown. The flow chart in Fig. 2 is identical to the flow chart in Fig. 1 except for that the maximum quantized coefficients are obtained and shifted in the steps 109 and 111 as is shown in the steps 209 and 211.

Using the method and encoder according to the second embodiment, less memory is needed for storing the shifted coefficients, because the quantized coefficients are smaller than the initial coefficients.

In Fig. 3, a flow chart illustrating the steps carried out in a third embodiment of an encoder is shown. The flow chart in Fig. 3 differs from the flow charts in Fig. 1 and Fig. 2 in that the number of bytes which were needed for encoding the coefficients in the ROI mask is not stored in the bitstream. Therefore, the flow chart in Fig. 3 does not comprise a step 317 corresponding to the steps 117 and 217. Even if it is stored it is not used at

Substitute sheet

7

the decoder. Therefore, the third embodiment is similar to the first and second embodiment, but requires less information to be stored in the bitstream.

Below the decoder operations corresponding to the different encoding schemes described above in conjunction with Figs 1 - 3 are described. Thus, in Fig. 4 a flow chart illustrating the steps carried out at a decoder according to the first and second embodiment is shown.

First, in a step 401 the header of the bitstream encoded according to the algorithm described above in conjunction with Figs. 1 and 2 is received. The information about the shifting value used and the number of bytes (ROI_bytes) corresponding to the ROI coefficients (i.e. those that were shifted) is obtained. Next, in a step 403, the rest of the bitstream is received. If number of bytes received is less than ROI_bytes, after an entropy decoding of coefficients, they are down-shifted by the shifting value, step 405. It should be noticed that BG coefficients up to this stage were coded to zero, so down-shifting does not affect them. Intermediate reconstructed images can be obtained by inverse wavelet transform.

Next, in a step 407 it is evaluated if the number of bytes received is less than ROI_bytes, if yes then return to step 403. Else proceed to step 409. In step 409 the rest of the bitstream is received. This corresponds now to BG data and therefore from that stage on no coefficient will be downshifted. Finally, in step 411, an inverse wavelet transform gives the reconstructed image.

It should be noticed that the number of bytes is not really needed to be known to the decoder. This is because the decoder can scale down all coefficients that are above where shift_value is the shifting value used.

In Fig. 5, a flow chart illustrating the steps carried out in a decoder arranged to decode a bitstream encoded according to the third embodiment described above in conjunction with Fig. 3 is

Substitute sheet

8

shown.

Thus, first in a step 501, the encoded bitstream is received. Information about the shifting value used is obtained. Next, in step 503 it is evaluated if the received coefficient is larger than $2^{(\text{shift value})}$, if so then the coefficient is down shifted by the shift value in a step 504, else the process proceeds to step 505. In step 505 it is decided to not downshift the coefficient. Finally, in step 507, an inverse wavelet transform of the output values from the steps 504 and 505 gives the reconstructed image.

It should be noticed that some problems may appear in floating point wavelets where some coefficients might be between 0 and 1 and therefore they will never become larger than $2^{(\text{shift value})}$ after shifting at the encoder. This means that they will never be downshifted at the decoder. To avoid such a problem the encoder according to the second embodiment could be used, where the quantized coefficients are shifted, since the quantized coefficients are integers.

The methods described above requires that the encoder shifts the coefficients up, i.e. multiplies them with a certain factor. Although the computational complexity of such operation is small, an alternative way to avoid transmitting the shape information and minimising the computational complexity of the decoder, which also avoids the down shifting operation at the decoder is sometimes advantageous.

In Fig. 6, a flow chart illustrating the steps carried out in an encoder providing an encoded bit stream which does not need to be down shifted is shown.

Thus, first in a step 601, an input image to be encoded is received and its ROI is specified. The required bitrate or quality for ROI and BG is the received, step 603. Next, the image is transformed into the wavelet domain and stored to a first memory (MEM1), step 605.

Thereupon, an ROI mask is generated as described above, step

Substitute sheet

9

607. The content of the first memory (MEM1) is then copied to a second memory (MEM2), step 609. The step 609 is only required if the BG information is to be used in later stages. Then all coefficients of MEM1 outside the ROI mask are set to zero, step 611. The wavelet coefficients in MEM1 are then encoded using JPEG 2000 encoding methods, step 613.

It should be noticed that encoding is performed for all coefficients. However, since BG coefficients are zero, only ROI coefficients are actually coded (BG coefficients corresponding to zero are also coded but they don't occupy much bitrate).

Encoding of MEM1 is performed until the required bitrate (ROI rate of ROI quality), then the encoder will switch and start encoding from the coefficients in MEM2, step 615. All coefficients in the ROI mask are then set to zero in MEM2, step 617. This means that BG coefficients will be coded.

Finally, in step 619 encoding of MEM2 (where coefficients in ROI mask are set to zero) is done until the bitrate or quality specified for the BG is achieved.

In Fig. 7, a flow chart illustrating the steps carried out in a decoder arranged to decode a bitstream encoded according to the embodiment described above in conjunction with Fig. 6 is shown.

Thus, first in a step 701, a bitstream encoded according to the algorithm described above in conjunction with Fig. 6 is received until the number of bytes for MEM1 is obtained. Entropy decoding is performed and the MEM1 coefficients are obtained. It should be noticed that it is possible to perform an inverse wavelet transform to produce intermediate reconstructed image if this is required or desired.

Next, in a step 703, the rest of the bitstream is received until the total number of bytes is received. Entropy decoding is performed and the MEM2 coefficients are obtained. The MEM2 coefficients are added to the MEM1 coefficients. An inverse wavelet transform of the added coefficients produces the

Substitute sheet

10

reconstructed image with both ROI and BG.

It should be noticed that when bits corresponding to MEM2 coefficients are received, the receiver can do inverse wavelet transform to reconstruct only the image corresponding to MEM2 (i.e. has only the BG). Then it can add the reconstructed images of MEM1 and MEM2 together instead of adding the coefficients as described above.

The bitstream sent from the encoder has to have information on how many bytes (or bits) were coded for MEM1 image (where the BG were set to zero). This is because the receiver has to know when it starts receiving MEM2 (where the coefficients in ROI mask are set to zero). The receiver in this case adds the reconstructed MEM2 coefficients to the reconstructed MEM1 coefficients.

This method avoids completely shifting of coefficients at the encoder and decoder, avoids transmission of shape information, avoids the use of shape encoder and shape decoder and avoids the generation of the ROI mask at the decoder. The decoder must only know when it stops receiving MEM1 coefficients and starts receiving MEM2 coefficients, so it can add the coefficients together.

For real time encoding and transmission (encode and send simultaneously), the receiver might not know the total number of bytes spent for ROI coding. In this case the transmitter has to send a signal at the stage where the ROI coding has finished informing the receiver during the transmission that ROI coefficients were coded and after this stage he should not downshift any coefficient. This can be done by sending a codestream that can't be emulated from the arithmetic encoder.

It should be noticed that when the encoder according to the third embodiment is used, then no problem appears in real time applications. This is because the number of bytes for ROI coefficients is not really needed. What is needed is the shifting value, which is always included in the bitstream. Since

the decoder down-shifts the coefficients that are above, it doesn't need any signal at all. Therefore, encoder according to the third embodiment has advantages compared to the other schemes.

It should also be noticed that although the methods described above consider the existence of one ROI, in practice there could be more than one. The methods apply in similar manner. In such cases one could consider that for a first ROI (ROI 1) all coefficients are shifted with the method described in previous section. Then for a second ROI (ROI 2) all coefficients are shifted in a similar manner but in such a way so that they are larger than the shifted coefficients of ROI 1. Then the procedure continues in a similar manner. The decoder can find which coefficients belong to ROI 1 and which to ROI 2 by having the information about the shifting value for each ROI.

It is also to be noticed that for rectangular shapes, the mask generation in the decoder can be avoided if at each subband information, of the ROI shape is transmitted. Rectangular ROI's have the property of producing rectangular ROI shapes in each subband and therefore the information of the ROI shape can be sent for each subband (for example upper left and lower right corner). This avoids the generation of the ROI mask in the decoder, however, it applies only for rectangular ROI shapes. The methods presented above are valid for any shape.

If the encoding is performed in the manner as described above, no information about the ROI need be transmitted except the value by which the coefficients have been upshifted and the number of bits of the bitstream that contain information that has been upshifted.

This means:

- * No information about the shape of the ROI need be transmitted. This saves a lot of bits in the bitstream especially for complex shapes.
- * There is no need to generate the mask for the ROI coefficients in the decoder. This saves memory and computational complexity

Substitute sheet

12

of the decoder

- * Shape encoding is not required
- * Shape decoding is not required
- * Approximate extraction of the ROI shape information is possible from the bitstream.

* Because the ROI mask expands in the higher subbands, it will start covering some of the BG or even all at the higher subbands (for example the LL subband). This means that BG information will be coded together with ROI information. Therefore, during decoding the receiver at the early stages of the transmission will recover most of the image (especially when progression by resolution is implemented). This avoids the problem of having a black BG reconstructed at the early stages.

In Fig. 8, a possible bitstream syntax is given. It should be noticed that this can be part of the total bitstream syntax which includes information on image type, etc. If encoding is done first and transmission after, the bitstream syntax should contain the following information:

- ROI coding method (ROI_CM), transmitted in a first field 801
- Shift value (SV), transmitted in a second field 803
- Number of Bytes for ROI (NBytes_ROI) (if needed), transmitted in a third field 805
- Rest of header info and bits (see Charilaos Christopoulos (editor), ISO/IEC JTC1/SC29/WG1 N988 JPEG 2000 Verification Model Version 2.0/2.1., October 5, 1998, transmitted in a fourth field 807

Where,

- ROI_CM: specifies that the proposed ROI coding method is used
- SV: specifies the value that the ROI coefficients were shifted up
- Nbytes_ROI: specifies the total number of bytes spent for coding the ROI coefficients (not needed for all encoding schemes specified above).

If the real time case is considered, i.e. encoding and transmission is performed simultaneously), the bitstream could be modified to the syntax shown in Fig. 9.

- ROI coding method (ROI_CM), transmitted in a first field 901
- Shift value (SV), transmitted in a second field 903
- Bits corresponding to rest of header and coefficients, transmitted in a third field 905
- Signal, transmitted in a fourth field 907

Where,

- Signal = codestream that can't be emulated from the arithmetic encoder (not needed for all encoding schemes specified above).

Substitute Sheet

14

CLAIMS

1. A method of encoding and transmitting an image having at least one Region of Interest (ROI), wherein the image is transformed into the frequency domain thereby producing coefficients, the method comprising the steps of:
 - shifting the coefficients corresponding to the at least one ROI, thereby producing shifted coefficients, so that a smallest shifted coefficient corresponding to the at least one ROI becomes larger than a largest coefficient of the remaining coefficients; and
 - encoding the shifted coefficients corresponding to the at least one ROI, so that they are transmitted first and so that no information regarding a shape of the at least one ROI needs to be transmitted in order to perform a decoding of the at least one ROI.
2. A method according to claim 1, further comprising, before the step of shifting, the step of quantizing the coefficients corresponding to the at least one ROI.
3. A method according to any of claims 1 - 2, wherein the image is transformed into the frequency domain using a wavelet transform.
4. An apparatus for encoding an image comprising:
 - means for defining at least one Region of Interest (ROI) in the image; and
 - means for transforming the image into the frequency domain thereby producing coefficients, said means for transforming comprising:
 - means for shifting said coefficients corresponding to the at least one ROI, thereby producing shifted coefficients, so that a smallest shifted coefficient corresponding to the

Substitute Sheet

15

at least one ROI becomes larger than a largest coefficient of the remaining coefficients; and

- means for encoding the shifted coefficients corresponding to the at least one ROI, so that they are transmitted first and so that no information regarding a shape of the at least one ROI needs to be transmitted in order to perform the decoding of the at least one ROI.

5. An apparatus according to claim 4, further comprising means, connected to the shifting means, for quantizing the coefficients corresponding to the at least one ROI before shifting them.

6. An apparatus according to any of claims 4 - 5, further comprising means for transforming the image into the frequency domain using a wavelet transform.

7. A method of encoding an image, comprising the steps of:
- specifying at least one Region of Interest (ROI);
- transforming the image using a wavelet transform thereby producing wavelet coefficients;
- generating a ROI mask from the at least one ROI;
- determining a maximum wavelet coefficient in the image; and
- shifting the wavelet coefficients in the ROI mask, thereby producing shifted wavelet coefficients, so that a smallest shifted wavelet coefficient in the ROI mask becomes larger than a largest wavelet coefficient outside the ROI mask.

8. A method according to claim 7, wherein only the wavelet coefficients outside the ROI mask are searched when determining the maximum wavelet coefficient.

9. A method according to claim 7 or 8, further comprising the steps of:

Substitute Sheet

16

- determining a number of bits to use for encoding the image; and
- encoding the image until the determined number of bits is obtained.

10. A method according to claim 9, wherein the image is encoded using an entropy code.

11. A method according to any of claims 7 - 10, further comprising the step of quantizing the wavelet coefficients prior to determining the maximum wavelet coefficient.

12. A method of transmitting an image encoded according to any of claims 7 - 11, from a transmitter to a receiver, comprising the step of adding a value used to shift the wavelet coefficients in a transmitted bit stream.

13. A method according to claim 12, further comprising the step of adding a number of bytes used for encoding the wavelet coefficients in the ROI mask in the transmitted bit stream.

14. An encoder arranged to encode an image according to any of claims 7 - 11.

15. A decoder arranged to decode an image encoded according to any of claims 7 - 11.

16. A method of encoding an image, comprising the steps of:

- specifying at least one Region of Interest (ROI);
- transforming the image using a wavelet transform thereby producing wavelet coefficients;
- storing the wavelet coefficients in a first memory;
- generating a ROI mask from the at least one ROI;
- copying the wavelet coefficients of the first memory into a second memory;

Substitute Sheet

17

- setting the wavelet coefficients in the first memory outside the ROI mask to zero;
- setting the wavelet coefficients in the second memory inside the ROI mask to zero;
- encoding the wavelet coefficients in the first memory; and
- encoding the wavelet coefficients in the second memory when all the wavelet coefficients in the first memory are encoded.

17. A method according to claim 16, further comprising the steps of:

- determining a number of bits to use for encoding the image; and
- encoding the image until the determined number of bits is obtained.

18. An encoder arranged to encode an image according to any of claims 16 - 17.

19. A decoder arranged to decode an image encoded according to any of claims 16 - 17.

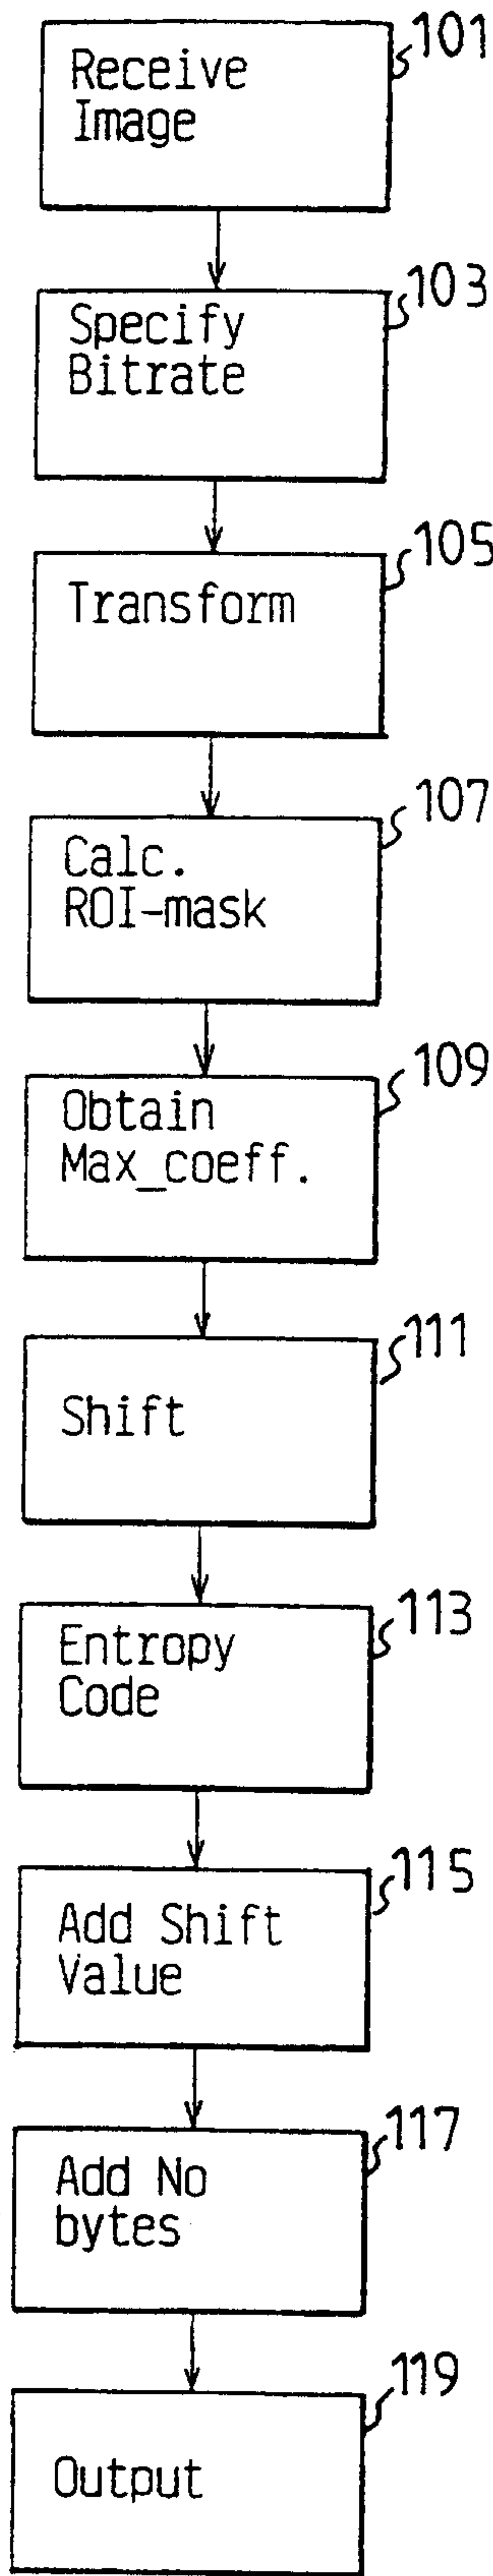


FIG.1




FIG. 2




FIG. 3




FIG. 4

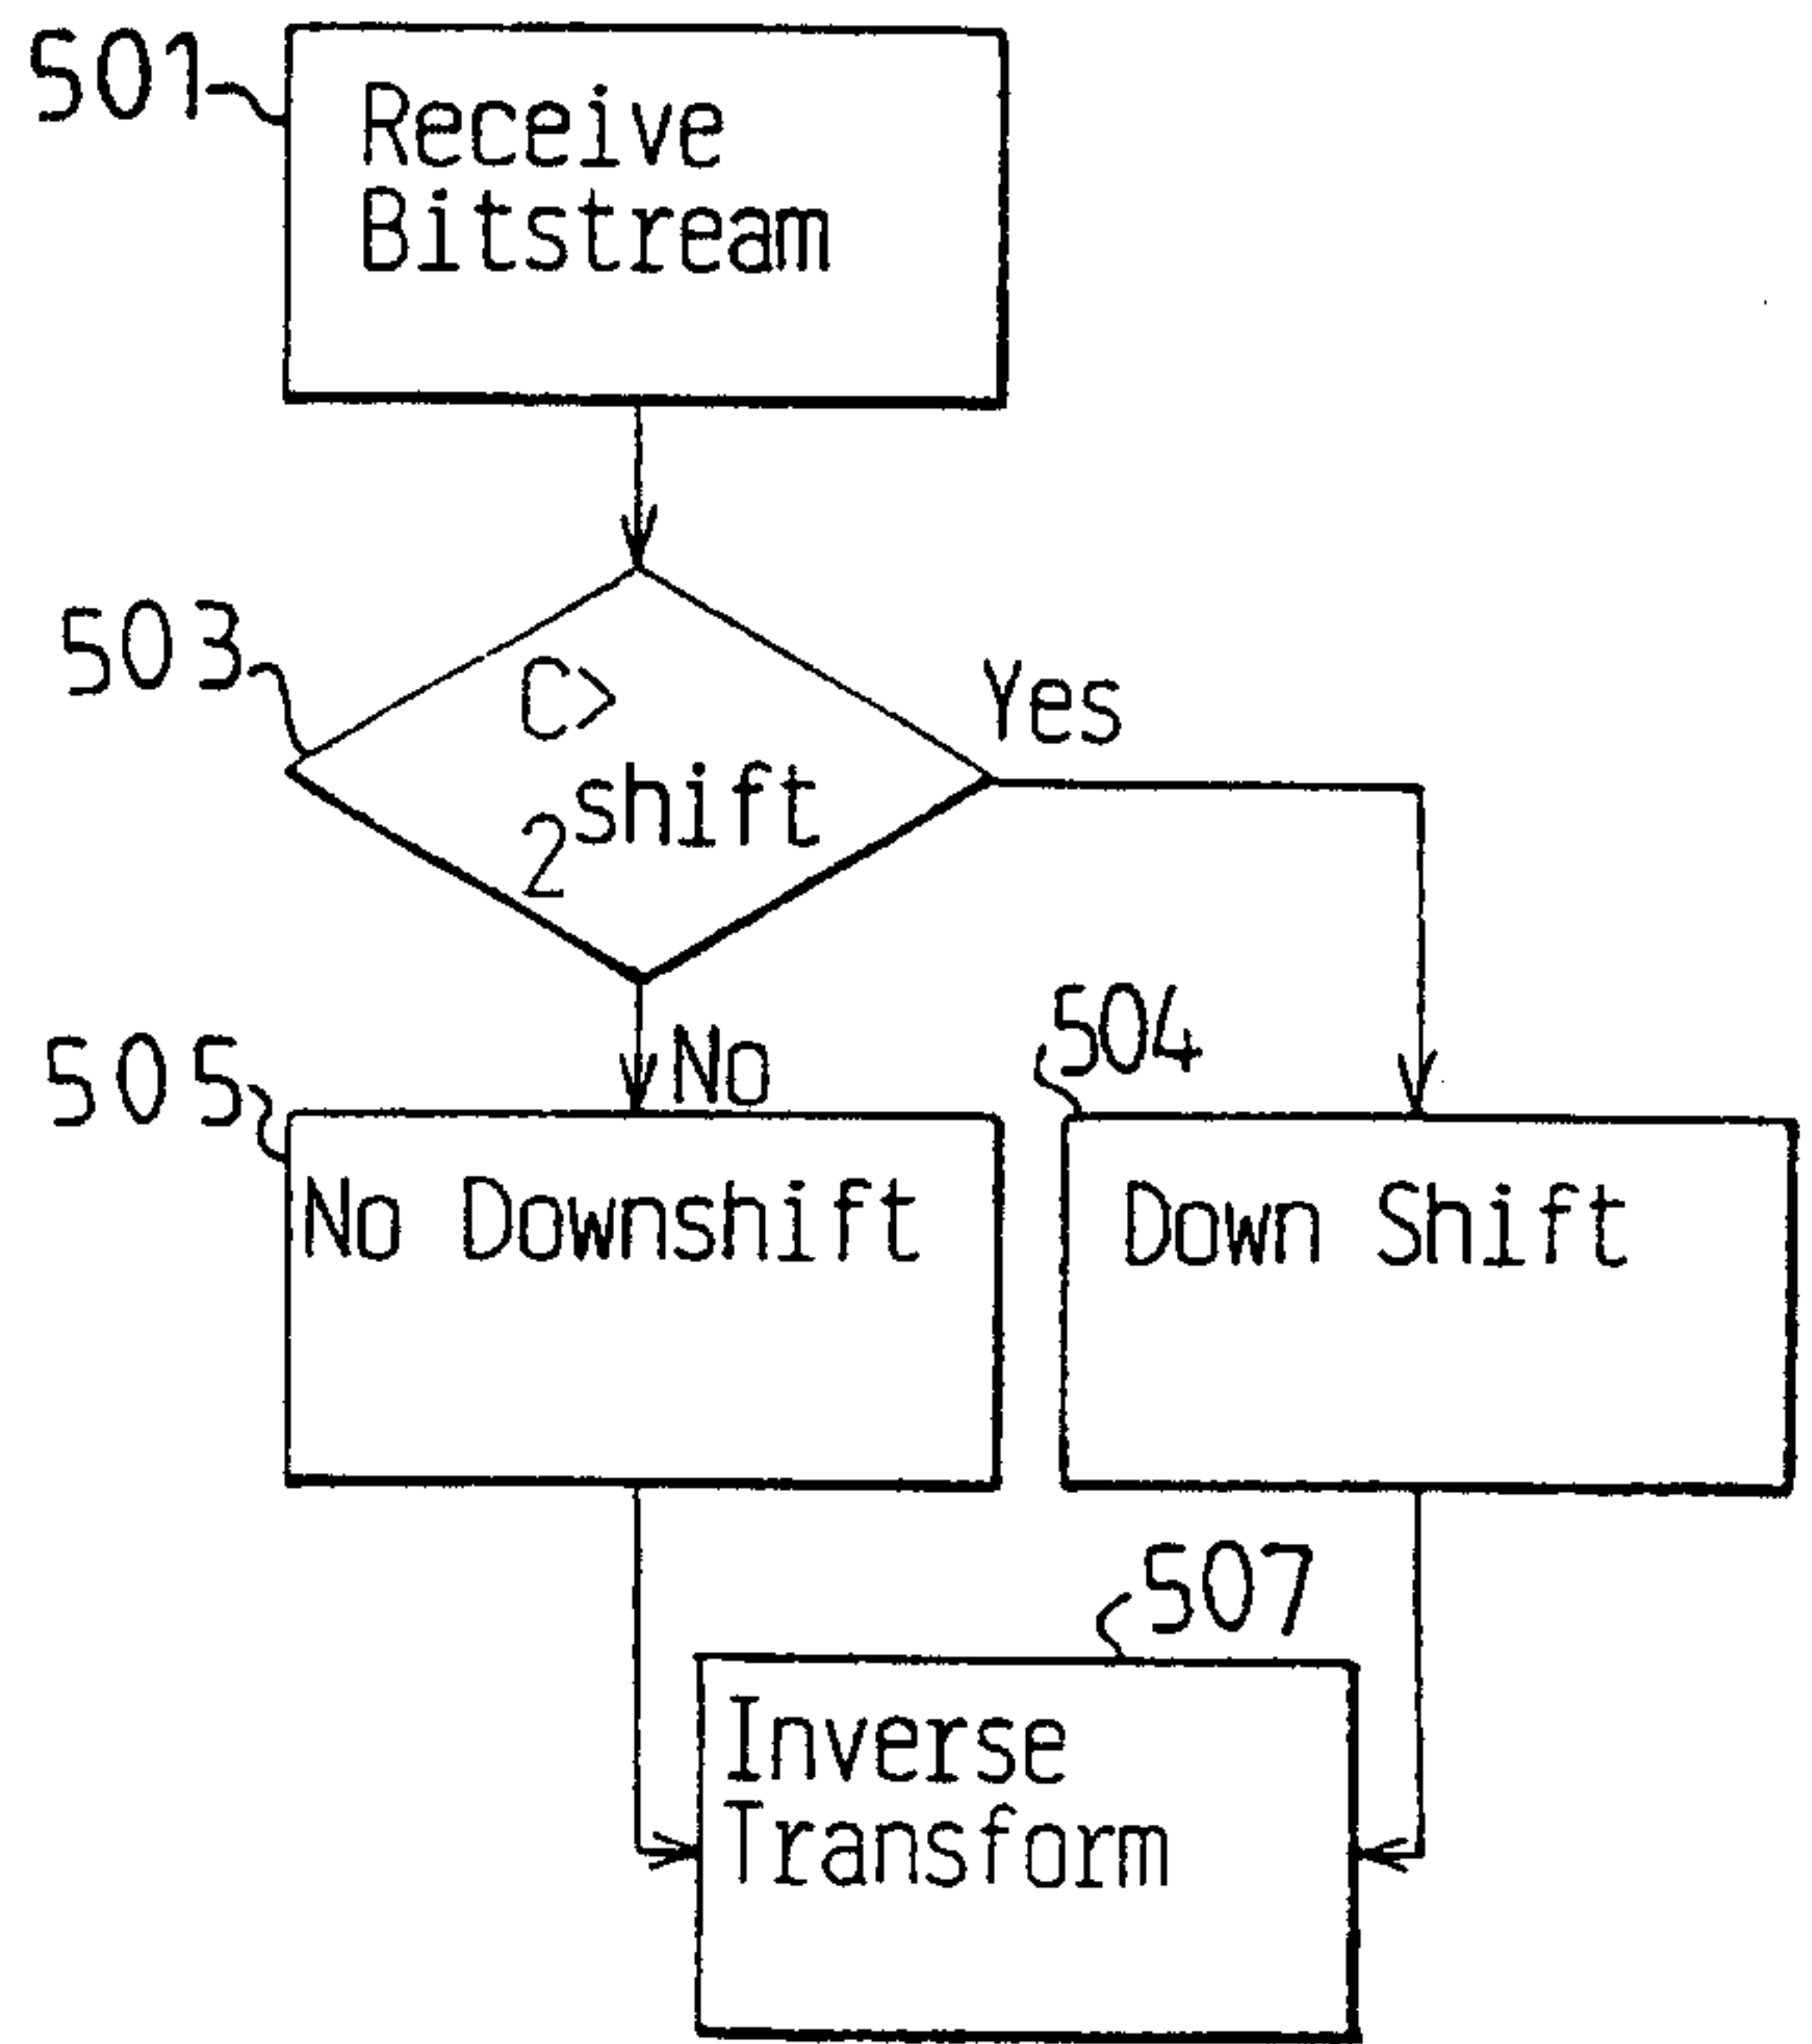


FIG. 5

Harish Chelak

Substitute sheet

3 / 4




FIG. 6

FIG. 7

	§ 801	§ 803	§ 805	§ 807
---	ROI coding method (ROI_CM)	Shift value (SV)	Number of Bytes for ROI (NBytes_ROI)(if needed)	Rest of header info and bits (see())

Where ROI_CM: specifies that the proposed ROI coding method is used

FIG. 8

SV: specifies the value that the ROI coefficients were shifted up.

Nbytes_ROI: specifies the total number of bytes spent for coding the ROI coefficients

	§ 901	§ 903	§ 905	§ 907
---	ROI coding method (ROI_CM)	Shift value (SV)	Bits corresponding to rest of header and coefficients	Signal

Where Signal=codestream that can't be emulated from the arithmetic encoder.

FIG. 9

