

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
8 July 2010 (08.07.2010)

PCT

(10) International Publication Number  
**WO 2010/077325 A2**

- (51) International Patent Classification:  
*H04N 7/26* (2006.01)
- (21) International Application Number:  
PCT/US2009/006653
- (22) International Filing Date:  
17 December 2009 (17.12.2009)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
61/203,805 29 December 2008 (29.12.2008) US  
61/203,807 29 December 2008 (29.12.2008) US
- (71) Applicant (for all designated States except US): **THOMSON LICENSING** [FR/FR]; 1 rue Jeanne d'Arc, F-92443 Issy-les moulineaux (FR).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): **JOSHI, Rajan, Laxman** [IN/US]; 838 S. Barrington Avenue, #205, Los Angeles, California 90049 (US).
- (74) Agents: **SHEDD, Robert, D.** et al.; Thomson Licensing LLC, Two Independence Way|Suite 200, Princeton, New Jersey 08540 (US).

- (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, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: METHOD AND APPARATUS FOR ADAPTIVE QUANTIZATION OF SUBBAND/WAVELET COEFFICIENTS

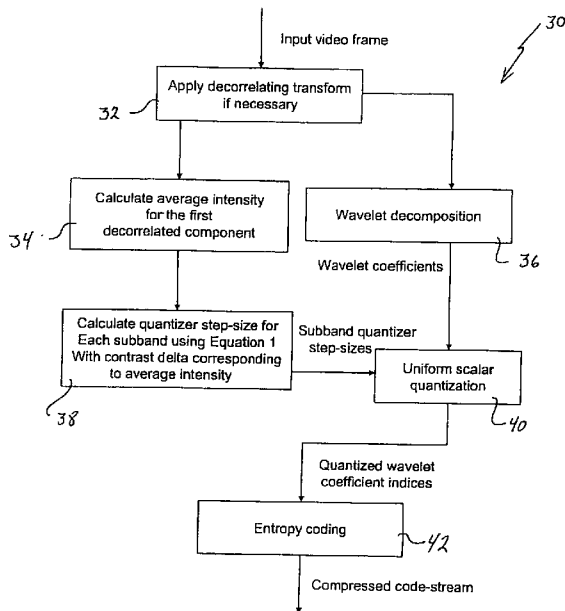


FIG. 3

(57) Abstract: According to one implementation, the present invention provides a method and apparatus to adapt the quantization steps-size used to quantize wavelet coefficients to the average brightness level of the corresponding pixels in a wavelet image or video coder. In another implementation, this method and apparatus produces a JPEG2000 Part 1 compliant code-stream.



WO 2010/077325 A2

**METHOD AND APPARATUS FOR ADAPTIVE QUANTIZATION OF  
SUBBAND/WAVELET COEFFICIENTS**

**CROSS REFERENCE TO RELATED APPLICATIONS**

5           This application claims priority from U.S. Provisional Patent Application  
Serial Nos. 61/203805 and 61/203807, both filed on December 29, 2008, the entire  
contents of which are incorporated herein by reference.

**TECHINICAL FIELD**

10           The present invention relates to image/video compression. More particularly,  
it relates to the quantization of wavelet coefficients in the compression of  
images/video.

**BACKGROUND**

15           When compressing an image or video frame using JPEG2000, in some  
scenarios, a goal is to achieve a certain visual quality without any restrictions on the  
compressed file size. One common way to achieve this is to use a two-dimensional  
contrast sensitivity function (2-D CSF) of the Human Visual System (HVS) as  
described in "Efficient JPEG2000 VBR compression with true constant quality," Paul  
20   W. Jones, SMPTE Technical Conference and Exhibition, Hollywood, CA, October  
2006. (hereinafter-referred to as "SMPTE – Paul W. Jones"). The entire contents of  
which is incorporated herein by reference. This method describes how to calculate  
the quantization step-size for each subband such that the resulting distortion in the  
reconstructed image or video frame is *just noticeable* under certain viewing  
25   conditions. The viewing conditions consist of parameters such as viewing distance,  
ambient light, display size, etc. The quantizer step-size calculated in this manner  
depends on the linear contrast produced on the displayed or projected image for one  
code value change in the subband domain. The contrast per code value varies  
depending on the average brightness level in the neighborhood of the contrast  
30   stimulus or the average brightness to which the observer is adapted. In the above  
paper, the authors approximate the contrast per codevalue by a constant value chosen  
from an appropriate mid-scale input level. But the observer may be adapted to  
different brightness levels for different frames. Additionally, the adaptation may be

different for different regions within an image or a frame. We describe a method to take this variation into account when determining the quantizer step-size.

### SUMMARY

5

According to an implementation, the method for compressing images or video frames using a wavelet encoder includes calculating an average intensity for each wavelet coefficient within a subband, and calculating a quantizer step size for each wavelet coefficient within the subband based on the calculated average intensity.

10

The method further includes performing wavelet decomposition to produce the wavelet coefficients, generating quantized wavelet coefficients using the calculated quantizer step sizes, and coding the quantized wavelet coefficients to produce a compressed video stream.

15

According to one implementation, the calculating of the average intensity includes applying a decorrelating transform for RGB or XYZ video frame, and calculating the average intensity on a first decorrelated component.

According to another implementation, the calculating of the average intensity is performed by calculating the average intensity from wavelet coefficients in subband

20 0.

According to yet a further implementation the compressing of images or video frames using a wavelet encoder is performed under the JPEG2000 standard, and further includes varying a default quantizer dead zone width for each subband, and storing the varied dead zone width information as a COM marker segment in a

25 JPEG2000 Part 1 file.

These and other aspects, features and advantages of the present principles will become apparent from the following detailed description of exemplary embodiments, which is to be read in connection with the accompanying drawings.

30

### BRIEF DESCRIPTION OF THE DRAWINGS

The present principles may be better understood in accordance with the following exemplary figures, in which:

Figure 1 is flow diagram of a method for wavelet encoding of images according to the prior art;

Figure 2 is graphical representation of the change in contrast (contrast delta) as a function of the codevalue for a subset of the valid codevalues according to an implementation;

Figure 3 shows a flow diagram of a method for compressing video frames using a wavelet encoder according to an implementation of the invention;

Figure 4 shows a frame representation with 3 levels of decomposition according to an implementation;

Figure 5 is flow diagram of a method for encoding video frames using a generic wavelet encoder according to an implementation of the invention;

Figure 6 is a flow diagram of a method for encoding video frames using a JPEG2000 encoder implementing the method of the present invention;

Figure 7 is a graphical representation of the reconstruction values in a Part 1 JPEG2000 quantizer in a JPEG2000 decoder showing dead-zone allocation according to the known JPEG2000 standard;

Figure 8 is a graphical representation of the reconstruction values in a Part 2 JPEG2000 quantizer in a JPEG2000 encoder showing dead-zone allocation according to an implementation of the present invention;

Figure 9 is a flow diagram of a method for changing of the default dead zone in an JPEG2000 encoder according to an implementation of the invention;

Figure 10 is a flow diagram of a method that a JPEG2000 decoder utilizes to change the default dead zone according to an implementation of the invention; and

Figure 11 is a block diagram of a standard video encoder as an example of a device implementing the present invention.

### DETAILED DESCRIPTION

The present principles are directed to image encoding and the adaptive quantization of wavelet coefficients designation and dead zone designation of the same. These principles can be applied to, and are shown in one embodiment to be directed to, JPEG2000 encoding.

The present description illustrates the present principles. It will thus be appreciated that those skilled in the art will be able to devise various arrangements that, although not explicitly described or shown herein, embody the present principles and are included within its spirit and scope.

5

All examples and conditional language recited herein are intended for pedagogical purposes to aid the reader in understanding the present principles and the concepts contributed by the inventor(s) to furthering the art, and are to be construed as being without limitation to such specifically recited examples and conditions.

10

Moreover, all statements herein reciting principles, aspects, and embodiments of the present principles, as well as specific examples thereof, are intended to encompass both structural and functional equivalents thereof. Additionally, it is intended that such equivalents include both currently known equivalents as well as equivalents developed in the future, i.e., any elements developed that perform the same function, regardless of structure.

15

Thus, for example, it will be appreciated by those skilled in the art that the block diagrams presented herein represent conceptual views of illustrative circuitry embodying the present principles. Similarly, it will be appreciated that any flow charts, flow diagrams, state transition diagrams, pseudocode, and the like represent various processes which may be substantially represented in computer readable media and so executed by a computer or processor, whether or not such computer or processor is explicitly shown.

25

The functions of the various elements shown in the figures may be provided through the use of dedicated hardware as well as hardware capable of executing software in association with appropriate software. When provided by a processor, the functions may be provided by a single dedicated processor, by a single shared processor, or by a plurality of individual processors, some of which may be shared. Moreover, explicit use of the term “processor” or “controller” should not be construed to refer exclusively to hardware capable of executing software, and may implicitly include, without limitation, digital signal processor (“DSP”) hardware, read-only

30

memory (“ROM”) for storing software, random access memory (“RAM”), and non-volatile storage.

Other hardware, conventional and/or custom, may also be included. Similarly,  
5 any switches shown in the figures are conceptual only. Their function may be carried out through the operation of program logic, through dedicated logic, through the interaction of program control and dedicated logic, or even manually, the particular technique being selectable by the implementer as more specifically understood from the context.

10

In the claims hereof, any element expressed as a means for performing a specified function is intended to encompass any way of performing that function including, for example, a) a combination of circuit elements that performs that function or b) software in any form, including, therefore, firmware, microcode or the  
15 like, combined with appropriate circuitry for executing that software to perform the function. The present principles as defined by such claims reside in the fact that the functionalities provided by the various recited means are combined and brought together in the manner which the claims call for. It is thus regarded that any means that can provide those functionalities are equivalent to those shown herein.

20

Reference in the specification to “one embodiment” or “an embodiment” of the present principles, as well as other variations thereof, means that a particular feature, structure, characteristic, and so forth described in connection with the embodiment is included in at least one embodiment of the present principles. Thus,  
25 the appearances of the phrase “in one embodiment” or “in an embodiment”, as well as any other variations, appearing in various places throughout the specification are not necessarily all referring to the same embodiment.

According to one implementation, the present invention describes a way to  
30 adapt the quantization steps-size used to quantize wavelet coefficients to the average brightness level of the corresponding pixels in a wavelet image or video coder. In another implementation, this method produces a JPEG2000 Part 1 compliant code-stream.

As mentioned above, the present invention improves on the known method for determining quantizer step-size for each subband for visually lossless JPEG2000 compression under certain viewing conditions.

5           Figure 1 shows a method 10 that a wavelet encoder can implement for encoding images according to the known prior art. First an image or a video frame undergoes wavelet decomposition/transformation 12 to produce wavelet coefficients. Next, the wavelet coefficients undergo uniform scalar quantization 14 (with or without dead zones). The scalar quantization may have a dead-zone, typically equal  
10   to twice the size of the quantizer step-size. The setting of the dead zones is discussed in further detail below. The resultant quantized coefficient indices undergo entropy coding 16 to produce a compressed code-stream. Most popular wavelet coders such as JPEG2000, use this basic structure.

15           Although we have described a generic wavelet image coding method in Figure 1, those skilled in the art will realize that it is equally applicable to methods used by image coders based on subband or wavelet packet decompositions. Another important point to note, is that in case of compressing video frames, the wavelet decomposition/transform may be applied to the prediction residual of a video frame  
20   after applying temporal prediction. In other cases, a motion adaptive or motion compensated 3D wavelet transform may be applied to a group of video frames to produce wavelet coefficients. The present invention is applicable to these scenarios as well.

25           One important problem for such wavelet encoders is to determine a quantizer step-size for each subband so as to guarantee a specific visual quality for the reconstructed image under certain viewing conditions. One example is for digital cinema applications. In this scenario, the viewing conditions such as viewing distance, display size and characteristics, ambient light, etc. are well controlled. One  
30   way to determine the quantization steps-size for each subband is proposed in the article discussed above in the background discussion. Those of skill in the art recognize that this method uses two-dimensional contrast sensitivity function (2-D CSF) of the human visual system (HVS). According to this method, the quantizer

step-size  $Q_b$  for a given subband  $b$  that produces *just noticeable* distortion in the reconstructed image can be calculated as

$$Q_b = \frac{[\Delta_b(l)] \cdot C_i(b)}{\Delta C_{cv}(l)}, \quad (1)$$

Where

- 5 •  $\Delta_b(l)$  is the quantizer step-size that in subband  $b$  that produces one codevalue change in the decompressed image.
- $C_i(b)$  is the threshold contrast for the observer for subband  $b$ . This is the Michelson contrast defined as

$$Contrast = \frac{L_{\max} - L_{\min}}{L_{\max} + L_{\min}} = \frac{\Delta L}{2 \cdot L_{\text{mean}}},$$

10 where  $L$  is luminance and  $\Delta L$  is the peak-to-peak luminance variation. It should be noted that the luminance is measured from a displayed or a projected image.

- $\Delta C_{cv}(l)$  is the contrast delta (change in contrast) on the display or projector for a one codevalue change in the decompressed image. The contrast delta is a function of the codevalue itself.

Figure 2 shows a graphical representation that plots the contrast delta as a function of code value for a subset of the valid code values. This is done for the digital cinema system as specified by Digital Cinema Initiative (DCI) specification which uses 12-bit XYZ colorspace and a gamma of 2.6 for the projector. As can be easily seen from Figure 2, the contrast delta per code value change changes significantly as a function of code value. In "SMPTE - Paul W. Jones" mentioned above, the authors approximate  $\Delta C_{cv}(l)$  by a single constant value corresponding to a mid-scale input code value. This prior art states that the observer is more likely to be adapted to this brightness level. However, the average brightness level may change from frame to frame in case of video. Also it may change locally from one region of an image to another. The present principles method uses the average brightness level of a neighborhood of the contrast stimulus to choose the correct  $\Delta C_{cv}(l)$  and vary the quantization step-size appropriately.

30

Figure 3 shows an embodiment of the method 30 of present invention for compressing video frames using a wavelet encoder. Initially, for each frame that is to be compressed, the average intensity value for all the pixels in the frame is calculated 34. For RGB or XYZ color images, prior to the calculation at step 34, a decorrelating transformation (32) is applied as specified in JPEG2000 Part 1. In case of RGB color image, this transformation yields YUV (or YCbCr) components. As such, the present invention always calculates 34 the average intensity based on the first component (Y) after any decorrelating transform is applied. Then, the quantizer step-size for each subband is calculated (38) using Equation (1), where  $\Delta C_{c,v}(I)$  corresponding to the average intensity for that frame is used. At step 36, the wavelet decomposition is performed to produce the wavelet coefficients, which are input into the uniform scalar quantization step 40. The Uniform Scaler quantization step 40 receives the subband quantizer step sizes and generates the quantized wavelet coefficients indices for entropy coding step 42. The result is the compressed code-stream.

Another embodiment of the present invention in the context of a generic wavelet encoder is described below. Consider an input image that has been wavelet transformed into subbands (after applying decorrelating transform if necessary). In this example, there are  $N_L$  levels of subband decomposition. Figure 4 shows an example with 3 levels of decomposition. An  $N_L$  level wavelet decomposition produces  $(3N_L + 1)$  subbands, where the subbands are indexed from 0 to  $3N_L$ , starting with the lowest frequency subband.

Figure 5 shows a flow diagram for this embodiment of the invention. First, if necessary, the decorrelating transform and  $N_L$  level wavelet decomposition is performed (52). Then, for subband 0, the  $N_L$  LL subband is scalar quantized (60) using a quantizer step-size determined from Equation (1) using a fixed  $\Delta C_{c,v}(I)$  for the entire subband. Then, the quantized wavelet coefficients from the LL subband 0 are used (at step 54) to calculate  $A_b(x,y)$  and derive the proper  $\Delta C_{c,v}(I)$ , while the quantizer step-sizes for the wavelet coefficients from remaining subbands are calculated as follows: A wavelet coefficient from subband  $b$  at level  $L$  is denoted by

$W_b(x, y)$ , where  $x$  and  $y$  denote the row and column indices within the subband grid. Now, a wavelet coefficient  $W_0(\hat{x}, \hat{y})$  from the  $N_l$  LL subband is associated with  $W_b(x, y)$  as follows.

$$\hat{x} = \left\lfloor \frac{x}{2^{(N_l-l)}} \right\rfloor, \hat{y} = \left\lfloor \frac{y}{2^{(N_l-l)}} \right\rfloor.$$

5 Here it is assumed that the image or video frame always starts at (0,0) and at each stage, the number of low pass filtered samples is greater than, or equal to, the number of high pass filtered samples. Now consider a neighborhood  $\Omega_0(\hat{x}, \hat{y})$  of the wavelet coefficient  $W_0(\hat{x}, \hat{y})$  in the  $N_l$  LL subband. The neighborhood  $\Omega_0(\hat{x}, \hat{y})$  is defined by 2 parameters,  $\delta_x$  and  $\delta_y$ , such that all the wavelet coefficients  $W_0(x, y)$  from  
 10 subband  $N_l$  LL, belonging to the neighborhood  $\Omega_0(\hat{x}, \hat{y})$  satisfy

$$|x - \hat{x}| \leq \delta_x, |y - \hat{y}| \leq \delta_y.$$

For subbands at different levels of decomposition, different values of  $\delta_x$  and  $\delta_y$  can be used. Then, for each coefficient  $W_b(x, y)$  from subband  $b$ , the average of the wavelet coefficients in  $\Omega_0(\hat{x}, \hat{y})$  for the first decorrelated component is calculated  
 15 (step 56) and denoted by  $A_b(x, y)$ . It is assumed that the wavelet analysis filters use a (1,1) normalization so that the nominal range of coefficients is the same as the range of input pixel values. The average of the wavelet coefficients  $A_b(x, y)$  is truncated to the valid range of codevalues, which in case of a 12-bit image is [0,4095]. If, before taking the wavelet transform, a DC value is subtracted from all the samples, it may be  
 20 necessary to add it back to the average  $A_b(x, y)$  before the truncation step. Then, the quantizer step-size for wavelet coefficient  $W_b(x, y)$  is calculated (56) using Equation (1) where  $\Delta C_{CV}(1)$  replaced by  $A_b(x, y)$  (suitably offset and truncated). It should be noted that since the quantized  $N_l$  LL subband coefficients are used for this calculation, the decoder can replicate these steps to derive the actual quantization step  
 25 size without any side information, provided that the compressed data corresponding to the  $N_l$  LL subband is included in its entirety before any compressed data from the other subbands. This also assumes that the relationship between contrast delta and codevalue is known to both the encoder and the decoder. Once calculated, each

wavelet coefficient from the other subband is quantized using the calculated step-size (58). Coding (62) can take place at this point once all wavelet coefficients have been quantized accordingly.

5 Those of skill in the art will recognize that there may be some difficulty in applying the above inventive concept to the JPEG2000 standard. The JPEG2000 standard mandates that the same quantizer step-size be used to quantize all the coefficients in a subband. The quantizer step-size can be varied by a power of 2 by discarding certain bit-planes or coding passes on a codeblock-by-codeblock basis. So  
 10 a slight modification of the method is needed to comply with the standard. First, analyzing the relationship between contrast delta and codevalues as shown in Figure 2, a contrast delta value of  $S_0$  is identified. The corresponding codevalue is denoted by  $C$ . Then, codevalues corresponding to  $S_n = 2^n S_0, n > 0$  are identified and denoted by  $CV_n$ . Here, it is assumed that the highest value of  $n$  is  $N$ . In Figure 2,  
 15  $S_0 = 0.5, S_1 = 1.0, S_2 = 2.0$ . In this example, the contrast deltas for codevalues below 500 are ignored, but they can be considered to find additional  $CV_n$  values if desired. Then, the codevalue threshold  $T_n$  is determined such that the contrast delta corresponding to codevalue of  $T_n$  is the average of contrast deltas for  $CV_n$  and  $CV_{n+1}$ .

20 The block diagram for JPEG2000 encoding method 70 according to an implementation of the present invention is shown in Figure 6. After decorrelating transform step 72 (if necessary), wavelet transformation is performed. In the quantization step 74, the quantizer step-size for each subband is determined using Equation (1), where contrast delta value corresponding to  $CV_N$  is used. Since for  
 25 smaller codevalues, the contrast delta is higher, this results in smaller quantizer step-sizes. The idea is to quantize with small step-sizes initially (step 76), and then, based on average intensity, determine whether in certain regions bit-planes can be discarded. This is accomplished as follows. For the  $N_L$  LL subband all bit-planes are encoded and retained in the final compressed code-stream. Typically the  $N_L$  LL subband is  
 30 very small. Hence, this has negligible impact on the overall bit-rate. Then, each codeblock  $B$  from each of the remaining subbands is associated with a set  $\Omega_B$  (step

78). The set  $\Omega_B$  consists of all the corresponding wavelet coefficients from the  $N_L$  LL subband for the coefficients in B. Then the average of the coefficients belonging to the set  $\Omega_B$ , denoted by  $A_B$ , is determined from the first decorrelated component. Then, two consecutive thresholds,  $T_{(n+1)}$  and  $T_n$  are found such that

5  $T_{n+1} \leq A_B < T_n$ . In that case,  $(N - (n + 1))$  bit-planes are discarded for codeblock B (step 80). If  $A_B < T_{N-1}$ , no bit-planes are discarded. If  $A_B \geq T_0$ ,  $N$  bit-planes are discarded for codeblock B (step 80). Then, these decisions regarding the discarding of the LSB bit-planes are passed on to the entropy coder 82 which produces a JPEG2000 compliant code-stream.

10

In addition to the method disclosed herein, it should be understood that hardware, software or any apparatus which performs these functions is also a part of the disclosed invention. Figure 11 shows a high level block diagram of a system 130 capable of implementing the above described methods of the invention. Although

15 shown as a stand along device, it is to be understood that this system 130 can be implemented as part of a multifunction, more complex device, such as, for example and any encoder, or a JPEG2000 compliant encoder. The system includes a processor 132 and one or more ROM memories 134, one or more RAM type memories 136 and a user interface 138 of any suitable known type (e.g., keyboard, mouse, touch screen,

20 etc.).

As mentioned above, the scalar quantization may have a dead-zone, typically equal to twice the size of the quantizer step-size. The following is a discussion of another implementation of the invention where “variable scalar quantization dead-

25 zones” feature from JPEG2000 Part 2 are incorporated into a JPEG2000 Part 1 compliant file. The main idea is to vary the default quantizer dead-zone width used in JPEG2000 Part 1, to improve the visual quality of the reconstructed images or video for certain textured regions and certain kind of imagery. One example is video with significant amount of film-grain. The present invention describes a way to store this

30 “dead-zone width” information as a COM marker segment inside a JPEG2000 Part 1 compliant file so that a JPEG2000 compliant decoder that is aware of this, can perform optimal dequantization to improve the visual quality of reconstructed images or video.

As is understood by those of skill in the art, the JPEG2000 compression standard mandates the use of a uniform quantizer that has a dead-zone around zero, to quantize the wavelet coefficients. Part 2 of the JPEG2000 standard allows the width of the dead-zone to vary for each subband, component, and tile. This results in better visual quality and sometimes, higher peak signal-to-noise ratio (PSNR), for certain textured regions and certain kind of imagery. One example of this is video frames with significant amount of film grain.

Unfortunately, there are hardly any existing JPEG2000 Part 2 implementations. On the other hand, due to adoption of JPEG2000, Part 1 by the Digital Cinema Initiative (DCI) committee (**The Digital Cinema Initiative (DCI) specification V1.0, July 2005**), the number of JPEG2000 Part 1 implementations is much higher. However, as noted above, Part 1 of the JPEG2000 standard uses a fixed dead-zone width that is equal to two times the quantization step-size. Thus, it is desirable to incorporate the capability to vary the dead-zone width while generating compressed files that are compliant with Part 1 of the JPEG2000 standard. The present implementation of the invention proposes a method to achieve this goal by using COM marker segment in a JPEG2000 Part 1 compliant file.

It should be noted that a JPEG2000 Part 1 compliant decoder that does not know how to parse or use the information stored in the COM marker segment, can still decode the compressed file, albeit at a higher distortion. But a JPEG2000 decoder that can take advantage of the COM marker segment information can perform optimal dequantization to improve the visual quality of the reconstructed images or video.

Part 1 of the JPEG2000 compression standard uses a uniform scalar quantizer with a dead-zone to quantize the wavelet coefficients as shown in Figure 7. In Figure 7, the quantizer step-size is  $\Delta$ . The range of input values that get quantized to quantizer bin 0 is referred to as the dead-zone. In this case, the size of the dead-zone is  $2\Delta$ . The vertical lines denote the boundaries of quantization intervals. The quantization rule is as follows:

$$q[n] = \text{sign}(y[n]) \left\lfloor \frac{|y[n]|}{\Delta} \right\rfloor, \quad (2)$$

where  $\lfloor \cdot \rfloor$  represents the truncation to the nearest integer towards zero. Here,  $y[n]$  represents the input sample and  $q[n]$  represents the corresponding quantizer index.

At the decoder, the reconstructed value,  $\hat{y}[n]$ , is generated using the dequantization

5 rule

$$\hat{y}[n] = \begin{cases} (q[n] + \gamma)\Delta, & \text{if } q[n] > 0, \\ (q[n] - \gamma)\Delta, & \text{if } q[n] < 0, \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Here  $0 \leq \gamma < 1$ , is a reconstruction parameter arbitrarily chosen by the JPEG2000 decoder. A value of  $\gamma = 0.50$ , which is the most commonly used, results in midpoint reconstruction. In Figure 7, we have assumed  $\gamma = 0.50$  for determining the

10 reconstruction values.

As mentioned above, the JPEG2000 standard does not mandate the use of a specific dead-zone on the encoder side, but a JPEG2000 Part 1 compliant decoder assumes that the JPEG2000 encoder has used a dead-zone of  $2\Delta$ . If the encoder uses

15 a different dead-zone, this can result in a mismatch between the encoder and the decoder resulting in higher distortion. A large dead-zone such as  $2\Delta$  has a disadvantage. If the input image contains flat areas with significant amount of film-grain, the wavelet coefficients corresponding to that area tend to have small magnitudes. Due to the large dead-zone, all the wavelet coefficients having small

20 non-zero magnitudes get quantized to zero. This has the effect of wiping out or introducing large distortions in the film-grain structure. This leads to visually annoying and objectionable artifacts.

To overcome this problem, in Part 2 of the JPEG2000 standard, the width of

25 the dead-zone can be varied from one subband to another. Figure 8 shows such a uniform scalar quantizer with a modified dead-zone of  $2(1 - \varepsilon)\Delta$  where  $-1 \leq \varepsilon < 1$ . In Figure 8,  $\varepsilon > 0$ , resulting in reduced width for the dead-zone compared to Part 1 of the JPEG2000 standard. For purposes of this disclosure, we refer to  $\varepsilon$  herein as a “dead zone modifier coefficient”.] This has the effect that some of the small

magnitude wavelet coefficients that were quantized to zero are now quantized to  $\pm 1$ . This provides better reconstruction for the flat areas with a significant amount of film-grain. In this case, the encoder uses the following quantization rule:

$$q[n] = \text{sign}(y[n]) \left\lfloor \frac{|y[n]|}{\Delta} + \varepsilon \right\rfloor. \quad (4)$$

5 The corresponding dequantization rule is

$$\hat{y}[n] = \begin{cases} (q[n] + \gamma - \varepsilon)\Delta, & \text{if } q[n] > 0, \\ (q[n] - \gamma + \varepsilon)\Delta, & \text{if } q[n] < 0, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

Here  $\gamma$  has the same interpretation as before. It should be noted that by reducing the width of the dead-zone, more samples are quantized to non-zero values. This leads to a lower distortion; but at the same time, may increase the bit-rate. Typically, when  
 10 using the modified dead-zone, it is desirable to use a higher quantizer step-size to achieve the same bit-rate as for the case of original dead-zone width of  $2\Delta$ . Thus, there is a trade-off between the reconstruction quality of the flat areas with significant film-grain and the rest of the image or video frame.

15 It should be noted that JPEG2000 Part 1 quantizer is a special case of JPEG2000 Part 2 quantizer, with  $\varepsilon = 0$ . Another point of note is that the dequantization rules for the Part 1 and Part 2 quantizers are identical except that the dequantization parameter  $\gamma$  is replaced with  $\gamma - \varepsilon$ . This means that a JPEG2000 Part 1 decoder can be used to dequantize the quantization indices generated by a  
 20 JPEG2000 Part 2 quantizer, provided the Part 1 decoder knows the value of  $\varepsilon$  used by the Part 1 quantizer. But JPEG2000 file format does not have any explicit provision for storing this information. In the absence of any information about  $\varepsilon$ , the JPEG2000 decoder is forced to use Equ. (3) as the dequantization rule, resulting in higher distortion. To overcome this, the present invention proposes to store the value  
 25 of  $\varepsilon$  in a COM segment marker in a JPEG2000 file. As in Part 2 of JPEG2000, the value of  $\varepsilon$  can be different for each tile, component, and subband.

In JPEG2000, comment (COM) marker segment provides a facility for including unstructured comment information in the code-stream. The first two bytes

comprise of the comment marker, **FF64<sub>h</sub>**. This is followed by a two byte parameter, LCOM, specifying the length of the comment marker segment, excluding the first two bytes. This is followed by a two byte parameter TY. TY = 0 means that the comment data is in binary format. TY = 1 means that the comment data is in the form of (Latin) character data. The TY parameter is followed by the actual comment data. In a preferred embodiment, the comment data is in the form of characters. The comment data consists of one or more groups. A group represents the  $\epsilon$  values for the subbands from a particular tile-component. A group consists of a number of fields as shown below in table 1, and as referred to in Figure 9. Figure 9 is refers to table 1 below, which provides some detail about how  $\epsilon$  value for each subband is stored in the COM marker. Different entries within a field and the fields themselves are separated by spaces.

Tile Index	Component index	A list of $\epsilon$ values	End of group symbol
Integer	Integer	Floating point	Integer

Table 1

A tile index of -1 signifies that the same  $\epsilon$  values will be used in all tiles. Similarly, a component index of -1 signifies that the same  $\epsilon$  values will be used in all components. The number of  $\epsilon$  values in a group is less than or equal to the number of subbands in that tile-component. The  $\epsilon$  values are listed starting with the highest frequency subband (1HH) and proceeding towards the lowest frequency subband (LL). If the number of entries is less than the number of subbands in that tile-component, the last  $\epsilon$  value is repeated for the remaining subbands. The end of group symbol is mandatory for every group except the last one.

Figure 9 shows the method 90 for changing the default dead zone in a JPEG2000 encoder according to an implementation of the invention. Initially, an input image or video frame is wavelet decomposed (92) into N subbands, thus generating wavelet coefficients grouped into N subbands. The generated wavelet coefficients are used at the scaler quantization step 94, along with quantization parameters  $\Delta_b$ ,  $\epsilon_b$  that are provided for each subband b, where  $0 \leq b < N$ . Here, the

uniform scalar quantization of subband  $b$  with step size  $\Delta_b$  and a dead zone of  $2(1 - \epsilon_b)$   $\Delta_b$  is performed. The outputs at this stage produce the indices for the quantized wavelet coefficients and entropy coding and JPEG200 tier-2 coding is performed (96) to generate the code stream. As shown, the COM marker segment is generated at step  
5 98 based on  $\epsilon_b, 0 \leq b < N$ . Finally the code stream is combined with the COM marker segment (100) and the JPEG200 Part 1 compliant bit stream is produced.

Figure 10 shows a decoding method in which the input is a JPEG2000 Part 1 compliant bit-stream and the first step is to extract the code stream and COM segment  
10 marker (112). Entropy decoding (114) is performed on the code stream, while  $\epsilon_b, 0 \leq b < N$  is extracted from the COM marker segment and used at the dequantization of the wavelet coefficient step (118). The output of the dequantization step 118 results in the reconstructed wavelet coefficients grouped into  $N$  subbands, and the inverse wavelet transform is applied  
15 120 to produce the reconstructed image or video.

These and other features and advantages of the present principles may be readily ascertained by one of ordinary skill in the pertinent art based on the teachings herein. It is to be understood that the teachings of the present principles may be implemented in various forms of hardware, software, firmware, special purpose  
20 processors, or combinations thereof.

Most preferably, the teachings of the present principles are implemented as a combination of hardware and software. Moreover, the software may be implemented as an application program tangibly embodied on a program storage unit. The  
25 application program may be uploaded to, and executed by, a machine comprising any suitable architecture. Preferably, the machine is implemented on a computer platform having hardware such as one or more central processing units ("CPU"), a random access memory ("RAM"), and input/output ("I/O") interfaces. The computer platform may also include an operating system and microinstruction code. The various  
30 processes and functions described herein may be either part of the microinstruction code or part of the application program, or any combination thereof, which may be executed by a CPU. In addition, various other peripheral units may be connected to the computer platform such as an additional data storage unit and a printing unit.

It is to be further understood that, because some of the constituent system components and methods depicted in the accompanying drawings are preferably implemented in software, the actual connections between the system components or the process function blocks may differ depending upon the manner in which the present principles are programmed. Given the teachings herein, one of ordinary skill in the pertinent art will be able to contemplate these and similar implementations or configurations of the present principles.

Although the illustrative embodiments have been described herein with reference to the accompanying drawings, it is to be understood that the present principles is not limited to those precise embodiments, and that various changes and modifications may be effected therein by one of ordinary skill in the pertinent art without departing from the scope or spirit of the present principles. All such changes and modifications are intended to be included within the scope of the present principles as set forth in the appended claims.

## CLAIMS

1. A method for compressing images or video frames using a wavelet encoder, the method comprising the steps of:
  - calculating an average intensity (34) for each wavelet coefficient within a subband;
  - calculating a quantizer step size (38) for each wavelet coefficient within the subband based on the calculated average intensity; and
  - performing encoding of wavelet coefficients using said quantizer step size.
  
2. The method of claim 1, said encoding of wavelet coefficients further comprising:
  - performing wavelet decomposition (36) to produce the wavelet coefficients;
  - generating quantized wavelet coefficients (40) using the calculated quantizer step sizes; and
  - coding (42) the quantized wavelet coefficients to produce a compressed video stream.
  
3. The method of claim 1, wherein said calculating an average intensity further comprises:
  - applying a decorrelating transform (32) for RGB or XYZ video frames; and
  - calculating the average intensity (34) on a first decorrelated component.
  
4. The method of claim 1, wherein the compressing is performed in accordance with the JPEG2000 standard, said method further comprising:
  - varying a default quantizer dead zone width for each subband; and
  - storing said varied dead zone width information as a COM marker segment in a JPEG2000 Part 1 file.
  
5. The method of claim 1, wherein said calculating the average intensity is performed by calculating the average intensity from wavelet coefficients in subband 0.

6. A method for compressing images or video frames using a wavelet encoder, the method comprising the steps of:

calculating an average intensity (54) for each wavelet coefficient in each of one or more subbands;

calculating a quantizer step size (56) for each wavelet coefficient using the calculated average intensity for the corresponding wavelet coefficient;

quantizing each wavelet coefficient (58) from each of the one or more subbands using the calculated step size; and

coding (62) quantized wavelet coefficients to produce a compressed code stream

7.,The method of claim 6, further comprising the step of performing uniform scalar quantization on a first of the one or more subbands using a fixed quantizer step size (60) to produce quantized wavelet coefficient indices.

8. The method of claim 7, wherein said coding (62) further comprising coding the quantized wavelet coefficients with the quantized wavelet coefficient indices.

9. A method for compressing images or video frames to produce a JPEG2000 part 1 compliant stream, the method comprising the steps of:

wavelet decomposing (92) of the input image or video frame into N subbands to produce wavelet coefficients grouped into N subbands;

performing uniform scalar quantization (94) of each subband with a predetermined quantization step size and dead zone parameter to produce indices for quantized wavelet coefficients;

entropy coding and JPEG2000 tier coding (96) of the indices for quantized wavelet coefficients to generate a code stream.

10. The method of claim 9, further comprising the steps of:

generating a COM marker segment (98) based on the dead zone parameter;

combining (100) the code stream and COM marker segment to produce the JPEG2000 Part 1 compliant bit-stream.

11. An apparatus for compressing images or video frames using a wavelet encoder, the method comprising the steps of:

means for calculating (130, 132) an average intensity for each wavelet coefficient within a subband ; and

means for calculating (130, 132) a quantizer step size for each wavelet coefficient within the subband based on the calculated average intensity.

12. The apparatus of claim 12, further comprising:

means for performing wavelet decomposition (130, 132) to produce the wavelet coefficients;

means for generating quantized wavelet coefficients (130, 132) using the calculated quantizer step sizes; and

means for coding (130, 132) the quantized wavelet coefficients to produce a compressed video stream.

13. The apparatus of claim 12, wherein said means for calculating an average intensity further comprises:

means for applying a decorrelating transform (130,132) for RGB or XYZ video frames; and

means for calculating the average intensity (130,132) on a first decorrelated component.

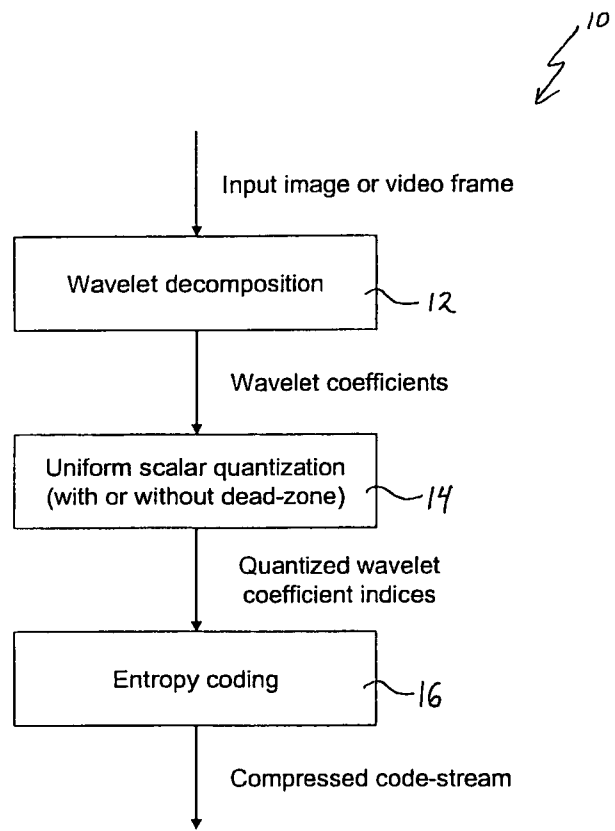
14. An apparatus for compressing images or video frames using a wavelet encoder, the apparatus comprising:

a processor (132) in signal communication with at least one memory device (134, 136), wherein said processor and said at least one memory device is configured to calculate an average intensity for each wavelet coefficient within a subband, and calculate a quantizer step size for each wavelet coefficient within the subband based on the calculated average intensity.

15. The apparatus of claim 14, wherein said processor and said at least one memory device are further configured to perform wavelet decomposition to produce the wavelet coefficients, generate quantized wavelet coefficients (40) using the calculated quantizer step sizes, and code the quantized wavelet coefficients to produce a compressed video stream.

16. The apparatus of claim 14, wherein during the calculation of the average intensity, said processor is further configured to apply a decorrelating transform for RGB or XYZ video frames, and calculate the average intensity on a first decorrelated component.

17. The apparatus of claim 14, wherein the compressing is performed under the JPEG2000 standard, and said processor varies a default quantizer dead zone width for each subband, and stores the varied dead zone width information as a COM marker segment in a JPEG2000 Part 1 file.



**FIG. 1**  
(PRIOR ART)




FIG. 2

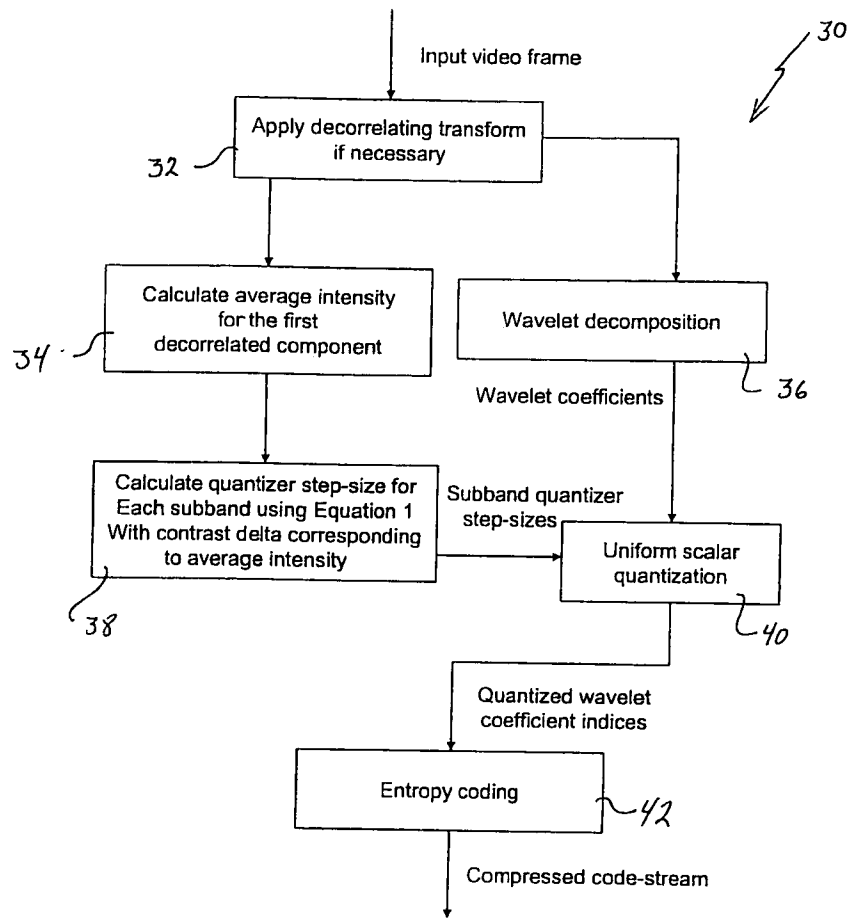


FIG. 3

3LL	3HL	2HL	1HL
3LH	3HH		
2LH	2HH		
1LH		1HH	

FIG. 4

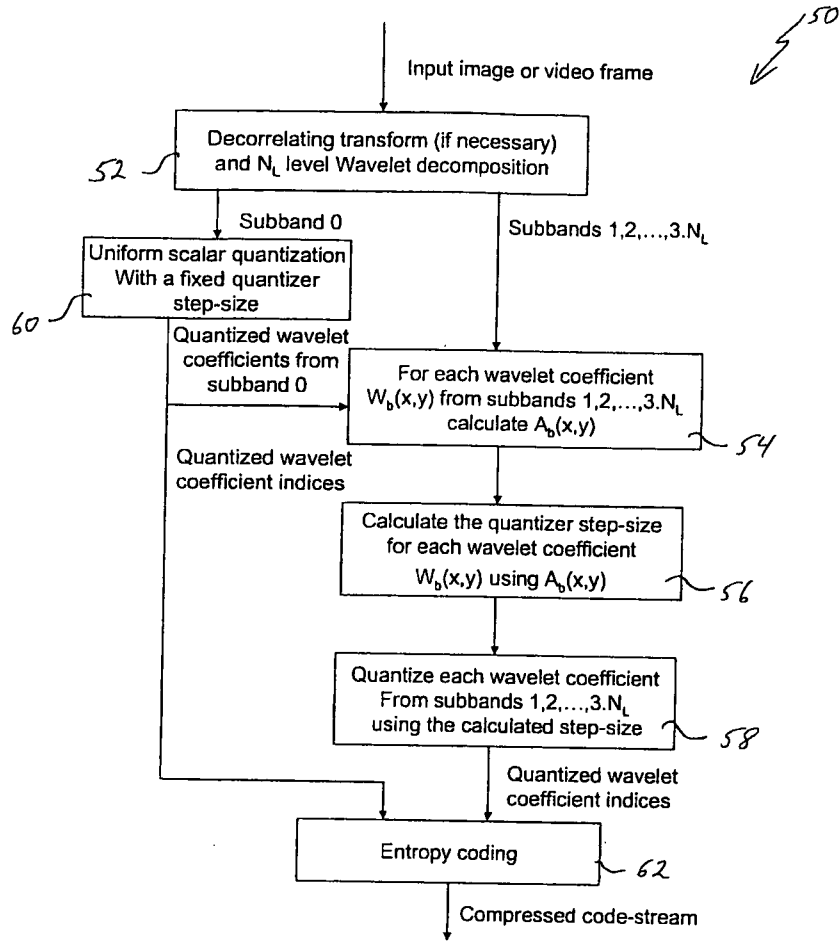


FIG. 5

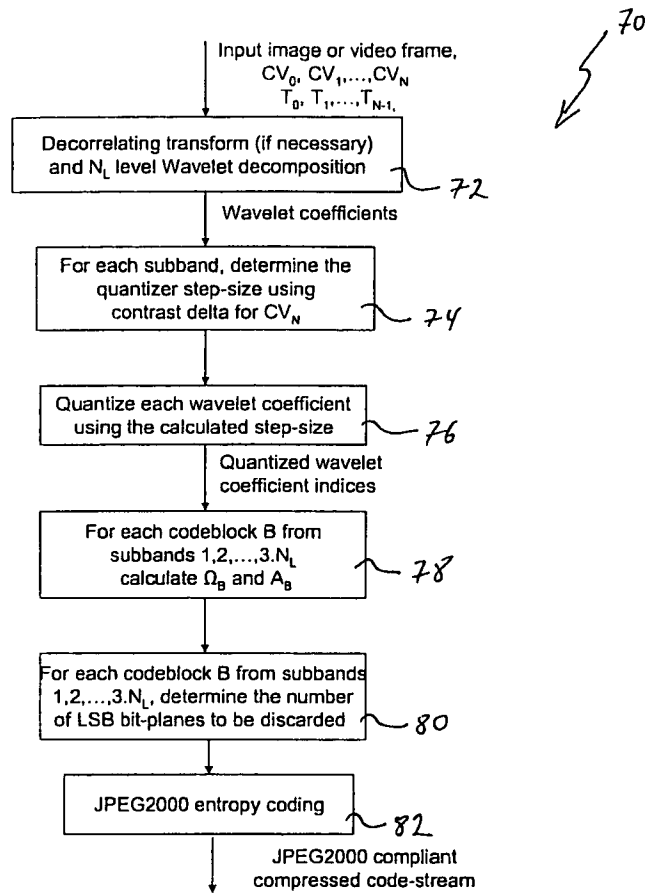


FIG. 6




FIG. 7




FIG. 8



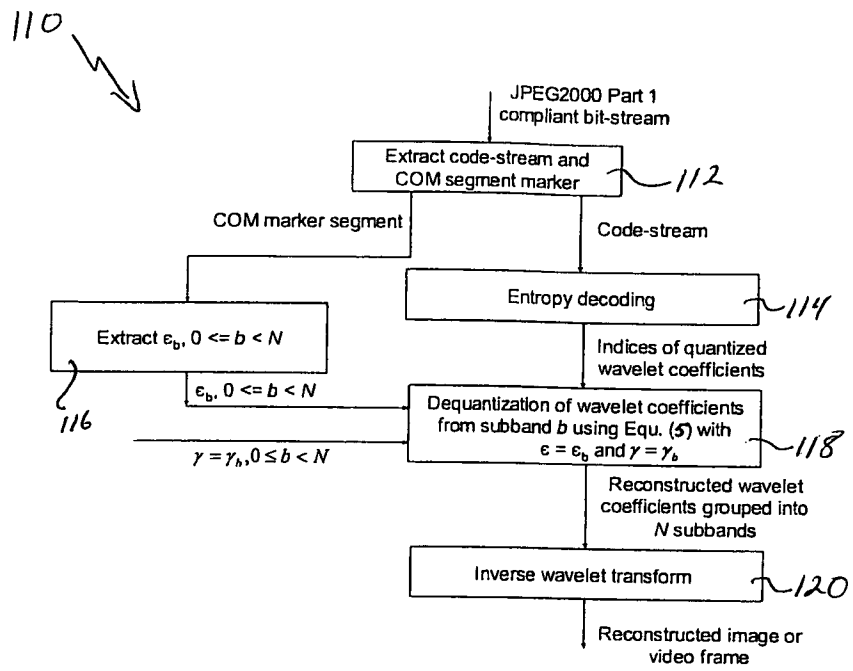


FIG. 10




FIG. 11