

(19) World Intellectual Property Organization  
International Bureau(43) International Publication Date  
18 October 2007 (18.10.2007)

PCT

(10) International Publication Number  
WO 2007/117707 A1

## (51) International Patent Classification:

H04N 7/24 (2006.01)

## (21) International Application Number:

PCT/US2007/008937

## (22) International Filing Date:

9 April 2007 (09.04.2007)

## (25) Filing Language:

English

## (26) Publication Language:

English

## (30) Priority Data:

11/400,788 7 April 2006 (07.04.2006) US

(71) Applicant (for all designated States except US): MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventors: ZUO, Xiping; One Microsoft Way, Redmond, Washington 98052-6399 (US). LIN, Chih-Lung; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH,

CN, CO, CR, CU, CZ, DE, DK, DM, 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, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

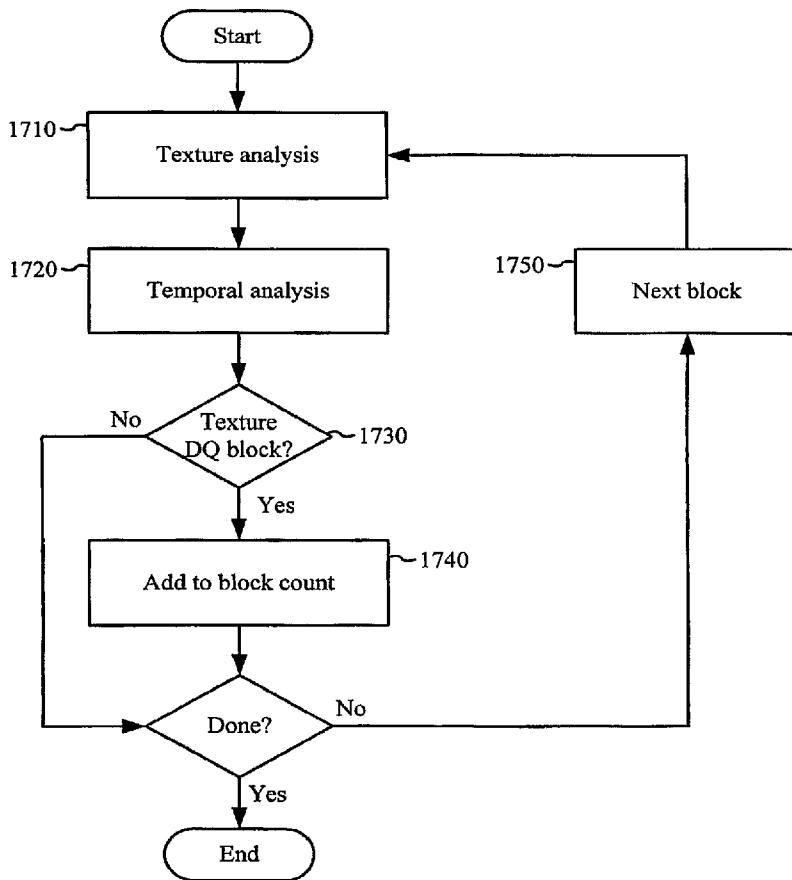
(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

## Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

[Continued on next page]

## (54) Title: QUANTIZATION ADJUSTMENT BASED ON TEXTURE LEVEL



(57) Abstract: A video encoder identifies a current smooth region of a current picture in a sequence and performs temporal analysis by determining whether a corresponding region in at least one previous and/or future picture is smooth. Based at least in part on the temporal analysis, the encoder adjusts quantization in the current smooth region. An encoder determines a differential quantization interval for a sequence, the interval comprising an interval number. The interval constrains the encoder to skip differential quantization for at least the interval number of predicted pictures after a predicted differentially quantized picture. An encoder analyzes texture in a current picture and sets a smoothness threshold. The encoder compares texture data with the smoothness threshold and adjusts differential quantization for at least part of the current picture based on a finding of at least one smooth region in the current picture according to the smoothness threshold.



**Published:**

- *with international search report*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

## QUANTIZATION ADJUSTMENT BASED ON TEXTURE LEVEL BACKGROUND

With the increased popularity of DVDs, music delivery over the Internet, and digital cameras, digital media have become commonplace. Engineers use a variety of 5 techniques to process digital audio, video, and images efficiently while still maintaining quality. To understand these techniques, it helps to understand how the audio, video, and image information is represented and processed in a computer.

### I. Representation of Media Information in a Computer

A computer processes media information as a series of numbers representing that 10 information. For example, a single number may represent the intensity of brightness or the intensity of a color component such as red, green or blue for each elementary small region of a picture, so that the digital representation of the picture consists of one or more arrays of such numbers. Each such number may be referred to as a *sample*. For a color image, it is conventional to use more than one sample to represent the color of each 15 elemental region, and typically three samples are used. The set of these samples for an elemental region may be referred to as a *pixel*, where the word "pixel" is a contraction referring to the concept of a "picture element." For example, one pixel may consist of three samples that represent the intensity of red, green and blue light necessary to represent the elemental region. Such a pixel type is referred to as an *RGB pixel*. Several 20 factors affect quality of media information, including sample depth, resolution, and frame rate (for video).

Sample depth is a property normally measured in bits that indicates the range of 25 numbers that can be used to represent a sample. When more values are possible for the sample, quality can be higher because the number can capture more subtle variations in intensity and/or a greater range of values. Resolution generally refers to the number of samples over some duration of time (for audio) or space (for images or individual video pictures). Images with higher resolution tend to look crisper than other images and contain more discernable useful details. Frame rate is a common term for temporal 30 resolution for video. Video with higher frame rate tends to mimic the smooth motion of natural objects better than other video, and can similarly be considered to contain more detail in the temporal dimension. For all of these factors, the tradeoff for high quality is the cost of storing and transmitting the information in terms of the bit rate necessary to represent the sample depth, resolution and frame rate, as Table 1 shows.

Bits Per Pixel (sample depth times samples per pixel)	Resolution (in pixels, Width x Height)	Frame Rate (in frames per second)	Bit Rate (in millions of bits per second)
8 (value 0-255, monochrome)	160x120	7.5	1.2
24 (value 0-255, RGB)	320x240	15	27.6
24 (value 0-255, RGB)	640x480	30	221.2
24 (value 0-255, RGB)	1280x720	60	1327.1

**Table 1: Bit rates for different quality levels of raw video**

Despite the high bit rate necessary for storing and sending high quality video (such as HDTV), companies and consumers increasingly depend on computers to create, distribute, and play back high quality content. For this reason, engineers use compression 5 (also called source coding or source encoding) to reduce the bit rate of digital media. Compression decreases the cost of storing and transmitting the information by converting the information into a lower bit rate form. Compression can be lossless, in which quality of the video does not suffer but decreases in bit rate are limited by the complexity of the video. Or, compression can be lossy, in which quality of the video suffers but decreases in 10 bit rate are more dramatic. Decompression (also called decoding) reconstructs a version of the original information from the compressed form. A “codec” is an encoder/decoder system.

In general, video compression techniques include “intra” compression and “inter” or predictive compression. For video frames, intra compression techniques compress 15 individual frames, typically called I-frames or key frames. Inter compression techniques compress frames with reference to preceding and/or following frames, and inter-compressed frames are typically called predicted frames, P-frames, or B-frames.

## **II. Inter and Intra Compression in Windows Media Video, Versions 8 and 9**

Microsoft Corporation’s Windows Media Video, Version 8 [“WMV8”] includes a 20 video encoder and a video decoder. The WMV8 encoder uses intra and inter compression, and the WMV8 decoder uses intra and inter decompression. Windows Media Video, Version 9 [“WMV9”] uses a similar architecture for many operations.

### **A. Intra Compression**

Figure 1 illustrates block-based intra compression 100 of a block 105 of samples in 25 a key frame in the WMV8 encoder. A block is a set of samples, for example, an 8x8 arrangement of samples. The WMV8 encoder splits a key video frame into 8x8 blocks and applies an 8x8 Discrete Cosine Transform [“DCT”] 110 to individual blocks such as

the block 105. A DCT is a type of frequency transform that converts the 8x8 block of samples (spatial information) into an 8x8 block of DCT coefficients 115, which are frequency information. The DCT operation itself is lossless or nearly lossless. Compared to the original sample values, however, the DCT coefficients are more efficient for the 5 encoder to compress since most of the significant information is concentrated in low frequency coefficients (conventionally, the upper left of the block 115) and many of the high frequency coefficients (conventionally, the lower right of the block 115) have values of zero or close to zero.

The encoder then quantizes 120 the DCT coefficients, resulting in an 8x8 block of 10 quantized DCT coefficients 125. Quantization is lossy. Since low frequency DCT coefficients tend to have higher values, quantization typically results in loss of precision but not complete loss of the information for the coefficients. On the other hand, since high frequency DCT coefficients tend to have values of zero or close to zero, quantization of the high frequency coefficients typically results in contiguous regions of zero values. In 15 addition, in some cases high frequency DCT coefficients are quantized more coarsely than low frequency DCT coefficients, resulting in greater loss of precision/information for the high frequency DCT coefficients.

The encoder then prepares the 8x8 block of quantized DCT coefficients 125 for 20 entropy encoding, which is a form of lossless compression. The exact type of entropy encoding can vary depending on whether a coefficient is a DC coefficient (lowest frequency), an AC coefficient (other frequencies) in the top row or left column, or another AC coefficient.

The encoder encodes the DC coefficient 126 as a differential from the DC 25 coefficient 136 of a neighboring 8x8 block, which is a previously encoded neighbor (e.g., top or left) of the block being encoded. (Figure 1 shows a neighbor block 135 that is situated to the left of the block being encoded in the frame.) The encoder entropy encodes 140 the differential.

The entropy encoder can encode the left column or top row of AC coefficients as a differential from a corresponding left column or top row of the neighboring 8x8 block. 30 This is an example of AC coefficient prediction. Figure 1 shows the left column 127 of AC coefficients encoded as a differential 147 from the left column 137 of the neighboring (in reality, to the left) block 135. The differential coding increases the chance that the differential coefficients have zero values. The remaining AC coefficients are from the block 125 of quantized DCT coefficients.

The encoder scans 150 the 8x8 block 145 of quantized AC DCT coefficients into a one-dimensional array 155 and then entropy encodes the scanned AC coefficients using a variation of run length coding 160. The encoder selects an entropy code from one or more run/level/last tables 165 and outputs the entropy code.

5        **B.     Inter Compression**

Inter compression in the WMV8 encoder uses block-based motion compensated prediction coding followed by transform coding of the residual error. Figures 2 and 3 illustrate the block-based inter compression for a predicted frame in the WMV8 encoder. In particular, Figure 2 illustrates motion estimation for a predicted frame 210 and Figure 3 10 illustrates compression of a prediction residual for a motion-compensated block of a predicted frame.

For example, in Figure 2, the WMV8 encoder computes a motion vector for a macroblock 215 in the predicted frame 210. To compute the motion vector, the encoder searches in a search area 235 of a reference frame 230. Within the search area 235, the 15 encoder compares the macroblock 215 from the predicted frame 210 to various candidate macroblocks in order to find a candidate macroblock that is a good match. The encoder outputs information specifying the motion vector (entropy coded) for the matching macroblock. The motion vector is differentially coded with respect to a motion vector predictor.

20        After reconstructing the motion vector by adding the differential to the motion vector predictor, a decoder uses the motion vector to compute a prediction macroblock for the macroblock 215 using information from the reference frame 230, which is a previously reconstructed frame available at the encoder and the decoder. The prediction is rarely perfect, so the encoder usually encodes blocks of pixel differences (also called the error or 25 residual blocks) between the prediction macroblock and the macroblock 215 itself.

Figure 3 illustrates an example of computation and encoding of an error block 335 in the WMV8 encoder. The error block 335 is the difference between the predicted block 315 and the original current block 325. The encoder applies a DCT 340 to the error block 335, resulting in an 8x8 block 345 of coefficients. The encoder then quantizes 350 the 30 DCT coefficients, resulting in an 8x8 block of quantized DCT coefficients 355. The encoder scans 360 the 8x8 block 355 into a one-dimensional array 365 such that coefficients are generally ordered from lowest frequency to highest frequency. The encoder entropy encodes the scanned coefficients using a variation of run length coding

370. The encoder selects an entropy code from one or more run/level/last tables 375 and outputs the entropy code.

Figure 4 shows an example of a corresponding decoding process 400 for an inter-coded block. In summary of Figure 4, a decoder decodes (410, 420) entropy-coded information representing a prediction residual using variable length decoding 410 with one or more run/level/last tables 415 and run length decoding 420. The decoder inverse scans 430 a one-dimensional array 425, storing the entropy-decoded information into a two-dimensional block 435. The decoder inverse quantizes and inverse DCTs (together, 440) the data, resulting in a reconstructed error block 445. In a separate motion compensation path, the decoder computes a predicted block 465 using motion vector information 455 for displacement from a reference frame. The decoder combines 470 the predicted block 465 with the reconstructed error block 445 to form the reconstructed block 475. An encoder also performs the inverse quantization, inverse DCT, motion compensation and combining to reconstruct frames for use as reference frames.

### 15 III. Lossy Compression and Quantization

The preceding section mentioned quantization, a mechanism for lossy compression, and entropy coding, also called lossless compression. Lossless compression reduces the bit rate of information by removing redundancy from the information without any reduction in fidelity. For example, a series of ten consecutive pixels that are all exactly the same shade of red could be represented as a code for the particular shade of red and the number ten as a "run length" of consecutive pixels, and this series can be perfectly reconstructed by decompression from the code for the shade of red and the indicated number (ten) of consecutive pixels having that shade of red. Lossless compression techniques reduce bit rate at no cost to quality, but can only reduce bit rate up to a certain point. Decreases in bit rate are limited by the inherent amount of variability in the statistical characterization of the input data, which is referred to as the source entropy.

30 In contrast, with lossy compression, the quality suffers somewhat but the achievable decrease in bit rate is more dramatic. For example, a series of ten pixels, each being a slightly different shade of red, can be approximated as ten pixels with exactly the same particular approximate red color. Lossy compression techniques can be used to reduce bit rate more than lossless compression techniques, but some of the reduction in bit rate is achieved by reducing quality, and the lost quality cannot be completely recovered. Lossy compression is often used in conjunction with lossless compression – in a system design in which the lossy compression establishes an approximation of the information

and lossless compression techniques are applied to represent the approximation. For example, the series of ten pixels, each a slightly different shade of red, can be represented as a code for one particular shade of red and the number ten as a run-length of consecutive pixels. In general, an encoder varies quantization to trade off quality and bit rate. Coarser 5 quantization results in greater quality reduction but allows for greater bit rate reduction. In decompression, the original series would then be reconstructed as ten pixels with the same approximated red color.

According to one possible definition, quantization is a term used for an approximating non-reversible mapping function commonly used for lossy compression, in 10 which there is a specified set of possible output values, and each member of the set of possible output values has an associated set of input values that result in the selection of that particular output value. A variety of quantization techniques have been developed, including scalar or vector, uniform or non-uniform, and adaptive or non-adaptive quantization.

#### 15 A. Scalar Quantizers

According to one possible definition, a scalar quantizer is an approximating functional mapping  $x \rightarrow Q[x]$  of an input value  $x$  to a quantized value  $Q[x]$ , sometimes called a reconstructed value. Figure 5 shows a "staircase" I/O function 500 for a scalar quantizer. The horizontal axis is a number line for a real number input variable  $x$ , and the 20 vertical axis indicates the corresponding quantized values  $Q[x]$ . The number line is partitioned by thresholds such as the threshold 510. Each value of  $x$  within a given range between a pair of adjacent thresholds is assigned the same quantized value  $Q[x]$ . For example, each value of  $x$  within the range 520 is assigned the same quantized value 530. (At a threshold, one of the two possible quantized values is assigned to an input  $x$ , 25 depending on the system.) Overall, the quantized values  $Q[x]$  exhibit a discontinuous, staircase pattern. The distance the mapping continues along the number line depends on the system, typically ending after a finite number of thresholds. The placement of the thresholds on the number line may be uniformly spaced (as shown in Figure 5) or non-uniformly spaced.

30 A scalar quantizer can be decomposed into two distinct stages. The first stage is the classifier stage, in which a classifier function mapping  $x \rightarrow A[x]$  maps an input  $x$  to a quantization index  $A[x]$ , which is often integer-valued. In essence, the classifier segments an input number line or data set. Figure 6A shows a generalized classifier 600 and thresholds for a scalar quantizer. As in Figure 5, a number line for a real number variable

$x$  is segmented by thresholds such as the threshold 610. Each value of  $x$  within a given range such as the range 620 is assigned the same quantized value  $Q[x]$ . Figure 6B shows a numerical example of a classifier 650 and thresholds for a scalar quantizer.

5 In the second stage, a reconstructor functional mapping  $k \rightarrow \beta[k]$  maps each quantization index  $k$  to a reconstruction value  $\beta[k]$ . In essence, the reconstructor places steps having a particular height relative to the input number line segments (or selects a subset of data set values) for reconstruction of each region determined by the classifier. The reconstructor functional mapping may be implemented, for example, using a lookup table. Overall, the classifier relates to the reconstructor as follows:

10 
$$Q[x] = \beta[A[x]] \quad (1)$$

In common usage, the term “quantization” is often used to describe the classifier stage, which is performed during encoding. The term “inverse quantization” is similarly used to describe the reconstructor stage, whether performed during encoding or decoding.

15 The distortion introduced by using such a quantizer may be computed with a difference-based distortion measure  $d(x-Q[x])$ . Typically, such a distortion measure has the property that  $d(x-Q[x])$  increases as  $x-Q[x]$  deviates from zero; and typically each reconstruction value lies within the range of the corresponding classification region, so that the straight line that would be formed by the functional equation  $Q[x] = x$  will pass through every step of the staircase diagram (as shown in Figure 5) and therefore  $Q[Q[x]]$  20 will typically be equal to  $Q[x]$ . In general, a quantizer is considered better in rate-distortion terms if the quantizer results in a lower average value of distortion than other quantizers for a given bit rate of output. More formally, a quantizer is considered better if, for a source random variable  $X$ , the expected (i.e., the average or statistical mean) value of the distortion measure  $\bar{D} = E_X\{d(X-Q[X])\}$  is lower for an equal or lower entropy  $H$  of 25  $A[X]$ . The most commonly-used distortion measure is the squared error distortion measure, for which  $d(|x-y|) = |x-y|^2$ . When the squared error distortion measure is used, the expected value of the distortion measure ( $\bar{D}$ ) is referred to as the mean squared error.

#### B. Dead Zone + Uniform Threshold Quantizers

30 A non-uniform quantizer has threshold values that are not uniformly spaced for all classifier regions. According to one possible definition, a dead zone plus uniform threshold quantizer [“DZ+UTQ”] is a quantizer with uniformly spaced threshold values for all classifier regions except the one containing the zero input value (which is called the dead zone [“DZ”]). In a general sense, a DZ+UTQ is a non-uniform quantizer, since the DZ size is different than the other classifier regions.

A DZ+UTQ has a classifier index mapping rule  $x \rightarrow A[x]$  that can be expressed based on two parameters. Figure 7 shows a staircase I/O function 700 for a DZ+UTQ, and Figure 8A shows a generalized classifier 800 and thresholds for a DZ+UTQ. The parameter  $s$ , which is greater than 0, indicates the step size for all steps other than the DZ.

5 Mathematically, all  $s_i$  are equal to  $s$  for  $i \neq 0$ . The parameter  $z$ , which is greater than or equal to 0, indicates the ratio of the DZ size to the size of the other steps. Mathematically,  $s_0 = z \cdot s$ . In Figure 8A,  $z$  is 2, so the DZ is twice as wide as the other classification zones. The index mapping rule  $x \rightarrow A[x]$  for a DZ+UTQ can be expressed as:

$$A[x] = \text{sign}(x) * \max\left(0, \left\lfloor \frac{|x|}{s} - \frac{z}{2} + 1 \right\rfloor\right) \quad (2),$$

10 where  $\lfloor \cdot \rfloor$  denotes the smallest integer less than or equal to the argument and where  $\text{sign}(x)$  is the function defined as:

$$\text{sign}(x) = \begin{cases} +1, & \text{for } x \geq 0, \\ -1, & \text{for } x < 0. \end{cases} \quad (3).$$

Figure 8B shows a numerical example of a classifier 850 and thresholds for a DZ+UTQ with  $s = 1$  and  $z = 2$ . Figures 5, 6A, and 6B show a special case DZ+UTQ with 15  $z = 1$ . Quantizers of the UTQ form have good performance for a variety of statistical sources. In particular, the DZ+UTQ form is optimal for the statistical random variable source known as the Laplacian source.

20 In some system designs (not shown), an additional consideration may be necessary to fully characterize a DZ+UTQ classification rule. For practical reasons there may be a need to limit the range of values that can result from the classification function  $A[x]$  to some reasonable finite range. This limitation is referred to as clipping. For example, in some such systems the classification rule could more precisely be defined as:

$$A[x] = \text{sign}(x) * \min\left[g, \max\left(0, \left\lfloor \frac{|x|}{s} - \frac{z}{2} + 1 \right\rfloor\right)\right] \quad (4),$$

where  $g$  is a limit on the absolute value of  $A[x]$ .

25 Different reconstruction rules may be used to determine the reconstruction value for each quantization index. Standards and product specifications that focus only on achieving interoperability will often specify reconstruction values without necessarily specifying the classification rule. In other words, some specifications may define the functional mapping  $k \rightarrow \beta[k]$  without defining the functional mapping  $x \rightarrow A[x]$ . This 30 allows a decoder built to comply with the standard/specification to reconstruct information

correctly. In contrast, encoders are often given the freedom to change the classifier in any way that they wish, while still complying with the standard/specification.

Numerous systems for adjusting quantization thresholds have been developed. Many standards and products specify reconstruction values that correspond to a typical 5 mid-point reconstruction rule (e.g., for a typical simple classification rule) for the sake of simplicity. For classification, however, the thresholds can in fact be adjusted so that certain input values will be mapped to more common (and hence, lower bit rate) indices, which makes the reconstruction values closer to optimal.

In many systems, the extent of quantization is measured in terms of quantization 10 step size. Coarser quantization uses larger quantization step sizes, corresponding to wider ranges of input values. Finer quantization uses smaller quantization step sizes. Often, for purposes of signaling and reconstruction, quantization step sizes are parameterized as multiples of a smallest quantization step size.

### **C. Quantization Artifacts**

15 As mentioned above, lossy compression tends to cause a decrease in quality. For example, a series of ten samples of slightly different values can be approximated using quantization as ten samples with exactly the same particular approximate value. This kind of quantization can reduce the bit rate of encoding the series of ten samples, but at the cost of lost detail in the original ten samples.

20 In some cases, quantization produces visible artifacts that tend to be more artificial-looking and visually distracting than simple loss of fine detail. For example, smooth, un-textured content is susceptible to contouring artifacts – artifacts that appear between regions of two different quantization output values – because the human visual system is sensitive to subtle variations (particularly luma differences) in smooth content.

25 Using the above example, consider a case where the luma values of the series of ten samples change gradually and consistently from the first sample to the tenth sample. Quantization may approximate the first five sample values as one value and the last five sample values as another value. While this kind of quantization may not create visible artifacts in textured areas due to masking effects, in smooth regions it can create a visible 30 line or step in the reconstructed image between the two sets of five samples.

### **IV. Differential Quantization in VC-1**

In differential quantization, an encoder varies quantization step sizes (also referred to herein as quantization parameters or QPs in some implementations) for different parts of a picture. Typically, this involves varying QPs on a macroblock level or other sub-

picture level. The encoder makes decisions on how to vary the QPs, and signals those decisions, as appropriate, to a decoder.

For example, a VC-1 encoder optionally chooses differential quantization for compression. The encoder sends a bitstream element (DQUANT) at a syntax level above picture level to indicate whether or not the QP can vary among the macroblocks in individual pictures. The encoder sends a picture-level bitstream element, PQINDEX, to indicate a picture QP. If DQUANT = 0, the QP indicated by PQINDEX is used for all macroblocks in the picture. If DQUANT = 1 or 2, different macroblocks in the same picture can use different QPs.

The VC-1 encoder can use more than one approach to differential quantization. In one approach, only two different QPs are used for a picture. This is referred to as bi-level differential quantization. For example, one QP is used for macroblocks at picture edges and another QP is used for macroblocks in the rest of the picture. This can be useful for saving bits at picture edges, where fine detail is less important for maintaining overall visual quality. Or, a 1-bit value signaled per macroblock indicates which of two available QP values to use for the macroblock. In another approach, referred to as multi-level differential quantization, a larger number of different QPs can be used for individual macroblocks in a picture.

The encoder sends a picture-level bitstream element, VOPDQUANT, when DQUANT is non-zero. VOPDQUANT is composed of other elements, potentially including DQPROFILE, which indicates which parts of the picture can use QPs other than the picture QP. When DQPROFILE indicates that arbitrary, different macroblocks can use QPs other than the picture QP, the bitstream element DQBILEVEL is present. If DQBILEVEL = 1, each macroblock uses one of two QPs (bi-level quantization). If DQBILEVEL = 0, each macroblock can use any QP (multi-level quantization).

The bitstream element MQDIFF is sent at macroblock level to signal a 1-bit selector for a macroblock for bi-level quantization. For multi-level quantization, MQDIFF indicates a differential between the picture QP and the macroblock QP or escape-coded absolute QP for a macroblock.

## 30 **V. Other Standards and Products**

Numerous international standards specify aspects of video decoders and formats for compressed video information. Directly or by implication, these standards also specify certain encoder details, but other encoder details are not specified. Some standards address still image compression/decompression, and other standards address audio

compression/decompression. Numerous companies have produced encoders and decoders for audio, still images, and video. Various other kinds of signals (for example, hyperspectral imagery, graphics, text, financial information, etc.) are also commonly represented and stored or transmitted using compression techniques.

5 Various video standards allow the use of different quantization step sizes for different picture types, and allow variation of quantization step sizes for rate and quality control.

10 Standards typically do not fully specify the quantizer design. Most allow some variation in the encoder classification rule  $x \rightarrow A[x]$  and/or the decoder reconstruction rule  $10 k \rightarrow \beta[k]$ . The use of a DZ ratio  $z = 2$  or greater has been implicit in a number of encoding designs. For example, the spacing of reconstruction values for predicted regions in some standards implies use of  $z \geq 2$ . Reconstruction values in these examples from standards are spaced appropriately for use of DZ+UTQ classification with  $z = 2$ . Designs based on  $z = 1$  (or at least  $z < 2$ ) have been used for quantization in several standards. In 15 these cases, reconstruction values are equally spaced around zero and away from zero.

Given the critical importance of video compression to digital video, it is not surprising that video compression is a richly developed field. Whatever the benefits of previous video compression techniques, however, they do not have the advantages of the following techniques and tools.

## 20 SUMMARY

The present application describes techniques and tools for adjusting quantization based on texture levels in video. For example, a video encoder improves the perceptual quality of video using adaptive smoothness thresholds for smooth regions and temporal analysis of smooth regions when allocating bits during encoding.

25 This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter.

30 In one aspect, a video encoder identifies a current smooth region (e.g., a gradient slope region) of a current video picture in a video picture sequence. The sequence has a display order in which display of previous pictures occurs before display of the current picture and display of future pictures occurs after the display of the current picture. The encoder performs temporal analysis on the current smooth region. The temporal analysis comprises determining whether a corresponding region in at least one of the previous

and/or future pictures is smooth. Based at least in part on the temporal analysis, the encoder adjusts quantization (e.g., reducing a quantization step size to avoid introduction of a contouring artifact) in the current smooth region. The temporal analysis can be performed on an adjustable number of future pictures. The temporal analysis can attribute 5 greater weight to smoothness in a corresponding region of a future picture that is temporally nearer to the current picture. The encoder can analyze a texture map to identify the current smooth region, and can analyze mean luminance values in the temporal analysis.

In another aspect, a video encoder determines a differential quantization interval 10 (e.g., a fixed or adaptively adjustable interval) for a video picture sequence, the interval comprising an interval number. The encoder uses the interval when performing differential quantization for predicted differentially quantized pictures in the sequence. The interval constrains the encoder to skip differential quantization for at least the interval number of predicted pictures after one of the predicted differentially quantized pictures. 15 For example, the differential quantization comprises selecting one or more differential quantization step sizes for a first predicted picture and changing the one or more differential quantization step sizes for a second predicted picture, where the one or more differential quantization step sizes differ from a picture quantization step size for the first predicted picture, and the second predicted picture is outside the interval from the first 20 predicted picture.

In another aspect, a video encoder analyzes texture in a current video picture (e.g., by analyzing a texture map) and sets a smoothness threshold for the current picture based 25 at least in part on the analyzed texture. The encoder compares texture data of the current picture with the smoothness threshold and adjusts differential quantization for at least part of the current picture based on a finding of at least one smooth region in the current picture according to the smoothness threshold. The encoder can analyze texture by applying a sliding window to a gradient value histogram of block gradient values. To 30 adjust differential quantization, the encoder can determine a percentage of flat blocks in the current picture and compare the percentage to one or more percentage thresholds. Or, the encoder can identify an isolated flat block in a texture region in the current picture and skip differential quantization for the isolated flat block.

The foregoing and other objects, features, and advantages will become more apparent from the following detailed description, which proceeds with reference to the accompanying figures.

## BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a diagram showing block-based intraframe compression of an 8x8 block of samples.

Figure 2 is a diagram showing motion estimation in a video encoder.

5 Figure 3 is a diagram showing block-based compression for an 8x8 block of prediction residuals in a video encoder.

Figure 4 is a diagram showing block-based decompression for an 8x8 block of prediction residuals in a video decoder.

Figure 5 is a chart showing a staircase I/O function for a scalar quantizer.

10 Figures 6A and 6B are charts showing classifiers and thresholds for scalar quantizers.

Figure 7 is a chart showing a staircase I/O function for a DZ+UTQ.

Figures 8A and 8B are charts showing classifiers and thresholds for DZ+UTQs.

15 Figure 9 is a block diagram of a suitable computing environment in conjunction with which several described embodiments may be implemented.

Figure 10 is a block diagram of a generalized video encoder system in conjunction with which several described embodiments may be implemented.

Figure 11 is a diagram of a macroblock format used in several described embodiments.

20 Figure 12 is a flow chart of an adaptive video encoding method.

Figure 13 is a diagram showing computation of a pixel gradient using luminance and chrominance data for a block.

Figure 14 is a histogram graph of plural pixel gradients for the block of Figure 13.

Figure 15 is a graph of an example block value characterization framework.

25 Figure 16 is a flow chart showing a generalized technique for applying differential quantization based on texture information.

Figure 17 is a flow chart showing a technique for using temporal analysis to make texture DQ decisions.

30 Figure 18 is a flow chart showing a technique for making a texture DQ decision using percentage thresholds and isolated smooth block filtering.

Figure 19 is a flow chart showing a technique for selectively adjusting texture level thresholds for high-texture pictures.

Figure 20 is a code diagram showing example pseudo-code for determining an adaptive texture-level threshold.

Figure 21 is a diagram showing two examples of gradient slope regions.

Figure 22A is a diagram showing an example frame with a gradient slope region, a textured region, a sharp-edge region and a flat region. Figure 22B is a diagram showing a contouring artifact in the gradient slope region of Figure 22A. Figure 22C shows 5 macroblock-level detail of a contouring artifact of Figure 22B.

Figure 23 is a flow chart showing a generalized region-based gradient slope detection technique.

Figure 24 is a block diagram of an example gradient slope detector according to one implementation.

10 Figure 25 is a diagram that depicts 4-to-1 down-sampling of a gradient slope region with film grains that potentially cause anomalous gradient slope directions.

Figure 26 is an equation diagram for 16x16 compass operators  $K_H$  and  $K_V$ .

Figure 27 is a code diagram showing example pseudo-code for computing the gradient direction for a region using the compass operators of Figure 26.

15 Figure 28 is a flow chart showing a technique for performing consistency checking for gradient slope regions.

Figure 29 is a diagram that depicts buckets in a bucket voting technique.

Figure 30 is a flow chart showing an example technique for selecting a macroblock QP to help preserve one or more non-zero AC coefficients.

20 Figure 31 is a diagram showing a DC shift in three neighboring blocks in a gradient slope region after quantization and inverse quantization.

Figure 32 is a flow chart showing a generalized technique for adjusting quantization to reduce or avoid introduction of contouring artifacts in DC shift areas.

25 Figure 33 is a flow chart showing a combined technique for tailoring quantization in DC shift areas to reduce or avoid introduction of quantization artifacts.

#### DETAILED DESCRIPTION

The present application relates to techniques and tools for efficient compression of video. In various described embodiments, a video encoder incorporates techniques for encoding video, and corresponding signaling techniques for use with a bitstream format or 30 syntax comprising different layers or levels. Some of the described techniques and tools can be applied to interlaced or progressive frames.

Various alternatives to the implementations described herein are possible. For example, techniques described with reference to flowchart diagrams can be altered by changing the ordering of stages shown in the flowcharts, by repeating or omitting certain

stages, etc. For example, initial stages an analysis (e.g., obtaining texture information for a picture or performing texture analysis in detecting smooth regions) can be completed before later stages (e.g., making encoding decisions for the picture or performing temporal analysis in detecting smooth regions) begin, or operations for the different stages can be 5 interleaved on a block-by-block, macroblock-by-macroblock, or other region-by-region basis. As another example, although some implementations are described with reference to specific macroblock formats, other formats also can be used.

The various techniques and tools can be used in combination or independently. Different embodiments implement one or more of the described techniques and tools. 10 Some techniques and tools described herein can be used in a video encoder, or in some other system not specifically limited to video encoding.

### **I. Computing Environment**

Figure 9 illustrates a generalized example of a suitable computing environment 900 in which several of the described embodiments may be implemented. The computing 15 environment 900 is not intended to suggest any limitation as to scope of use or functionality, as the techniques and tools may be implemented in diverse general-purpose or special-purpose computing environments.

With reference to Figure 9, the computing environment 900 includes at least one processing unit 910 and memory 920. In Figure 9, this most basic configuration 930 is 20 included within a dashed line. The processing unit 910 executes computer-executable instructions and may be a real or a virtual processor. In a multi-processing system, multiple processing units execute computer-executable instructions to increase processing power. The memory 920 may be volatile memory (e.g., registers, cache, RAM), non-volatile memory (e.g., ROM, EEPROM, flash memory, etc.), or some combination of the 25 two. The memory 920 stores software 980 implementing a video encoder with one or more of the described techniques and tools.

A computing environment may have additional features. For example, the computing environment 900 includes storage 940, one or more input devices 950, one or more output devices 960, and one or more communication connections 970. An 30 interconnection mechanism (not shown) such as a bus, controller, or network interconnects the components of the computing environment 900. Typically, operating system software (not shown) provides an operating environment for other software executing in the computing environment 900, and coordinates activities of the components of the computing environment 900.

The storage 940 may be removable or non-removable, and includes magnetic disks, magnetic tapes or cassettes, CD-ROMs, DVDs, or any other medium which can be used to store information and which can be accessed within the computing environment 900. The storage 940 stores instructions for the software 980 implementing the video encoder.

5 The input device(s) 950 may be a touch input device such as a keyboard, mouse, pen, or trackball, a voice input device, a scanning device, or another device that provides input to the computing environment 900. For audio or video encoding, the input device(s) 950 may be a sound card, video card, TV tuner card, or similar device that accepts audio 10 or video input in analog or digital form, or a CD-ROM or CD-RW that reads audio or video samples into the computing environment 900. The output device(s) 960 may be a display, printer, speaker, CD-writer, or another device that provides output from the computing environment 900.

15 The communication connection(s) 970 enable communication over a communication medium to another computing entity. The communication medium conveys information such as computer-executable instructions, audio or video input or output, or other data in a modulated data signal. A modulated data signal is a signal that has one or more of its characteristics set or changed in such a manner as to encode 20 information in the signal. By way of example, and not limitation, communication media include wired or wireless techniques implemented with an electrical, optical, RF, infrared, acoustic, or other carrier.

25 The techniques and tools can be described in the general context of computer-readable media. Computer-readable media are any available media that can be accessed within a computing environment. By way of example, and not limitation, with the computing environment 900, computer-readable media include memory 920, storage 940, communication media, and combinations of any of the above.

30 The techniques and tools can be described in the general context of computer-executable instructions, such as those included in program modules, being executed in a computing environment on a target real or virtual processor. Generally, program modules include routines, programs, libraries, objects, classes, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The functionality of the program modules may be combined or split between program modules as desired in various embodiments. Computer-executable instructions for program modules may be executed within a local or distributed computing environment.

For the sake of presentation, the detailed description uses terms like "decide" and "analyze" to describe computer operations in a computing environment. These terms are high-level abstractions for operations performed by a computer, and should not be confused with acts performed by a human being. The actual computer operations corresponding to these terms vary depending on implementation.

## **II. Generalized Video Encoder**

Figure 10 is a block diagram of a generalized video encoder 1000 in conjunction with which some described embodiments may be implemented. The encoder 1000 receives a sequence of video pictures including a current picture 1005 and produces compressed video information 1095 as output to storage, a buffer, or a communication connection. The format of an output bitstream can be a Windows Media Video or VC-1 format, MPEG-x format (e.g., MPEG-1, MPEG-2, or MPEG-4), H.26x format (e.g., H.261, H.262, H.263, or H.264), or other format.

The encoder 1000 processes video pictures. The term picture generally refers to source, coded or reconstructed image data. For progressive video, a picture is a progressive video frame. For interlaced video, a picture may refer to an interlaced video frame, the top field of the frame, or the bottom field of the frame, depending on the context. The encoder 1000 is block-based and uses a 4:2:0 macroblock format for frames. As shown in Figure 11, macroblock 1100 includes four 8x8 luminance (or luma) blocks (Y1 through Y4) and two 8x8 chrominance (or chroma) blocks (U and V) that are co-located with the four luma blocks but half resolution horizontally and vertically, following the conventional 4:2:0 macroblock format. For fields, the same or a different macroblock organization and format may be used. The 8x8 blocks may be further sub-divided at different stages, e.g., at the frequency transform and entropy encoding stages. The encoder 1000 can perform operations on sets of samples of different size or configuration than 8x8 blocks and 16x16 macroblocks. Alternatively, the encoder 1000 is object-based or uses a different macroblock or block format.

Returning to Figure 10, the encoder system 1000 compresses predicted pictures and intra-coded, key pictures. For the sake of presentation, Figure 10 shows a path for key pictures through the encoder system 1000 and a path for predicted pictures. Many of the components of the encoder system 1000 are used for compressing both key pictures and predicted pictures. The exact operations performed by those components can vary depending on the type of information being compressed.

A predicted picture (e.g., progressive P-frame or B-frame, interlaced P-field or B-field, or interlaced P-frame or B-frame) is represented in terms of prediction (or difference) from one or more other pictures (which are typically referred to as reference pictures or anchors). A prediction residual is the difference between what was predicted 5 and the original picture. In contrast, a key picture (e.g., progressive I-frame, interlaced I-field, or interlaced I-frame) is compressed without reference to other pictures.

If the current picture 1005 is a predicted picture, a motion estimator 1010 estimates motion of macroblocks or other sets of samples of the current picture 1005 with respect to one or more reference pictures, for example, the reconstructed previous picture 1025 10 buffered in the picture store 1020. If the current picture 1005 is a bi-predictive picture, a motion estimator 1010 estimates motion in the current picture 1005 with respect to up to four reconstructed reference pictures (for an interlaced B-field, for example). Typically, a motion estimator estimates motion in a B-picture with respect to one or more temporally previous reference pictures and one or more temporally future reference pictures, but B- 15 pictures need not be predicted from different temporal directions. The encoder system 1000 can use the separate stores 1020 and 1022 for multiple reference pictures.

The motion estimator 1010 can estimate motion by full-sample,  $\frac{1}{2}$ -sample,  $\frac{1}{4}$ -sample, or other increments, and can switch the precision of the motion estimation on a picture-by-picture basis or other basis. The motion estimator 1010 (and compensator 20 1030) also can switch between types of reference picture sample interpolation (e.g., between bicubic and bilinear) on a per-frame or other basis. The precision of the motion estimation can be the same or different horizontally and vertically. The motion estimator 1010 outputs as side information motion information 1015 such as differential motion vector information. The encoder 1000 encodes the motion information 1015 by, for 25 example, computing one or more predictors for motion vectors, computing differentials between the motion vectors and predictors, and entropy coding the differentials. To reconstruct a motion vector, a motion compensator 1030 combines a predictor with differential motion vector information.

The motion compensator 1030 applies the reconstructed motion vector to the 30 reconstructed picture(s) 1025 to form a motion-compensated current picture 1035. The prediction is rarely perfect, however, and the difference between the motion-compensated current picture 1035 and the original current picture 1005 is the prediction residual 1045. During later reconstruction of the picture, the prediction residual 1045 is added to the motion compensated current picture 1035 to obtain a reconstructed picture that is closer to

the original current picture 1005. In lossy compression, however, some information is still lost from the original current picture 1005. Alternatively, a motion estimator and motion compensator apply another type of motion estimation/compensation.

A frequency transformer 1060 converts the spatial domain video information into 5 frequency domain (i.e., spectral) data. For block-based video pictures, the frequency transformer 1060 applies a DCT, variant of DCT, or other block transform to blocks of the sample data or prediction residual data, producing blocks of frequency transform coefficients. Alternatively, the frequency transformer 1060 applies another conventional frequency transform such as a Fourier transform or uses wavelet or sub-band analysis. 10 The frequency transformer 1060 may apply an 8x8, 8x4, 4x8, 4x4 or other size frequency transform.

A quantizer 1070 then quantizes the blocks of spectral data coefficients. The quantizer applies uniform, scalar quantization to the spectral data with a step-size that varies on a picture-by-picture basis or other basis (e.g., a macroblock-by-macroblock 15 basis). Alternatively, the quantizer applies another type of quantization to the spectral data coefficients, for example, a non-uniform, vector, or non-adaptive quantization, or directly quantizes spatial domain data in an encoder system that does not use frequency transformations. Techniques and tools relating to quantization in some implementations are described in detail below.

20 In addition to adaptive quantization, the encoder 1000 can use frame dropping, adaptive filtering, or other techniques for rate control.

The encoder 1000 may use special signaling for a skipped macroblock, which is a macroblock that has no information of certain types (e.g., no differential motion vectors for the macroblock and no residual information).

25 When a reconstructed current picture is needed for subsequent motion estimation/compensation, an inverse quantizer 1076 performs inverse quantization on the quantized spectral data coefficients. An inverse frequency transformer 1066 then performs the inverse of the operations of the frequency transformer 1060, producing a reconstructed prediction residual (for a predicted picture) or a reconstructed key picture. 30 If the current picture 1005 was a key picture, the reconstructed key picture is taken as the reconstructed current picture (not shown). If the current picture 1005 was a predicted picture, the reconstructed prediction residual is added to the motion-compensated current picture 1035 to form the reconstructed current picture. One or both of the picture stores 1020, 1022 buffers the reconstructed current picture for use in motion compensated

prediction. In some embodiments, the encoder applies a de-blocking filter to the reconstructed frame to adaptively smooth discontinuities and other artifacts in the picture.

The entropy coder 1080 compresses the output of the quantizer 1070 as well as certain side information (e.g., motion information 1015, quantization step size (QP)).

5 Typical entropy coding techniques include arithmetic coding, differential coding, Huffman coding, run length coding, LZ coding, dictionary coding, and combinations of the above. The entropy coder 1080 typically uses different coding techniques for different kinds of information (e.g., DC coefficients, AC coefficients, different kinds of side information), and can choose from among multiple code tables within a particular coding technique.

10 The entropy coder 1080 provides compressed video information 1095 to the multiplexer ("MUX") 1090. The MUX 1090 may include a buffer, and a buffer level indicator may be fed back to a controller. Before or after the MUX 1090, the compressed video information 1095 can be channel coded for transmission over the network. The channel coding can apply error detection and correction data to the compressed video information 1095.

15 A controller (not shown) receives inputs from various modules such as the motion estimator 1010, frequency transformer 1060, quantizer 1070, inverse quantizer 1076, entropy coder 1080, and buffer 1090. The controller evaluates intermediate results during encoding, for example, estimating distortion and performing other rate-distortion analysis.

20 The controller works with modules such as the motion estimator 1010, frequency transformer 1060, quantizer 1070, and entropy coder 1080 to set and change coding parameters during encoding. When an encoder evaluates different coding parameter choices during encoding, the encoder may iteratively perform certain stages (e.g., quantization and inverse quantization) to evaluate different parameter settings. The

25 encoder may set parameters at one stage before proceeding to the next stage. Or, the encoder may jointly evaluate different coding parameters. The tree of coding parameter decisions to be evaluated, and the timing of corresponding encoding, depends on implementation:

30 The relationships shown between modules within the encoder 1000 indicate general flows of information in the encoder; other relationships are not shown for the sake of simplicity. In particular, Figure 10 usually does not show side information indicating the encoder settings, modes, tables, etc. used for a video sequence, picture, macroblock, block, etc. Such side information, once finalized, is sent in the output bitstream, typically after entropy encoding of the side information.

Particular embodiments of video encoders typically use a variation or supplemented version of the generalized encoder 1000. Depending on implementation and the type of compression desired, modules of the encoder can be added, omitted, split into multiple modules, combined with other modules, and/or replaced with like modules.

5 For example, the controller can be split into multiple controller modules associated with different modules of the encoder. In alternative embodiments, encoders with different modules and/or other configurations of modules perform one or more of the described techniques.

### **III. Characterization of Video Content Using a Perceptual Model**

10 Video content can be characterized using a perceptual model. This can help an encoder to make appropriate encoding decisions for different kinds of video content. An encoder can analyze a picture before encoding to provide characterizations for the content in different parts of the picture (e.g., blocks, macroblocks, etc.). For example, relatively smooth parts of a video picture, such as a blue sky, may be characterized as less acceptable for introducing distortion because certain kinds of quality degradation (e.g., quantization artifacts) are more easily perceived by humans in smooth regions. In contrast, distortion is generally less noticeable (and thus more acceptable) in texture regions.

15

With reference to Figure 12, a video encoder such as one described above with reference to Figure 10 implements adaptive encoding techniques in a process 1200 that characterizes portions (e.g., blocks of macroblocks, macroblocks, or other regions) of a video picture (e.g., as a smooth region, edge region, texture region, etc.) and adapts one or more encoding techniques according to the characterization. Many of the described techniques provide adaptive encoding within a picture, such as on a block, macroblock or other region. The techniques use information to classify different parts of the image and to encode them accordingly. More particularly, a video encoder characterizes portions of the picture to classify content based on its perceptual characteristics.

20

25

At 1210, the video encoder characterizes one or more portions of a video picture. For example, the encoder characterizes a block of the video picture based on intensity variance within the block. In one implementation, the encoder computes a sum of the differences between a pixel and its adjacent pixels for the pixels in the block or a down-sampled version of the block. This sum of differences value measures intensity variance between a pixel and its surrounding pixels. For example, surrounding pixels comprise two or more other pixels adjacent to or nearly adjacent to a pixel, such as above or below, to

the left or right, or diagonal to a pixel. The difference between a pixel's intensity and the intensities of its surrounding pixels is computed based on differences in luma and/or chroma data. In other words, the differences are computed with luma samples and/or chroma samples. An average computed difference value is assigned to the pixel (e.g., a 5 pixel gradient). A difference value is computed in this way for pixels in a block (e.g., a block gradient), or for some sub-sampled set thereof. The difference values assigned to pixels in a block are evaluated to determine a characterization or classification (e.g., smooth, edge, or texture; texture or non-texture; smooth or non-smooth; etc.) for the block, which can be expressed a *block value*. In one example, the pixel gradients for pixels in a 10 block are evaluated to determine a median difference value for the block gradient (e.g., a block median). Thus, intensity differences between pixels within a block provide a measure of intensity variance for a block, macroblock, or other video picture region.

A block median is not required to determine a block value. An intensity variance or block characterization may also be based on an average value for difference values 15 assigned to pixels in the block (e.g., a block average). The block median or average can be used to classify the block and/or can be used as input to a finer-grained control function. The characterization or control function adaptively varies one or more aspects of encoding.

Alternatively, instead of computing an intensity variance to characterize a block, 20 the encoder uses another metric.

At 1220, the encoder adaptively encodes the video picture based on the characterizations. In one implementation, encoding techniques for removal or reduction of 25 contouring artifacts are performed based on block characterization. For example, gradient slope detection, DC shift detection, AC coefficient preservation, and adaptive differential quantization are performed for certain smooth regions, and textured regions are quantized more strongly to conserve bit rate.

Although Figure 12 shows the characterizing stage 1210 preceding the adaptive encoding stage 1220 for multiple portions of a picture, these stages may also occur iteratively on a block-by-block basis in the picture or be ordered on some other basis.

30 At 1230, the encoder signals the adaptively encoded bit stream. When differential quantization is used by the encoder to encode based on block characterization, for example, the video encoder encodes information in the compressed bit stream using a signaling scheme for signaling the differential quantization to a video decoder.

At 1240, a corresponding video decoder reads the adaptively encoded bit stream, including the encoded data for the video picture. For example, the video decoder reads signaled differential quantization information. At 1250, the decoder decodes the compressed bit stream, for example, dequantizing blocks according to signaled differential 5 quantization information.

#### A. Example Block-based Characterization

Figure 13 is a diagram showing block-based operations for characterizing blocks using luma and/or chroma data. The luma block "Y" (1302) is an 8x8 block of a macroblock in a 4:2:0 macroblock format. Although not required, in this example, 10 corresponding chroma blocks 1304, 1306 for the pixel block are also used in computing a gradient block 1308. Although not required, as shown in this example, the luma block 1302 is down-sampled 1312 by a factor of 2 horizontally and vertically (e.g., by simple averaging of pairs of samples) to create a luma block 1310 that matches the 4x4 dimensions of the chroma blocks.

15 As shown in the down-sampled luma block 1310, the intensity value of a luma sample for a pixel 1314 is compared to samples for four pixels near it in the down-sampled luma block 1310, and an average sum of the difference between the sample for the pixel 1314 and the samples for its surrounding vertical and horizontal pixels is computed. In this example, the pixel 1314 is located at position  $Y'(r, c)$ . The average sum 20 of the differences for the luma intensity value for this pixel 1314 as compared to its surrounding pixels is:

$$L_l(r, c) = [ |Y'(r, c) - Y'(r, c-1)| + |Y'(r, c) - Y'(r-1, c)| + |Y'(r, c) - Y'(r, c+1)| + |Y'(r, c) - Y'(r+1, c)| ] / 4 \quad (5)$$

25 As shown,  $Y'(r, c)$  is the luma component of the pixel 1314 at row  $r$  and column  $c$  in the down-sampled block  $Y'$ .  $L_l(r, c)$  provides an indication of how the pixel 1314 differs in luma intensity from its neighbors within the block  $Y'$ . This luma intensity difference measurement is an example of a pixel gradient.

30 Optionally, chroma data 1304, 1306 may be considered alone instead of luma data, or may be considered together with luma data to determine intensity differences. The average sum of the differences for luma intensity values and chroma intensity values for pixel 1314 can be represented as the average of the differences in intensity values of samples for the surrounding pixels as shown in the following equation:

$$G_l(r, c) = \{ [ |Y'(r, c) - Y'(r, c-1)| + |Y'(r, c) - Y'(r-1, c)| + |Y'(r, c) - Y'(r, c+1)| + |Y'(r, c) - Y'(r+1, c)| ] + [ |U(r, c) - U(r, c-1)| + |U(r, c) - U(r-1, c)| + |U(r, c) - U(r+1, c)| ] \} / 8 \quad (6)$$

$$|U(r, c) - U(r, c+1)| + |U(r, c) - U(r+1, c)| + [ |V(r, c) - V(r, c-1)| + |V(r, c) - V(r-1, c)| + |V(r, c) - V(r+1, c)| ] \} / 12 \quad (6)$$

$G_I(r, c)$  is an example of a pixel gradient for the pixel located at  $(r, c)$  in the down-sampled block, and the pixel gradient provides an indication of how the pixel 1314 differs 5 in luma and chroma intensity from its surrounding pixel neighbors. In this example, the pixel gradient value  $G_I(r, c)$  is based on pixels that are immediately vertical or horizontal, but does not consider other pixels in the neighborhood. It is contemplated that other pixel data may also be considered in creation of a pixel gradient in other variations. For example, diagonal pixels could be considered as part of, or instead of the provided 10 arrangement. Or, intensity differences across a longer stretch (e.g., 2 or 3 pixels) could be considered.

$G_I(r, c)$  provides an indication of how a single pixel differs from its neighbors in luma and chroma intensity. In order to characterize the intensity variance for an entire block, the same analysis is performed on plural or all pixels within the block. In one such 15 example, a block 1308 of pixel gradients is created, and a block gradient is derived therefrom. As noted, computing a pixel gradient or a block gradient may include luma comparisons alone, chroma comparisons alone, or both luma and chroma comparisons together.

If desirable, the above equation for finding  $G_I(r, c)$  may be varied to account for 20 missing block boundary values. For example, samples outside the block may be extrapolated or assumed to be the same as other adjacent samples within the block when adapting the equation  $G_I(r, c)$  to account for boundary values. Or, the denominator of the equations may be reduced and surrounding samples in certain directions ignored in the comparisons, for example, where those surrounding samples are outside of the block. As 25 shown, a block 1308 of pixel gradients may provide pixel gradient data for all pixels in the block. Or, a block 1308 of pixel gradients may include pixel gradient data for less than all pixels in the block.

Figure 14 is a histogram of plural pixel gradients in the block 1308 of Figure 13. More specifically, the histogram 1400 provides a visualization of how the block is 30 characterized or valued. In this example, there are eight pixel gradient values below 30, and eight pixel gradient values above 30. Thus, a median value for this block gradient is 30. (For an even number of candidates, the median can be computed as the average of the two middle candidate values, or as one or the other of the two middle candidate values.) The median value may be used to characterize the block as smooth, texture, or edge. Of

course, other metrics may be used to characterize blocks once the pixel gradients or blocks of pixel gradients are obtained. For example, blocks may be characterized according to an average of pixel gradient values. Once a block value is assigned it can be used in a characterization scheme (e.g., smooth or non-smooth; smooth, texture, edge; etc.) or in a finer grained control function. The block value can be used to determine how the block is treated in an adaptive encoding strategy.

A block value may be selected by ordering plural pixel gradients and selecting a median gradient value from the ordered values. For example, a set of pixel gradients within a block, such as {10, 14, 28, 36, 38}, has a block value assigned equal to the median pixel gradient in the set, or 28. In another example, a block value is determined based on the average gradient in the set, or 25.2 for the preceding numerical example. Of course, the set may be obtained from a complete block gradient, or a subset thereof.

### C. Example Use of Characterization Information

Figure 15 is a graph of an example block characterization framework, continuing the example of Figures 13 and 14. As shown, a block with a block value in the range from 0 up to and including 30 will be characterized as a smooth block. A block with a block value in the range of greater than 30 but less than or equal to 60 will be characterized as a texture block, and a block with a block value greater than 60 will be characterized as an edge block.

Alternatively, an encoder uses another characterization framework, for example, one including other and/or additional characterizations for blocks or other portions of video pictures. For different gradients and metrics, the framework can change in scale and/or number of dimensions.

An encoder can use the characterizations of the blocks or other portions of video pictures when making encoding decisions. Table 2 relates features of an example adaptive coding scheme to block characterizations as described with reference to Figure 15. As shown, differently characterized blocks are treated differently in terms of one or more adaptive features.

<u>Characterization</u>	<u>DC Shift Detection</u>	<u>Gradient Slope Detection</u>	<u>Quantization</u>
Smooth	Yes	Yes	Lower QP
Edge	No	No	Higher QP
Texture	No	No	Higher QP

**Table 2: Adaptive Encoding Features**

The various adaptive features shown in Table 2 are discussed throughout this document and will be further discussed below. Alternatively, an encoder uses another mapping of adaptive feature decisions to block characterizations. Moreover, some features described herein need not take into account characterizations of video content.

5 **IV. Differential Quantization Based on Texture Level**

In differential quantization, an encoder varies quantization step sizes (also referred to herein as quantization parameters or QPs in some implementations) for different parts of a picture. Typically, this involves varying QPs on a macroblock or other sub-picture level. An encoder makes decisions on how to vary the QPs and can signal those decisions, 10 as appropriate, to a decoder.

Previous encoders have used bi-level differential quantization (varying between two QPs) and multi-level differential quantization (varying between three or more QPs). For example, in one bi-level differential quantization approach, one QP is used for macroblocks at picture edges and another QP is used for macroblocks in the rest of the 15 picture. This can be useful for saving bits at picture edges, where fine detail is less important for maintaining overall visual quality. In a multi-level differential quantization approach, a larger number of different QPs can be used for individual macroblocks in a picture. For example, an encoder can choose a QP for a macroblock and signal a differential between the QP for the current picture and the QP for the macroblock.

20 Perceptual sensitivity to quantization artifacts is highly related to the texture level of the video in both the spatial and temporal domain. High texture levels often result in masking effects that can hide quality degradation and quantization artifacts. However, in regions with lower texture levels (e.g., smooth regions), degradation and quantization artifacts are more visible. Although previous encoders have made quantization 25 adjustments for some parts of video pictures (e.g., picture edges), a more comprehensive content-based differential quantization strategy as described herein provides improved rate-distortion performance in many scenarios.

Accordingly, many of the described techniques and tools use texture-based differential quantization (referred to herein as *texture DQ*) to allocate bits based on various 30 texture levels to achieve better perceptual quality. In texture DQ, different QPs are chosen to code video based on texture information and, in some cases, based on other information such as temporal analysis information. An encoder analyzes texture information (and possibly other information) and applies texture DQ to appropriate regions (*texture DQ regions*), such as 8x8 blocks or macroblocks in a picture. Many of the described

techniques and tools focus on *smooth* regions as potential texture DQ regions. Smooth regions include flat regions (areas of constant or nearly constant color) and gradient slope regions (areas of color that vary at a constant or nearly constant rate across the region). Smooth regions may be considered smooth even when interrupted by small areas of noise, 5 film grains, or other color variations.

Figure 16 is a flow chart showing a generalized technique 1600 for applying differential quantization based on texture information. An encoder such as the encoder 1000 of Figure 10 or other tool performs the technique 1600.

At 1610, an encoder obtains texture information (e.g., characterizations or block 10 values that indicate whether different regions are smooth, edge, or texture regions) for a current picture. At 1620, the encoder finds a texture DQ region (e.g., a smooth region in which contouring artifacts may be present) or texture DQ regions in the current picture. At 1630, the encoder applies texture DQ to the texture DQ region(s) and encodes the picture. For example, smooth regions are coded with smaller QPs than high texture 15 regions. If there are more pictures to encode, the encoder takes the next picture at 1640 and selectively applies texture DQ to the next picture, as appropriate. The encoder outputs encoded data for the video picture, for example, to storage, a communication connection, or a buffer.

Different texture DQ region detection techniques can be used to determine whether 20 a region should be treated as a smooth region. For example, an encoder can use different texture metrics and/or different texture thresholds (and can adjust thresholds adaptively) to determine whether a particular region should be considered a texture DQ region. Adaptive quantization value mapping can be used to allocate bits for better perceptual video quality. Differential quantization decisions also can be based on temporal analysis (i.e., looking at 25 future pictures to make decisions based on characteristics of a region over time).

Differential quantization decisions can be made for both intra pictures and predicted pictures. For predicted pictures, P- and B-picture differential quantization intervals between differentially quantized pictures can be controlled. Further, by observing the texture of a picture when dominant high texture areas are present, the 30 smooth region texture threshold can be relaxed to code a relatively smooth region (compared to the dominant high texture areas) with a smaller QP.

Techniques similar to those described with reference to Figures 12-15 in Section III, above, can be used to generate a *texture map* for a current picture. For example, the encoder calculates gradients for the texture levels for the picture as the first derivatives

(differences) in the Y, U and V channels for the picture, as described in section III. When the macroblock format is 4:2:0, to speed up the calculation process, the encoder can downsample the Y channel by a factor of 2:1 horizontally and vertically. The encoder sums the gradients of Y, U and V for each pixel in both horizontal and vertical direction.

5 For an 8x8 block in full resolution, the encoder computes the mean of the sum of the gradients in the corresponding 4x4 block in the downsampled picture to use as the block gradient value. Computing the mean of the gradients has a lower computational complexity than computing the median as described in section III.

10 Alternatively, an encoder obtains texture information for the picture in some other way. For example, an encoder chooses different gradient directions for calculating gradients, calculates gradients only for the luma channel, etc. However the texture information is obtained or calculated, it can then be used to make texture DQ decisions.

15 The texture map indicates the texture levels of the different parts of the picture. For example, the texture map can be used to identify smooth regions (e.g., blocks, macroblocks, edges, or other areas) and textured regions in the picture. Described differential quantization techniques can be performed on appropriate parts of the picture based on the information in the texture map. Alternatively, an encoder use texture information without first creating a texture map.

#### A. Temporal Analysis

20 In addition to texture information from a current video picture, temporal analysis can be used to make accurate differential quantization decisions. One reason for using temporal analysis is that the impact of using a smaller QP on a smooth region will be greater if the smooth region remains smooth over several pictures, especially when the other pictures reference the smooth region in motion compensation. Conversely, one 25 benefit of using a smaller QP will be lost if smooth blocks are replaced with high texture or edge blocks in future pictures. Accordingly, an encoder looks at future pictures after finding a smooth region in a current picture and makes differential quantization decisions based on how smoothness of the region changes in the future pictures. The encoder can also look at previous pictures, for example, B-pictures that precede a current video picture 30 in display order but reference the current video picture in motion compensation.

Figure 17 shows an example technique 1700 for using temporal analysis to make texture DQ decisions. An encoder such as the encoder 1000 of Figure 10 or other tool performs the technique 1700.

At 1710, an encoder performs texture analysis on a current block in a current picture in a video sequence. For example, the encoder looks at gradient information for the block. The encoder can compare the gradient information to a gradient threshold for the block and classify the block as smooth or non-smooth (e.g., texture, edge), where the 5 gradient threshold is fixed or set dynamically for the current picture or other part of the video sequence. Alternatively, the encoder performs texture analysis for some other portion in the current picture.

At 1720, the encoder performs temporal analysis. The encoder can perform the temporal analysis automatically or only if the current block is classified as a smooth block. 10 For example, the encoder determines if a smooth block in a current picture stays smooth in future pictures. If so, the smooth region in the current picture is later coded with a smaller QP. Or, the encoder determines if a smooth block in the current picture was also smooth in previous pictures, or in both previous and future pictures.

15 The number of previous and/or future pictures that the encoder analyzes can vary depending on implementation. If the smooth region is replaced in a future picture (e.g., the next picture or some other temporally close picture) by a textured region, the smooth region in the current picture might be coded with a larger QP, since the advantages of using a smaller QP are likely not as persistent. In one implementation, temporally closer pictures are weighted more heavily than more distant pictures in making the differential 20 quantization decision. The weighting and the number of previous and/or future pictures that the encoder looks at can vary depending on implementation.

25 To simplify the calculations, the encoder can find a single value to compare the current block and the corresponding block in a future picture. For example, since luma values are fairly consistent within smooth blocks, the mean of the luma values for the block is calculated to measure the similarity of corresponding blocks in future pictures. In the following example equation, the "strength"  $S(t)$  of the future smoothness of corresponding blocks in a future picture is calculated by a sum of the weighted absolute difference between the mean luma values of the current block and the corresponding block in the future picture, the mean luma values of the corresponding blocks in the two future 30 pictures, and so on.

$$S(t) = C(n) * \sum_{i=1}^n (n - i + 1) * |(M(t+i) - M(t+i-1))| \quad (7)$$

where  $n$  is the total number of temporal "look-ahead" pictures,  $C(n)$  is normalization factor, which is defined to be  $2/(n*(n+1))$ , and  $M(t)$  is the mean of luma values for the

block (or corresponding block) in the picture at time  $t$ . The encoder can also measure past smoothness instead of or in addition to future smoothness. Alternatively, the encoder uses another weighting system and/or smoothness metric in the temporal analysis of smoothness.

5 Referring again to Figure 17, at 1730 the encoder uses results of the texture analysis and the temporal analysis to determine whether to classify the block as a texture DQ block. For example, the encoder computes a smoothness strength  $S(t)$  for a smooth block (but not other blocks) and compares the smoothness strength  $S(t)$  to a temporal smoothness threshold. The temporal smoothness threshold can be fixed or dynamically 10 set.

In Figure 17, if the encoder finds that the current block is a smooth block and that the corresponding block in previous and/or future pictures is also smooth, the encoder adds the current block to a count of texture DQ blocks at 1740. The encoder can use the count of texture DQ blocks to determine whether to perform texture DQ on the picture. 15 Alternatively, an encoder uses temporal analysis in some other way to make a texture DQ decision.

If there are more blocks to analyze, the encoder takes the next block at 1750 and repeats the process shown in Figure 17. This continues until the encoder has evaluated the blocks of the current video picture. At that point, the encoder uses the count of smooth 20 blocks or other results of the temporal analysis in an encoding decision.

Although Figure 17 shows an encoder performing temporal analysis on a block-by-block, alternatively, the encoder performs temporal analysis on a macroblock-by-macroblock basis or some other region-by-region basis.

#### B. Texture DQ Thresholds and Isolated Smooth Block Filtering

25 Whether or not the encoder uses temporal analysis, the encoder can use several other mechanisms in deciding when to apply texture DQ. An encoder can use one or more prevalence thresholds (e.g., percentages of smooth blocks in the picture) to make decisions on whether to perform DQ and, if so, how fine the QPs for texture DQ regions should be. For example, if the number or percentage of smooth blocks in a picture is above a 30 threshold, the encoder can choose a coarser step size in order to avoid spending too many bits encoding smooth content with small QPs. The encoder also may have a lower threshold to determine whether the number or percentage of smooth blocks is enough to use texture DQ in the picture at all.

Another way to reduce bit rate is to treat certain smooth blocks as texture blocks when the smooth blocks are in predominantly textured regions. This can be referred to as isolated smooth block filtering (although a smooth block need not be completely “isolated” to be filtered in this way). For example, a smooth block surrounded by textured 5 blocks need not be coded with a smaller QP than the textured blocks, since quantization artifacts in the smooth block are likely to be masked by the surrounding textured content. As a result, an encoder can choose not to perform texture DQ on isolated smooth blocks. The encoder also can disregard isolated smooth blocks when calculating the number or percentage of smooth blocks in a picture.

10 Figure 18 shows an example technique 1800 for making a texture DQ decision using thresholds and isolated smooth block filtering. An encoder such as the encoder 1000 of Figure 10 or other tool performs the technique 1800.

15 At 1810, the encoder finds smooth blocks in the current picture. For example, the encoder performs texture analysis and temporal analysis as described with reference to Figure 17. Alternatively, the encoder finds the smooth blocks in the current picture in some other way.

20 At 1820, the encoder performs isolated smooth block filtering. For example, the encoder removes single smooth blocks that are surrounded in the current picture by non-smooth blocks. An encoder can use many different decision models to perform isolated smooth block filtering. For example, an encoder can choose to treat a smooth block as a textured block only when all its neighboring blocks are textured blocks. Or, an encoder can choose to treat a smooth block as a textured block if a certain number of its neighboring blocks are textured. Or, the encoder removes isolated smooth blocks in larger groups (e.g., 2 or 3) and/or using some other test for whether block(s) are isolated.

25 At 1830, the encoder checks the percentage of smooth blocks in the picture against a low threshold (e.g., 1-2% of the total blocks in the picture). If the percentage of smooth blocks falls below the low threshold, the encoder determines that texture DQ will not be used for this picture (1840). If the percentage of smooth blocks is above the low threshold, the encoder checks the percentage against a high threshold at 1850. This higher 30 threshold is used to pick a QP for the smooth blocks. If the percentage is higher than the high threshold, the encoder performs texture DQ but chooses a coarser QP (1860) for the smooth blocks to reduce bit rate. Otherwise, the encoder chooses a finer QP (1870) for the smooth blocks. If there are more pictures to analyze (1880), the encoder can repeat the

process for other pictures. The number of thresholds and the threshold percentage values can vary depending on implementation.

Alternatively, an encoder performs isolated smooth block filtering without using texture DQ thresholds, or uses texture DQ thresholds without isolated smooth block

5 filtering. Or, an encoder performs texture DQ without isolated smooth block filtering or using DQ thresholds.

### C. Adaptive Texture Level Threshold

An encoder can use a fixed texture-level or smoothness threshold to determine whether a given block should be considered a texture DQ block (e.g., a smooth block).

10 Taking into account the bit rate cost of DQ signaling (e.g., one bit per macroblock in an “all macroblock” bi-level DQ signaling scenario) and the bit rate cost of quantizing some parts of a picture at a smaller QP, the threshold acts as a check on the costs of texture DQ. For example, an encoder obtains a block value (using a technique described with reference to Figures 13 and 14 or some other technique) for a block and compares the block value to 15 a fixed texture-level/smoothness threshold value (e.g., a threshold value described with reference to Figure 15).

An encoder also can adaptively change texture-level/smoothness threshold values. For example, since the perceptibility of smooth blocks may change in pictures with a lot of high-texture content, the texture-level threshold for classifying a block as a smooth block 20 can be relaxed in a medium-texture or high-texture picture. This is an example of an adaptive texture-level threshold. An encoder may allow several different thresholds to be selected within a range of thresholds. In one implementation, an adaptive texture-level threshold for smooth blocks can be varied between a block value of 14 and a block value of 30. Different differential quantization mappings can be used for different texture-level 25 thresholds. An adaptive texture level threshold can be useful for allocating bits to smoother regions in higher-texture frames to improve quality in the smoother regions.

Figure 19 shows a technique 19 for selectively adjusting texture level thresholds for high-texture pictures. An encoder such as the encoder 1000 of Figure 10 or other tool performs the technique 1900. The encoder determines whether to adjust texture level 30 thresholds by detecting the presence of dominant high-texture content in a picture. In one implementation, the detection of high-texture content is implemented by calculating the texture “energy” in a sliding window with size of 10 in a texture histogram.

Referring to Figure 19, an encoder obtains a texture information (e.g., a texture level histogram) for a picture at 1910 in an adaptive texture-level threshold technique

1900. For example, the encoder obtains a texture map as described above and creates a texture level histogram from the information.

At 1920, the encoder checks whether the picture is a high-texture picture. If the picture is a high-texture picture, the encoder adjusts the texture level threshold for the 5 picture at 1930. If the picture is not a high-texture picture, the encoder processes the picture without adjusting the texture level threshold (1940). The encoder then can analyze and choose texture level thresholds for other pictures (1950). Alternatively, the encoder applies a sliding scale of different texture level thresholds for different levels of high-texture content in the picture.

10 For example, to check the extent of dominant high-texture content in a picture, an encoder computes a texture histogram for the picture. The encoder applies a sliding window in the texture histogram to calculate texture energy and determine a peak or prominent high-texture band. Equation (8) shows one way for the encoder to calculate the texture energy in the window. The sliding window starts sliding from the minimum 15 texture level threshold  $g_0$  (which is by default 30), and the encoder computes the window value  $W(g)$  at  $g_0$ . The sliding window shifts 1 to the right after calculation of texture energy for that window, and the encoder computes the next window value  $W(g)$  starting at the new value of  $g_0$ . This continues until the encoder reaches the maximum of the texture levels represented in the histogram.

20 Let  $F(g)$  be the histogram of texture level per pixel. Let  $E(g)$  be the texture energy for the texture level, where  $E(g) = F(g) * g$ . The encoder calculates the texture energy of the sliding window  $W(g)$  as follows:

$$W(g) = \sum_{g=g_0}^{g_0+10} (F(g) * g) \quad (8).$$

25 If the maximum sliding window energy  $W(g)$  exceeds a certain percentage threshold of overall picture energy,  $g_0$  for that maximum sliding window energy  $W(g)$  is used to adjust the threshold for smooth regions.

Figure 20 shows pseudo-code 2000 used to determine a new adaptive smooth region threshold from  $g_0$ . If  $g_0$  is over 100, the adaptive threshold is set to 30. The encoder also checks if  $g_0$  is less than 30 and, if so, sets the adaptive threshold to 14. 30 Otherwise, if  $30 \leq g_0 < 100$ , the adaptive threshold is set to a value from the table  $g\_iFlatThTable$ . To help maintain video quality, the maximum difference of a new adaptive threshold from the last adaptive threshold is capped at  $\pm 4$  for all pictures except scene change key pictures. The adaptive smooth threshold should not exceed the

threshold used to identify textured blocks – for example, in Figure 20 the highest adaptive threshold value is 30.

Alternatively, an encoder adaptively adjusts texture level thresholds in some other way (e.g., with a different texture strength or energy metric, without a sliding window,

5 with a differently configured sliding window, with different threshold values in a table or other data structure, without capping differences between adaptive thresholds, capping differences in adaptive thresholds in some other way, etc.).

#### **D. I-Picture and P-Picture Differential Quantization**

Described differential quantization techniques and tools can be used separately or 10 in combination on intra pictures and predicted pictures. The term *I-picture differential quantization* (I-picture DQ) refers to application of differential quantization to I-pictures, and the term *P-picture differential quantization* (P-picture DQ) refers to application of differential quantization to P-pictures. The use of I-picture DQ results in higher quality I-pictures, and the quality improvement can be maintained longer for predicted pictures that 15 depend from those I-pictures. P-picture DQ can further improve P-picture quality in both intra and inter blocks, but the quality of those P-pictures will also depend on the quality of the pictures from which they are predicted. Similarly, the impact of P-picture DQ on the quality of later predicted pictures will depend the similarity of the later predicted pictures to the pictures from which they are predicted.

20 **E. Differential Quantization Intervals**

Both I-picture DQ and P-picture DQ use one or more of the techniques described herein to decide whether to apply different QPs for different texture-level blocks. To balance quality and bit usage, a P-picture DQ interval can be used to control the amount of bits that are spent on P-picture DQ. For example, an encoder chooses to use P-picture DQ 25 on one in every  $n$  P-pictures, where  $n \geq 1$ , but skips P-picture DQ for pictures in the interval between differentially quantized P-pictures. The encoder spends bits on differential quantization to improve the perceptual quality of some P-pictures, and those quality improvements carry over into other predicted pictures. At the same time, the DQ interval helps constrain the overall number of bits the encoder spends on differential 30 quantization of predicted pictures.

Alternatively, the encoder selects another interval. For example, the encoder may choose to use P-picture DQ on only one P-picture per I-picture, or choose some other interval. The interval may be fixed or adaptive. For example, the encoder may adaptively adjust the P-picture DQ interval based on the type of content being encoded.

## V. Gradient Slope Detection

Among various visual artifacts introduced in video compression, *contouring* is one particular artifact that can be caused by quantization. Contouring artifacts are perceived by human eyes as structured, gradient discontinuities in what are otherwise continuous, 5 very smooth regions such as sky, water, etc. Such discontinuities can be very distracting and may lead a human observer to conclude that a whole picture is badly distorted even if other parts of the picture are coded with little visual distortion.

Gradient slope regions can give rise to contouring artifacts. According to one definition, a region is considered to be a *gradient slope region* if the region is smooth or 10 relatively smooth but pixel values change gradually within the region. Thus, while both gradient slope regions and flat regions are considered to be smooth regions, gradient slope regions differ from flat regions. According to one definition, a flat region is characterized by constant or relatively constant pixel values throughout the flat region. Gradient slope regions typically lack strong edges and extensive texture detail.

15 Figure 21 shows two examples of gradient slope regions. The gradient slope direction in each region is represented by arrows. In gradient slope region 2100, luma values increase gradually from the top to the bottom of the region. The direction of the slope in gradient slope region 2100 is the same in each part of the region. In gradient slope region 2110, luma values increase gradually from the center to the edges of the 20 region. The direction of the gradient slope varies within the gradient slope region 2110. However, within small neighborhoods, the gradient slope direction at each point is within a small angle  $\theta$  of the gradient slope direction at other points in the neighborhood, except for the neighborhood that includes the center point. As shown in Figure 21, gradient slope regions include regions where the gradient slope direction is constant throughout the 25 region, and regions where the gradient slope direction has small variations within a neighborhood.

Figure 22A is a diagram showing an example picture 2200 with a gradient slope region 2210, a textured region 2220, a sharp-edge region 2230 and a flat region 2240.

Figure 22B is a diagram showing results of quantization in the gradient slope region 2210. 30 The banding effect that is now visible (e.g., within macroblock 2250) is a contour artifact. Figure 22C shows detail of the macroblock 2250. Quantization of transform coefficients for the top half of the luma samples in macroblock 2250 results in uniform values stemming from a DC value of 68. Quantization of transform coefficients for the bottom half of the luma samples in macroblock 2250 results in uniform values stemming from the

DC value of 70. Thus, the quantization of the transform coefficients for the luma samples has created a visible contour artifact between the top-half 8x8 blocks and the bottom-half 8x8 blocks in macroblock 2250.

Many existing video encoders use techniques that are applied to a whole video 5 picture in an attempt to reduce contouring artifacts in the picture. Such techniques may result in over-spending bits, especially in regions that contain little or no contouring artifacts. Accordingly, several described techniques and tools allow an encoder to detect gradient slope regions, where contouring artifacts are likely to happen. When gradient slope regions are detected, an encoder can make coding decisions that reduce or avoid 10 introduction of contouring artifacts (e.g., adjustments of QPs) in the gradient slope regions. By doing so, an encoder can allocate bits more effectively and achieve better visual quality.

To detect gradient slope regions, an encoder can implement one or more of the following techniques:

- 15 1. Gradient slope region detection with coding decisions focused on reducing or removing introduction of contouring artifacts in the detected region(s).
2. Region-based gradient estimation and down-sampling to reduce computational cost and/or allow accurate gradient slope region detection despite the presence of anomalies such as film grains.
- 20 3. A gradient consistency check to detect gradual gradient change in local neighborhoods.
4. Bucket voting to make a binary decision regarding the presence of gradient slope region(s) in a picture.
- 25 5. The generation of a gradient slope mask (e.g., at macroblock-level) and gradient direction map to help an encoder to make appropriate coding decisions.

Figure 23 shows a generalized region-based gradient slope detection technique 2300. An encoder such as the encoder 1000 of Figure 10 or other tool performs the technique 2300. In some cases, the region-based gradient slope detection technique 2300 30 allows faster detection of gradient slope content by eliminating the need to find gradient slope directions for each pixel in a picture. For example, the picture is divided into non-overlapping rectangular regions of the same size. The size of the regions can vary depending on implementation. In one implementation, a region is a 16x16 macroblock (four 8x8 blocks). Preferably, the region is of a size that allows macroblock alignment.

At 2310, an encoder checks whether a current region is a smooth region. For example, the encoder uses a texture map of the picture in which an 8x8 block is characterized as smooth if its assigned block gradient value is less than 30, or the encoder uses checks whether the current region is smooth using another technique described in section III or IV. When a region includes multiple blocks, the region is considered to be a smooth region if all blocks contained in the region are smooth (or, alternatively, if some minimum number of the blocks are smooth). Different implementations can use different criteria for determining whether a particular region or block is smooth. For example, the criteria for determining whether a region is smooth may be different if the picture is down-sampled.

If a region is not smooth, the next region is processed (2320). For a smooth region, the encoder finds a gradient direction at 2330. For example, the encoder finds a gradient direction using a technique such as the one described with reference to Figures 26 and 27. Alternatively, the encoder finds the gradient direction with some other technique.

At 2340, the encoder makes a gradient slope decision for the region, using thresholds and/or decision-making logic that depend on the technique and metrics used to find the gradient direction for the region. If there are more regions to be processed, the encoder processes the next region (2320). In one implementation, after computing initial gradient directions for different regions in a picture, the encoder generates a binary mask that indicates whether gradient slope is present in different regions by applying a sliding window in the picture. The information in the binary mask allows the encoder to make accurate gradient slope decisions.

Figure 24 is a block diagram of an example gradient slope region detector (GSR detector) 2400 in a video encoder such as the one shown in Figure 10. The GSR detector 2400 takes pixel data from a current picture 2405 as input.

Depending on picture size and potentially other factors, the GSR detector 2400 determines whether to perform down-sampling in down-sampling module 2410. Example down-sampling techniques are described below.

The gradient calculator 2420 takes (possibly down-sampled) pixel data and a texture map 2425 as input and calculates gradients for smooth regions. For example, the gradient calculator uses a technique such as the one described with reference to Figures 26 and 27 or uses some other technique. An example region size in the gradient calculation is 16x16, but the size of regions can vary depending on implementation. Depending on whether and how much down-sampling is applied, the region for which a gradient is

calculated can represent different amounts of area in the original picture 2405. The gradient calculator 2420 outputs a map or other data structure indicating the gradient directions for smooth regions.

5 The consistency checker 2430 takes the calculated gradients for smooth regions and checks the angular consistency of those gradients, for example, as described below. The consistency checker 24 produces a consistency map or other data structure indicating consistency information for the calculated gradients.

10 The decision module 2440 uses additional decision rules (after consistency checking) to determine whether smooth regions should be considered gradient slope regions. Example decision rules and criteria are described below. The decision module 2440 considers the consistency map or other data structure indicating consistency information, and can also consider the calculated gradient directions or other information. The decision module 2440 outputs decision information in a map or other data structure for regions of the same or different size than the region size used in the gradient 15 calculation.

15 The decision for each region is provided to mask generator 2450 which produces a gradient slope mask and/or a binary gradient slope decision mask 2495 that indicates gradient slope decisions for regions in the picture. For example, a mask 2495 comprises a bit equal to "1" for each gradient slope region and a bit equal to "0" for other regions. 20 Accepting calculated gradients as input, the mask generator 2450 can produce another mask 2495 that indicates final gradient slopes for different regions of the original picture, accounting for down-sampling and mask decisions. When the GSR detector 2400 performs down-sampling before gradient calculation, the mask generator 2450 can assign gradient slopes for down-sampled regions to corresponding regions of the original picture.

25 The components of the GSR detector 2400 are shown as separate modules in Figure 24, but the functions of these components can be rearranged, combined or split into different modules depending on implementation. Furthermore, components of gradient slope detector 2400 can be omitted in other implementations. For example, down-sampling is not required. A GSR detector need not take a texture map as input, and can instead get 30 an indication of whether a region is smooth or not from some other source. A GSR detector need not use a consistency checker. Although a GSR detector will make some kind of decision as to whether a region is a gradient slope region, the specifics of how decisions are made (including decision rules in the decision module) can vary depending

implementation. Gradient slope decisions need not be included in a binary mask and may be communicated to other parts of the encoder in some other way.

#### A. Region-based Gradient Direction Estimation with Down-sampling

Down-sampling can be used prior to finding gradient directions for regions in order to reduce computational cost. In one implementation, if the original picture width is greater than 1280 and the height is greater than 720, the original picture is 4-to-1 down-sampled. For example, in a 1080p arrangement with a picture width of 1920 pixels and a picture height of 1080 pixels, a decoder produces a down-sampled picture with a width of 480 pixels and a height of 270 pixels.

Typically, a down-sampled picture is divided into non-overlapping rectangular regions of the same size. For example, after downsampling, each 16x16 region corresponds to 4 macroblocks (16 blocks) of the original, full resolution picture. A region in the down-sampled picture is considered to be a smooth region if at least 12 blocks to which the region corresponds are smooth. Region sizes depend on implementation, and the relation between regions in gradient estimation and regions in original pictures varies depending on down-sampling ratio.

Down-sampling also is useful for improving accuracy of gradient slope region detection despite the presence of anomalies such as film grains. For example, consider a portion of a picture 2500 with DC values of blocks as shown in Figure 25. The majority of the picture portion 2500 has consistent gradient slope directions, as shown by the gradually increasing DC values from the top to the bottom of the picture portion. However, the white sample values represent DC values affected by film grains that create anomalous gradient slope directions at full resolution. With simple 2-to-1 down-sampling horizontally and vertically, the dark-bordered sample values are used to calculate the gradient slope direction. Because the down-sampled values maintain a consistent gradient slope, the film grains do not affect detection of the gradient slope.

Down-sampling can be used for other picture resolutions, and other down-sampling ratios also can be used.

#### B. Calculating Gradient Slope Direction

In one implementation, to calculate gradient slope direction for a smooth region, two 16x16 compass operators  $K_H$  and  $K_V$  (defined in Figure 26) are applied to the region. This produces two gradients  $g_X, g_Y$  for the region, one for the horizontal direction and one for the vertical direction. For a 16x16 region, the compass operators give positive weights

to some values of the region and negative weight to other values of the region.

Alternatively, the compass operators compute gradients in some other way.

An angular representation of the gradient direction, denoted as  $\theta$ , is derived from the two gradients and mapped to an integer in [0, 255]. The pseudo-code 2700 in Figure 5 27 shows an example routine for computing the gradient direction for a region ( $\otimes$  denotes a per-element product) using the compass operators of Figure 26. If the region is a textured region or edge region, the routine returns -2. If the region is smooth but flat (indicated by low absolute values for the gradients  $g_x$  and  $g_y$  for the region, the routine returns -1. Otherwise, the routine computes the gradient slope as the arctangent of the 10 vertical gradient  $g_y$  over the horizontal gradient  $g_x$ , using offsets to differentiate between slope directions for same arctangent values (e.g., whether a positive arctangent value indicates an above, right slope or a below, left slope) and represent the range of slope values as positive numbers.

Alternatively, the gradient direction is computed in some other way. For example, 15 the encoder uses different compass operators, different thresholds for slope regions, different logic to compute the slope, and/or a different representation for slope information.

### C. Neighborhood Gradient Consistency Check

An encoder can perform a gradient consistency check for regions in order to help 20 make an accurate decision about whether a region should be considered a gradient slope region. The gradient consistency check helps to avoid “false alarms” in gradient slope content detection. In one implementation, the gradient slope consistency check involves using a 3x3 sliding window (three regions by three regions) to determine gradient slope consistency.

25 Figure 28 shows a technique for performing consistency checking for gradient slope regions. An encoder such as the encoder 1000 of Figure 10 or other tool performs the technique 2800.

At 2810, the encoder positions a sliding window at a current region in the picture. At 2820, the encoder checks the gradient directions of regions in the sliding window. 30 Then, at 2830, the encoder makes a consistency decision for the current region. For example, given the gradient directions of detected smooth regions in a picture (potentially down-sampled), a gradient consistency check is performed with the sliding window containing 3x3 neighboring regions. The window is moved in raster scan order, positioning the window on a region in the picture (e.g., by centering the window on the

region, performing the consistency check, then moving the window from left to right across the picture). For a given window value, the consistency check requires the difference between the maximum and the minimum *gradientDirection* (see, e.g., Figure 27) of all 9 regions within the window to be less than 32 (equivalent to 45 degrees when slopes are represented by numbers from 0 to 255). If this condition is satisfied, the moving window value for the 3x3 set of regions is 1; otherwise it is 0. Alternatively, the encoder uses a different mechanism to check consistency of slope directions, for example, using a different size sliding window, different slope range threshold for maximum slope – minimum slope, different measure such as variance for slope consistency, and/or different 5 checking pattern, or computes a sliding window value for each region as opposed to sets of regions. The consistency check varies for different representations of slope 10 information.

The encoder can then process the next set of regions (2840). As output, the encoder produces a mask or other data structure indicating decision information. For 15 example, the encoder produces a binary consistency mask (referred to herein as *consistencyMask*) obtained by positioning the sliding window and performing the consistency check on sets of regions in the picture, and assigning each set of regions a decision of 1 (consistent slope) or 0.

Optionally, the encoder performs further processing on the decision information. 20 In some implementations, an encoder performs morphological operations on a consistency mask to help refine gradient consistency decisions for a picture. Two possible morphological operations are *Erode* and *Dilate*.

For example, an *Erode* operation is performed on every bit in the *consistencyMask*, followed by a *Dilate* operation. In the *Erode* operation, a bit initially marked as 1 is 25 marked as 0 if in the four closest pixels (here, values in the *consistencyMask*), more than one was initially marked as 0. In the *Dilate* operation, a bit initially marked as 0 is marked 1 if in the four closest pixels, more than one were initially marked as 1.

Alternatively, an encoder generates masks without using morphological operations or other post-processing of the decision information.

#### 30 D. Decision Rules and Bucket Voting

Even after performing consistency checking, the incidence of smooth regions may be so low, or the smooth regions may be so isolated, that they would be inefficient to encode specially. For example, even after applying morphological operations, there may still be gradient slope regions represented in *consistencyMask* that are isolated enough to

not need differential quantization. In some implementations, an encoder uses decision rules (including, for example, bucket voting) to help decide whether DQ should be applied to gradient slope regions in the picture. In the GSR detector 2400 of Figure 24, decision module 2440 makes such decisions.

5 In one implementation, the encoder makes one or more binary decisions regarding whether the current picture contains significant gradient slope based on *consistencyMask*. The mask *consistencyMask* is divided into 25 rectangular regions of the same size (called *buckets*) with 5 buckets in each row and 5 in each column. (The “bucket” regions are hence larger than the regions used for decisions and regions used for gradient 10 calculations.) The 1s within each bucket are counted. Let  $Buckets[i][j]$  be the number of 1s contained in the bucket at location  $(i, j)$ , where  $0 \leq i, j \leq 4$ . Horizontal and vertical bucket projections – the number of 1s in each column of buckets and the number of 1s in each row of buckets, respectively – also are calculated according to the following 15 relationship:

$$15 \quad \begin{aligned} BucketProjection\_H[i] &= \sum_{0 \leq j \leq 4} Buckets[i][j] \\ BucketProjection\_V[j] &= \sum_{0 \leq i \leq 4} Buckets[i][j] \end{aligned} \quad (9)$$

In this implementation, the picture is considered to contain significant gradient slope if any of the following conditions are satisfied:

20

1. At least 6% of the pixels in *consistencyMask* (regardless of bucket distribution) are marked as 1, OR
2. In one or more of the buckets, at least 75% of the pixels are marked as 1, OR
3. In one or more of the bucket projections, at least 20% of the pixels are marked as 1.

25 For example, 16 x 16 regions for a down-sampled picture of size 960x1440 are represented with a mask of size 20x30 (each value for a 3x3 set of regions of the down-sampled picture), which is in turn divided into 25 buckets, each bucket corresponding to a 24 regions of the consistency mask. Each bucket includes 24 bits from *consistencyMask*, for a total of  $25 \times 24 = 600$  bits. The encoder counts the number of 1s in each bucket, with a distribution as shown in Figure 29. The encoder checks whether the total number of 1s is more than 6% of all bits. In this case, the total number of 1s (as shown in Figure 29) is 30 83, which is more than 6% of all bits. Thus, the encoder in the case would skip bucket projection, due to satisfaction of condition 1, above. If the total number of 1s were below the threshold for condition 1, the encoder would whether 75% of the bits in any bucket

were 1s (condition 2), and, if necessary, check horizontal and vertical bucket projections (condition 3) to determine whether the regions indicated as being gradient slope regions are such that a gradient slope mask and decision mask should be generated, such as the macroblock-level gradient slope masks described below.

5        Alternatively, an encoder uses other decision rules for processing consistency information in a mask *consistencyMask* or other representation. For example, the percentage thresholds shown in conditions 1, 2 and 3 can vary depending on implementation. Or, one or more of the conditions is omitted, or the conditions are reordered, replaced or supplemented by other conditions (e.g., different directions for  
10      bucket projections, etc.). Aside from checking consistency information, the encoder can also consider gradient values and/or other information when deciding whether or how much DQ should be applied to gradient slope regions in the picture. As another alternative, an encoder can omit these decision rules altogether, and simply use the *consistencyMask* when generating a gradient slope mask.

15      **E.      Macroblock-level Gradient Slope Mask Generation**

To provide gradient slope information in a form useful for later encoder decision-making, the encoder puts the information in maps, masks, or other data structures. The information can include gradient slope region presence/absence information as well as actual gradient direction values for gradient slope regions.

20      For gradient slope presence/absence information, if gradient slope regions are detected, the encoder produces a gradient slope mask. For example, an encoder produces a macroblock-level gradient slope mask (referred to herein as *MBSlopeMask*) by converting a region-level mask (such as *consistencyMask*) back to macroblock-level for the original picture, considering possible down-sampling. Note that each value in  
25      *consistencyMask* corresponds to 9 macroblocks in the original picture, or 36 macroblocks if the picture is 4-to-1 down-sampled. For each bit with value 1 in *consistencyMask*, the encoder marks corresponding macroblocks as 1 in *MBSlopeMask* except for macroblock that are not smooth. Checking for smoothness again helps to avoid false alarms in gradient slope detection. For example, in one implementation an encoder uses a texture  
30      map to obtain texture information for blocks in a macroblock, and the macroblock is considered smooth only if all four blocks within the macroblock are smooth.

Alternatively, the encoder provides gradient decision information in some other form and/or uses some other decision for macroblock smoothness.

5 For gradient direction information, a gradient direction map is generated by assigning each region's gradient direction to all its corresponding macroblocks that are smooth. In doing so, the encoder accounts for possible size differences between macroblocks of the original picture and gradient regions due to down-sampling before gradient calculation.

10 The generated gradient slope mask and gradient direction map are then used in the encoder to make better coding decisions. Generally speaking, the results generated by a gradient slope region detector can be used by an encoder to make other coding decisions. For example, an encoder can make quantization decisions based on a generated gradient slope mask and/or gradient direction map. Some of the possible encoder decisions are described below.

#### **VI. Adjusting Quantization to Preserve Non-zero AC Coefficients**

15 Typically, a picture is assigned a picture-level quantization parameter by a rate control unit in an encoder. Using the same picture-level QP, the amount of bits used to represent a highly textured macroblock is typically much greater (as much as 10 to 50 times greater) than the amount of bits used to represent a low textured macroblock. Since the human visual system is less sensitive to distortion in a busy, highly textured area than in a smooth, low-textured area, however, it makes sense to use a smaller QP for low textured macroblocks and a larger QP for highly textured macroblocks.

20 This leads to the often-used strategy of classifying macroblocks according to human visual importance (usually using variance of the blocks or the strength of the gradients inside the blocks) and assigning a target number of bits proportional to some perceptual weighting. The quantization parameter for each macroblock to be modified is selected by modifying the picture level quantizer according to the weighting.

25 Experiments have shown that in smooth regions of very low variation, blocks are often quantized to have energy only in DC coefficients (with no non-zero AC coefficients remaining) even at a reasonably low QP. Surprisingly, when DC values in adjacent blocks in extremely smooth regions vary by only 1 from block-to-block, the perceived blocky, contouring artifact are a lot more severe than one would expect with such a small difference in absolute terms. The occurrence of this type of artifact in relatively small regions inside an otherwise well-coded picture can cause the overall perceived quality for the entire picture to be lowered.

30 Traditional rate-distortion-based and perceptual-based macroblock QP selection techniques do not handle this situation well. With rate-distortion optimization, the smooth

blocks would be considered well-coded because of the small distortion in absolute terms, and thus no further bits would be spent for these blocks. On the other hand, typical perceptual-based methods classify macroblocks into perceptual classes and assign a quantization parameter to each macroblock by adding or subtracting a pre-defined offset to 5 the picture-level quantization parameter according to the perceptual class of the macroblock. Unless the pre-defined offset is very aggressive (e.g., reducing QP for smooth regions to 1), such methods cannot guarantee that smooth blocks with small variations will not be quantized to a single non-zero DC coefficient, with all AC coefficients quantized to zero. But setting a very aggressive offset can increase bits spent 10 in macroblocks that may not need them to improve perceptual quality, raising bit rate inefficiently and conflicting with the picture-level quantization parameter selected by the encoder for rate control.

15 Accordingly, several techniques and tools described below selectively and judiciously allocate bits within pictures such that enough bits are allocated to smooth regions to reduce or remove introduction of blocking or contour artifacts.

For example, an encoder calculates QPs and selects a quantization parameter for 20 each macroblock within an I-picture to allocate enough bits to smooth blocks, thereby reducing perceived blocking artifacts in the I-picture. For each macroblock with one or more smooth blocks, a QP is selected such that there are at least  $N$  non-zero quantized AC coefficients per block of the macroblock, where  $N$  is an integer greater than or equal to 1. Often, the preserved AC coefficients are coefficients for the lowest frequency AC basis 25 functions of the transform, which characterize gradual value changes horizontally and/or vertically across a block. This tends to help perceived visual quality for each block, especially for smooth regions with low variation. In one implementation, an encoder selects the largest QP, not exceeding the picture QP, that still preserves AC coefficients as desired. There may be situations (e.g., very flat blocks) that non-zero AC coefficients are not preserved. In general, however, in this way, the encoder is not overly aggressive in spending bits with smaller QPs and reduces or avoids conflict with the picture QP.

With reasonable values of  $N$ , the selected QP does not change for most 30 macroblocks; it is the same as the picture QP for most macroblocks, and only a few smooth blocks are affected. Reasonable values of  $N$  are 1, 2, 3 or 4. The selected QP is more likely to change for macroblocks with low texture. In one implementation,  $N=1$  or 2 improves perceived quality without too much increase in the picture's bit rate.

Figure 30 shows an example technique 3000 for selecting a macroblock QP to help preserve one or more non-zero AC coefficients. An encoder such as the encoder 1000 of Figure 10 or other tool performs the technique 3000.

At 3010, the encoder finds the  $N^{\text{th}}$  largest AC coefficients of each luma block of the macroblock. For example, the encoder finds the second largest AC coefficient of each of the four 8x8 blocks of a 16x16 macroblock, if  $N = 2$ . Let  $AC(0), AC(1), AC(2), AC(3)$  be the  $N^{\text{th}}$  largest coefficients for the four luma blocks 0, 1, 2 and 3, respectively. For different block organizations in a macroblock, the  $N^{\text{th}}$  coefficients can come from more or fewer blocks in the macroblock.

At 3020, the encoder finds the minimum of these  $N^{\text{th}}$  coefficient values. For the  $N^{\text{th}}$  coefficients of four blocks,  $AC_{\min} = \min(AC(0), AC(1), AC(2), AC(3))$ . For other numbers of blocks,  $AC_{\min}$  is computed differently.

At 3030, the encoder sets a QP for the macroblock such that  $AC_{\min}$  is outside the dead zone threshold for that QP. The dead zone threshold is a “cut-off” threshold for quantizing an AC coefficient to zero when the value of QP is used for quantization. The dead zone threshold is usually predetermined for, and proportional to, a given QP. The dead zone threshold is selected at some point between 0 and the first reconstruction point. When the encoder uses either a uniform quantizer or non-uniform quantizer, the first reconstruction point depends on the QP value and whether uniform or non-uniform quantization is used. In one implementation, the first reconstruction point is the reconstructed value of quantized coefficient level = 1, which for uniform quantization is  $2 * QP$  and for non-uniform quantization is  $3 * QP$ . For uniform quantization, the cut-off threshold thus lies between 0 and  $2 * QP$ . For non-uniform quantization, the cut-off threshold thus lies between 0 and  $3 * QP$ . For example, the dead zone threshold  $Z(QP)$  is selected as  $Z(QP) = 6 * QP / 5$  for uniform quantization, and  $Z(QP) = 2 * QP$  for non-uniform quantization. Alternatively, other cut-off thresholds can be used.

An AC coefficient  $AC$  will be quantized to zero if:  $\text{Abs}(AC) < Z(QP)$ . To set (3030) the QP for a macroblock, an encoder can find the QP for the macroblock ( $QP_m$ ) that will preserve at least  $N$  AC coefficients by comparing  $AC_{\min}$  with  $Z(QP)$  for candidate values of QP, starting with the picture QP and decreasing QP until a minimum QP for the quantizer is reached (e.g.,  $QP = 1$ ) or the inequality  $\text{Abs}(AC_{\min}) \geq Z(QP)$  is satisfied. If the inequality  $\text{Abs}(AC_{\min}) \geq Z(QP)$  is satisfied, the encoder sets the threshold QP for the macroblock to be the first QP (i.e., highest qualifying QP) that satisfies the inequality.

Alternatively, the encoder uses other logic to compute the QP for the macroblock, for example, starting from the lowest QP or using a binary search of QP values.

The process of using  $QP_m$  to quantize all blocks in the macroblock can be referred to as *unconstrained* bit rate quantization. In a *constrained* bit rate quantization technique, 5 an encoder determines the maximum QP (not greater than the picture QP) needed to produce the desired number of non-zero AC coefficients for each of the luma blocks of the macroblock separately (e.g.,  $QP_0$ ,  $QP_1$ ,  $QP_2$ , and  $QP_3$  for blocks 0, 1, 2 and 3, respectively) as described above. It follows that  $QP_m$  equals the minimum of  $QP_0$ ,  $QP_1$ ,  $QP_2$ , and  $QP_3$ . To reduce bit usage, an encoder could use  $QP_i$  to quantize block  $i$  (where  $i = 0, 1, 2, 3$ , 10 etc.) in place of  $QP_m$ . In an encoder that specifies a single QP for an entire macroblock, the encoder can instead keep only those AC coefficients that are non-zero when quantized using  $QP_i$  for each block  $i$  when quantizing the block using  $QP_m$ , preserving only the top  $N$  non-zero AC coefficients in a given block even if other AC coefficients in the block would be preserved with quantization by  $QP_m$ . For the quantization process shown in Figure 30, 15 the quantization process for each luma block can be performed as a two-pass process. In the first pass, the encoder “thresholds” DCT coefficients to zero if the coefficient is less than  $Z(QP_i)$ , and otherwise keeps the same DCT coefficients. Then, the “thresholded” DCT coefficients are quantized in the same manner using  $QP_m$ .

Alternatively, an encoder preserves non-zero AC coefficients in some other way. 20 For example, an encoder can select a QP on a basis other than a macroblock-by-macroblock basis (e.g., block-by-block basis). The encoder can preserve AC coefficient for I-pictures, P-pictures, or B-pictures, or combinations thereof.

If at the minimum possible QP the number of non-zero quantized coefficients is less than  $N$ ,  $N$  can be adjusted accordingly.

## 25 **VII. Differential Quantization on DC Shift**

In a typical lossy encoding scenario, not all quantized DC and AC coefficients can be recovered exactly after inverse quantization. For example, in some video codecs, DC coefficient values shift by one (i.e., increase or decrease by one relative to their pre-quantization value) for some QPs and DC coefficient values. This phenomenon is an 30 example of *DC shift*. Representations of some DC coefficient values are lossless through quantization and inverse quantization at one or more lower QPs, but lossy in other, higher QPs.

A region with several blocks in which all the AC coefficients are quantized to 0 and the DC coefficients cannot be recovered exactly can exhibit visible contouring

artifacts in DC shift areas. Such regions with contouring artifacts are often smooth, gradient slope regions, such as sky, water or light rays. Figure 31 is a diagram showing a DC shift in three neighboring blocks in a gradient slope region after quantization and inverse quantization. The DC values of three neighboring blocks 3102, 3104, 3106 in a 5 gradient slope region are 68, 69, and 70, respectively, prior to quantization. After quantization and inverse quantization, the DC value of block 3104 is shifted to 70. As shown in Figure 31, the DC values of the three neighboring blocks are now 68, 70, and 70. When such blocks are in a gradient slope region, the quantized DC values may cause 10 perceptible contouring artifacts. For example, referring again to Figures 22A-C, the gradient slope region 2210 has been quantized, resulting in a visible contouring artifact in Figure 22B. As shown in Figure 22C, quantization of the DC coefficients for the top-half blocks of macroblock 2250 results in uniform values reconstructed from a DC value of 68, while quantization of DC coefficients for the bottom-half blocks results in uniform values reconstructed from a DC value of 70.

15 Accordingly, several techniques and tools described below are used by a video encoder to detect DC shift areas and adjust quantization to reduce or avoid introduction of contouring artifacts in the DC shift areas.

Figure 32 is a flow chart showing a generalized technique 3200 for adjusting 20 quantization to reduce or avoid introduction of contouring artifacts in DC shift areas. An encoder such as the encoder 1000 of Figure 10 or other tool performs the technique 3200.

At 3210, an encoder detects a shift area. The search for DC shift areas can be aided by previous gradient slope detection. For example, the encoder detects DC shift areas by detecting one or more gradient slope regions (or using previously computed 25 gradient slope detection information) then identifying DC shift blocks in the gradient slope region(s), as described below.

At 3220, the encoder adjusts quantization in the DC shift area. For example, an encoder can use differential quantization (DQ) to code DC shift blocks in order to reduce or avoid introduction of contouring artifacts caused by DC shift. The encoder reduces QP for some macroblocks (those with DC shift blocks) but does not change QP for other 30 blocks. Reducing QP for macroblocks having DC shift blocks can help keep DC values lossless for the macroblocks, thereby reducing or avoiding introduction of contouring artifacts. An encoder can use bi-level DQ or multi-level DQ to resolve DC shift problems and thereby improve visual quality while controlling bit usage. If there are more pictures to analyze, the encoder processes the next picture (3230).

Alternatively, the encoder adjusts quantization for DC shift areas on a macroblock-by-macroblock basis or some other basis.

#### A. Gradient Slope Detection

Gradient slope detection can be used to identify one or more gradient slope regions 5 in a picture. The gradient slope region(s) tend to exhibit contouring artifacts, especially when blocks in the region(s) have non-zero DC coefficient values and AC coefficients of only zero after quantization. Once found, such region(s) can be checked for DC shift blocks that may contribute to contouring artifacts.

For example, an encoder finds a gradient slope region using a technique described 10 herein (Section V) or some other technique. If the only non-zero coefficients in blocks are DC coefficients after quantization, the encoder treats the blocks as candidates for DC shift area adjustment. Alternatively, the encoder considers additional blocks as candidates for DC shift area adjustment.

#### B. Identifying DC-shift blocks

15 The encoder identifies certain candidate blocks as DC shift blocks. The identification of DC shift blocks depends on details of the quantizer and QPs used to compress the blocks. For example, some reconstructed DC coefficients will not shift from their original value at one QP, but will shift at a coarser QP.

Examples of DC shift coefficients for different QPs in one encoder are provided in 20 the following table. The table indicates DC coefficient values exhibiting DC shift for different values of QP, where QP is derived explicitly from the parameter PQIndex (and, potentially, a half step parameter) or implicitly from the parameter PQIndex (and, potentially, a half step parameter). DC values not listed in the table are lossless for the indicated QP in the example encoder; DC values for QPs under 3 (which are not shown in 25 the table) are all lossless. The example encoder does not perform DC shift adjustment for QPs higher than those shown in the table. In the example encoder, quantization of DC coefficients is the same for different quantizers (e.g., uniform, non-uniform). Which DC coefficient values are DC shift coefficients will differ in different video codecs. Different quantizers (e.g., uniform, non-uniform) can result in different shift patterns if quantization 30 of DZ coefficients is different in the different quantizers.

Table 3: Example DC-shift Coefficients

<u>PQIndex</u> <u>(Implicit</u> <u>QP)</u>	<u>PQIndex</u> <u>(Explicit</u> <u>QP)</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>
3-5.5	3-5.5	6	7	96	97	186	187
		15	16	105	106	195	196
		24	25	114	115	204	205
		33	34	123	124	213	214
		42	43	132	133	222	223
		51	52	141	142	231	232
		60	61	150	151	240	241
		69	70	159	160	249	250
		78	79	168	169		
		87	88	177	178		
6- 7.5, 9- 10.5	6-7.5	2	1	92	93	178	177
		6	7	97	96	183	182
		11	12	102	101	187	188
		16	15	106	107	192	193
		21	20	111	112	197	196
		25	26	116	115	202	201
		30	31	121	120	207	206
		35	34	126	125	211	212
		40	39	158	157	216	217
		45	44	130	131	221	220
		49	50	135	136	225	226
		54	53	140	139	230	231
		59	58	144	145	235	236
		63	64	149	150	240	239
		68	69	154	155	245	244
		73	74	159	158	249	250
		78	77	164	163	254	255
		83	82	168	169		
		87	88	173	174		

<u>PQIndex</u> <u>(Implicit</u> <u>QP)</u>	<u>PQIndex</u> <u>(Explicit</u> <u>QP)</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>
8, 11 -	8 - 9	2	1	88	89	171	172
12		5	6	92	91	175	174
		9	8	95	96	178	179
		12	11	99	98	182	181
		15	16	102	101	185	186
		19	18	105	106	189	188
		22	23	109	108	192	191
		26	25	112	113	195	196
		29	30	116	115	199	198
		33	32	119	120	202	203
		36	37	123	122	206	205
		40	39	126	127	209	210
		43	44	158	156	213	212
		47	46	130	129	216	217
		50	51	133	134	220	219
		54	53	137	136	223	224
		57	56	140	141	227	226
		60	61	144	143	230	231
		64	63	147	146	234	233
		67	68	150	151	237	236
		71	70	154	153	240	241
		74	75	157	158	244	243
		78	77	161	160	247	248
		81	82	164	165	251	250
		85	84	168	167	254	255

<u>PQIndex</u> <u>(Implicit</u> <u>QP)</u>	<u>PQIndex</u> <u>(Explicit</u> <u>QP)</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>
13-14	10-11	2	3	90	89	175	174
		5	4	93	92	177	178
		8	7	95	96	180	181
		11	10	98	99	183	184
		13	14	101	102	186	185
		16	17	104	103	189	188
		19	20	107	106	192	191
		22	21	110	109	194	195
		25	24	112	113	197	198
		27	28	115	116	200	201
		30	31	118	119	203	202
		33	34	121	120	206	205
		36	35	124	123	209	208
		39	38	126	127	211	212
		42	41	158	157	214	215
		45	44	129	130	217	218
		47	48	132	133	220	219
		50	51	135	136	223	222
		53	52	138	137	225	226
		56	55	141	140	228	229
		59	58	144	143	231	232
		61	62	146	147	234	235
		64	65	149	150	237	236
		67	68	152	151	240	239
		70	69	155	154	243	242
		73	72	158	157	245	246
		76	75	160	161	248	249
		78	79	163	164	251	250
		81	82	166	167	254	253
		84	85	169	168		
		87	86	172	171		

<u>PQIndex</u> <u>(Implicit</u> <u>QP)</u>	<u>PQIndex</u> <u>(Explicit</u> <u>QP)</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>
15-16	12-13	2	1	87	88	171	170
		4	3	90	89	173	174
		6	7	92	93	176	175
		9	8	95	94	178	179
		11	12	97	98	181	180
		14	13	100	99	183	184
		16	17	102	103	186	185
		19	18	105	104	188	189
		21	22	107	108	191	190
		24	23	110	109	193	194
		26	27	112	111	195	196
		29	28	114	115	198	197
		31	30	117	116	200	201
		33	34	119	120	203	202
		36	35	122	121	205	206
		38	39	124	125	208	207
		41	40	127	126	210	211
		43	44	158	157	213	212
		46	45	129	130	215	216
		48	49	132	131	218	217
		51	50	134	135	220	221
		53	54	137	136	222	223
		56	55	139	140	225	224
		58	57	141	142	227	228
		60	61	144	143	230	229
		63	62	146	147	232	233
		65	66	149	148	235	234
		68	67	151	152	237	238
		70	71	154	153	240	239
		73	72	156	157	242	243
		75	76	159	158	245	244
		78	77	161	162	247	248
		80	81	164	163	249	250
		83	82	166	167	252	251
		85	84	168	169	254	255

<u>PQIndex</u> <u>(Implicit</u> <u>QP)</u>	<u>PQIndex</u> <u>(Explicit</u> <u>QP)</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>	<u>Original</u> <u>DC</u>	<u>Shifted</u> <u>DC</u>
17-18	14-15	1	2	87	88	171	172
		3	4	89	90	173	174
		5	6	92	91	175	176
		8	7	94	93	178	177
		10	9	96	95	180	179
		12	13	98	99	182	183
		14	15	100	101	184	185
		16	17	103	102	186	187
		19	18	105	104	189	188
		21	20	107	106	191	190
		23	24	109	110	193	194
		25	26	111	112	195	196
		27	28	114	113	198	197
		30	29	116	115	200	199
		32	31	118	119	202	201
		34	35	120	121	204	205
		36	37	122	123	206	207
		39	38	125	124	209	208
		41	40	127	126	211	210
		43	42	158	157	213	212
		45	46	129	130	215	216
		47	48	131	132	217	218
		50	49	133	134	220	219
		52	51	136	135	222	221
		54	53	138	137	224	223
		56	57	140	141	226	227
		58	59	142	143	228	229
		61	60	144	145	231	230
		63	62	147	146	233	232
		65	66	149	148	235	236
		67	68	151	152	237	238
		69	70	153	154	239	240
		72	71	156	155	242	241
		74	73	158	157	244	243
		76	77	160	159	246	247
		78	79	162	163	248	249
		81	80	164	165	250	251
		83	82	167	166	253	252
		85	84	169	168	255	254

The example encoder with the DC shift coefficients shown in Table 3 generally uses different QPs for textured regions than for smooth regions. The example encoder typically uses a QP in the range of 3-5 to encode smooth regions. As shown in Table 3, above, for QP 3-5, all the shifted DC values are 1 more than the original DC value. Other encoders may use different QPs for smooth regions versus texture regions.

To help reduce or avoid introduction of contouring artifacts when DC shift blocks are detected, the example encoder changes the QP for macroblocks containing DC shift blocks to keep the DC values lossless in those macroblocks. In particular, the example encoder reduces the QP for macroblocks containing DC shift blocks to QP = 2. (Other 5 encoders may use some other QP for DC shift areas.) In general, an encoder can select the largest available QP that results in lossless treatment of the DC coefficients of the blocks of the macroblock.

An encoder calculates a mean luma value per block to determine DC shift blocks in the gradient slope region(s), since the mean luma value corresponds to the DC shift 10 value (after compensating for expansion in the transform). The mean luma value allows the encoder to estimate or predict which blocks have DC shifts. Alternatively, an encoder calculates real DC values and looks them up in the DC shift table to identify exactly which blocks will have shifts.

The encoder can perform additional processing to exclude certain isolated DC shift 15 blocks in the gradient slope region(s). In the example encoder, once a current block is identified as a DC shift block located in a gradient slope region, the surrounding four neighboring blocks are checked. If any of the surrounding four neighboring blocks is a smooth block and has a DC value lower than the shifted DC value of the current block, the encoder uses QP = 2 for the macroblock containing the current block in order to keep 20 the DC values lossless. Alternatively, an encoder does not do a check of neighboring blocks, or checks some other arrangement of neighboring blocks to determine whether to make a change in the QP for the DC shift area.

### **C. Multi-level Differential Quantization Cost Model**

Bi-level DQ and multi-level DQ typically have different bit rate costs. In one 25 implementation, 1 bit per macroblock is used to signal a picture QP or alternative QP in “all macroblock” bi-level DQ, and at least 3 bits per macroblock are used to signal a picture QP or alternative QPs in multi-level DQ.

Although an encoder can use multi-level DQ to allow for reducing QP in a smooth region that contains DC shift blocks, an encoder instead can choose to adjust the QP for all 30 smooth regions (e.g., to QP = 2) and use a coarser picture QP for the rest of the picture in a bi-level DQ scenario. For example, an encoder may do this where the signaling costs of multi-level DQ are found to be too expensive relative to the costs of bi-level DQ.

In one implementation, the following table is used to calculate the cost of smooth blocks that going from QP = 3, 4, 5, and 6, respectively, to QP = 2.

*g\_iSmoothBlockDiffQPCost[4] = {18, 22, 28, 36}.*

This table is used in the following example of bi-level DQ cost B(QP) cost calculation.

5           
$$B(QP) = \text{counts\_of\_total\_MBs} + (\text{counts\_of\_biLevel\_Dquan\_MBs} - \text{counts\_of\_DC\_Shift\_Blocks}) * g_iSmoothBlockDiffQPCost[QP-3];$$

The cost B(QP) accounts for the costs of per-macroblock bi-level cost signaling and estimates the increased bit cost of using QP = 2 instead of a 3, 4, 5, or 6 for a block. Multi-level DQ cost  $M(QP)$  is calculated as:

10           
$$M(QP) = (\text{counts\_of\_frameQP\_MBs} * 3) + (\text{counts\_of\_biLevel\_Dquan\_MBs} - \text{counts\_of\_DC\_Shift\_Blocks}) * 8 + (\text{counts\_of\_DC\_Shift\_Blocks} * 3);$$

The cost  $M(QP)$  accounts for signaling costs of multi-level DQ, assuming escape coding for some macroblock quantization parameters. If  $B(qp) < M(qp)$ , then bi-level DQ will be used and QP = 2 will be used for all smooth blocks. Otherwise, multi-level DQ will be used.

15           Alternatively, an encoder uses other costs models for different types or configurations of DQ. Or, an encoder reduces QP for the entire picture when DC shift blocks are detected, or uses some other technique to change quantization to reduce or avoid introduction contouring artifacts when DC shift blocks are detected.

#### D. Picture QP switching

20           In one example encoder, multi-level DQ requires 3 bits to signal any QP from picture QP to picture QP+6. Any QP outside of this range will be signaled with 8 bits through escape coding. Alternative QPs that are used for smooth regions are normally smaller than the picture QP, and hence require escape coding.

25           Switching picture QPs can thus save coding overhead for multi-level DQ. For example, an encoder can choose a picture QP using the multi-level DQ cost model described above. For example, for a three-level scenario (e.g., a picture QP, a smooth region QP, and a DC shift QP), multi-level DQ cost is computed for different candidate values for a picture QP. An encoder can select the picture QP with minimum overhead cost.

30           Alternatively, an encoder uses other criteria to switch picture QPs, or does not perform picture QP switching.

#### E. Coarse Quantization for High-texture Macroblocks

If a decision is made to use multi-level DQ, coarse quantization can be used for high-texture macroblocks by assigning them a higher (coarser) QP than the picture QP.

The decision to use multi-level DQ for the picture (e.g., in order to use smaller QP for DC shift macroblocks) means there is no additional overhead cost to signal a per macroblock coarse QP that is higher than the picture QP. For example, picture QP + 1 can be used as the coarse QP to avoid noticeable differences in the quantization level, or some other QP can 5 be used. A texture threshold can be used to determine which macroblocks will be quantized with the coarse QP, after the encoder has decided to use multi-level DQ for the current picture.

Alternatively, an encoder uses other criteria to determine whether certain regions (e.g., macroblocks) should use a coarse QP, or does not use coarse QPs.

#### 10 F. Example Technique for DC Shift Quantization

Figure 33 is a flow chart showing a combined technique 3300 for tailoring quantization in DC shift areas to reduce or avoid introduction of quantization artifacts. An encoder such as the encoder 1000 of Figure 10 or other tool performs the technique 3300. This combined technique is an example that combines several of the aspects described 15 above. Other techniques will not use all of the aspects described with reference to this example, or will perform them in a different order or in alternative ways.

At 3310, an encoder detects one or more gradient slope regions in a current picture, for example, as described in Section V. At 3320, the encoder detects DC shift blocks in the gradient slope region(s), for example, using a DC shift table.

20 The encoder then decides how to quantize the picture. At 3330, the encoder decides whether to use bi-level DQ for the picture. If it does, the encoder uses a QP smaller than the picture QP for DC shift areas (3340) and other smooth areas. Otherwise, at 3350, the encoder decides whether to use multi-level DQ for the picture. If it does, the encoder uses a QP smaller than the picture QP for DC shift areas (3360), can use a 25 different QP for other smooth areas, and, if high-texture macroblocks are present, uses a coarse QP (such as one that is larger than the picture QP) for the high-texture macroblocks (3370). If the encoder does not choose bi-level or multi-level DQ, the encoder reduces the picture QP and uses the reduced picture QP for DC shift areas (3380) as well as other areas. Or, the encoder skips QP reduction for the DC shift areas if the costs of bi-level DQ 30 and multi-level DQ are both too high. When the encoder has chosen a quantization scheme, the encoder compresses the picture at 3390, and processes the next picture if any pictures remain to be processed (3395).

Having described and illustrated the principles of our invention with reference to various embodiments, it will be recognized that the various embodiments can be modified

in arrangement and detail without departing from such principles. It should be understood that the programs, processes, or methods described herein are not related or limited to any particular type of computing environment, unless indicated otherwise. Various types of general purpose or specialized computing environments may be used with or perform operations in accordance with the teachings described herein. Elements of embodiments shown in software may be implemented in hardware and vice versa.

5 In view of the many possible embodiments to which the principles of our invention may be applied, we claim as our invention all such embodiments as may come within the scope and spirit of the following claims and equivalents thereto.

We claim:

1. In a video encoder, a method comprising:
  - identifying a current smooth region of a current video picture in a video picture sequence, the video picture sequence having a display order in which display of plural previous pictures occurs before display of the current video picture in the display order and display of plural future pictures occurs after the display of the current video picture in the display order;
  - performing temporal analysis on the current smooth region, wherein the temporal analysis comprises determining whether a corresponding region in at least one of the plural previous and/or future pictures is smooth;
  - based at least in part on the temporal analysis, adjusting quantization in the current smooth region; and
  - outputting encoded data for the current video picture.
2. The method of claim 1 wherein the adjusting quantization in the current smooth region comprises reducing a quantization step size for the current smooth region.
3. The method of claim 2 wherein the reducing the quantization step size for the current smooth region avoids introduction of a contouring artifact in the current smooth region when the current video picture is reconstructed.
4. The method of claim 1 wherein the temporal analysis is performed on an adjustable number of the plural future pictures, and wherein the adjustable number is one or greater.
5. The method of claim 1 wherein the temporal analysis attributes greater weight to smoothness in a corresponding region of a future picture that is temporally nearer to the current video picture than to smoothness in a corresponding region of a future picture that is temporally more distant from the current video picture.
6. The method of claim 1 wherein the current smooth region is a gradient slope region.
7. The method of claim 1 wherein the video encoder analyzes a texture map to identify the current smooth region, and wherein the video encoder analyzes mean luminance values in the temporal analysis.
8. One or more computer-readable media having stored thereon computer executable instructions to cause a computer to perform the method of claim 1.
9. In a video encoder, a method comprising:
  - determining a differential quantization interval for a video picture sequence, the interval comprising an interval number;

using the interval when performing differential quantization for plural predicted differentially quantized pictures in the video picture sequence, wherein the interval constrains the encoder to skip differential quantization for at least the interval number of predicted pictures after one of the plural predicted differentially quantized pictures; and

5                   outputting encoded data for the video picture sequence.

10.       The method of claim 9 wherein the differential quantization comprises:

                  selecting one or more differential quantization step sizes for a first predicted picture of the plural predicted differentially quantized predicted pictures, the one or more differential quantization step sizes differing from a picture quantization step size for the first predicted picture;

10                   changing the one or more differential quantization step size for a second predicted picture of the plural predicted differentially quantized predicted pictures, wherein the second predicted picture is outside the interval from the first predicted picture.

15.       The method of claim 9 wherein the plural predicted differentially quantized pictures comprise at least one P-picture.

12.       The method of claim 9 wherein the plural predicted differentially quantized pictures comprise at least one B-picture.

13.       The method of claim 9 wherein the differential quantization interval is a fixed 20 interval.

14.       The method of claim 9 wherein the differential quantization interval is an adaptively adjustable interval.

15.       One or more computer-readable media having stored thereon computer executable instructions to cause a computer to perform the method of claim 9.

25.       16.      In a video encoder, a method comprising:

                  analyzing texture in a current video picture;

                  setting a smoothness threshold for the current video picture based at least in part on the analyzed texture in the current video picture;

                  comparing texture data of the current video picture with the smoothness threshold;

30                   adjusting differential quantization for at least part of the current video picture based on a finding of at least one smooth region in the current video picture according to the smoothness threshold; and

                  outputting encoded data for the current video picture.

17. The method of claim 16 wherein the analyzing texture comprises analyzing a texture map.
18. The method of claim 16 wherein the analyzing texture comprises applying a sliding window to a gradient value histogram of block gradient values.
- 5 19. The method of claim 16 wherein the adjusting differential quantization comprises: determining a percentage of flat blocks in the current video picture; and comparing the percentage to one or more percentage thresholds, wherein the adjusting differential quantization depends on results of the comparing the percentage to the one or more percentage thresholds.
- 10 20. The method of claim 16 wherein the adjusting differential quantization comprises: identifying an isolated flat block in a texture region in the current video picture; and skipping the differential quantization for the isolated flat block.

1 / 23

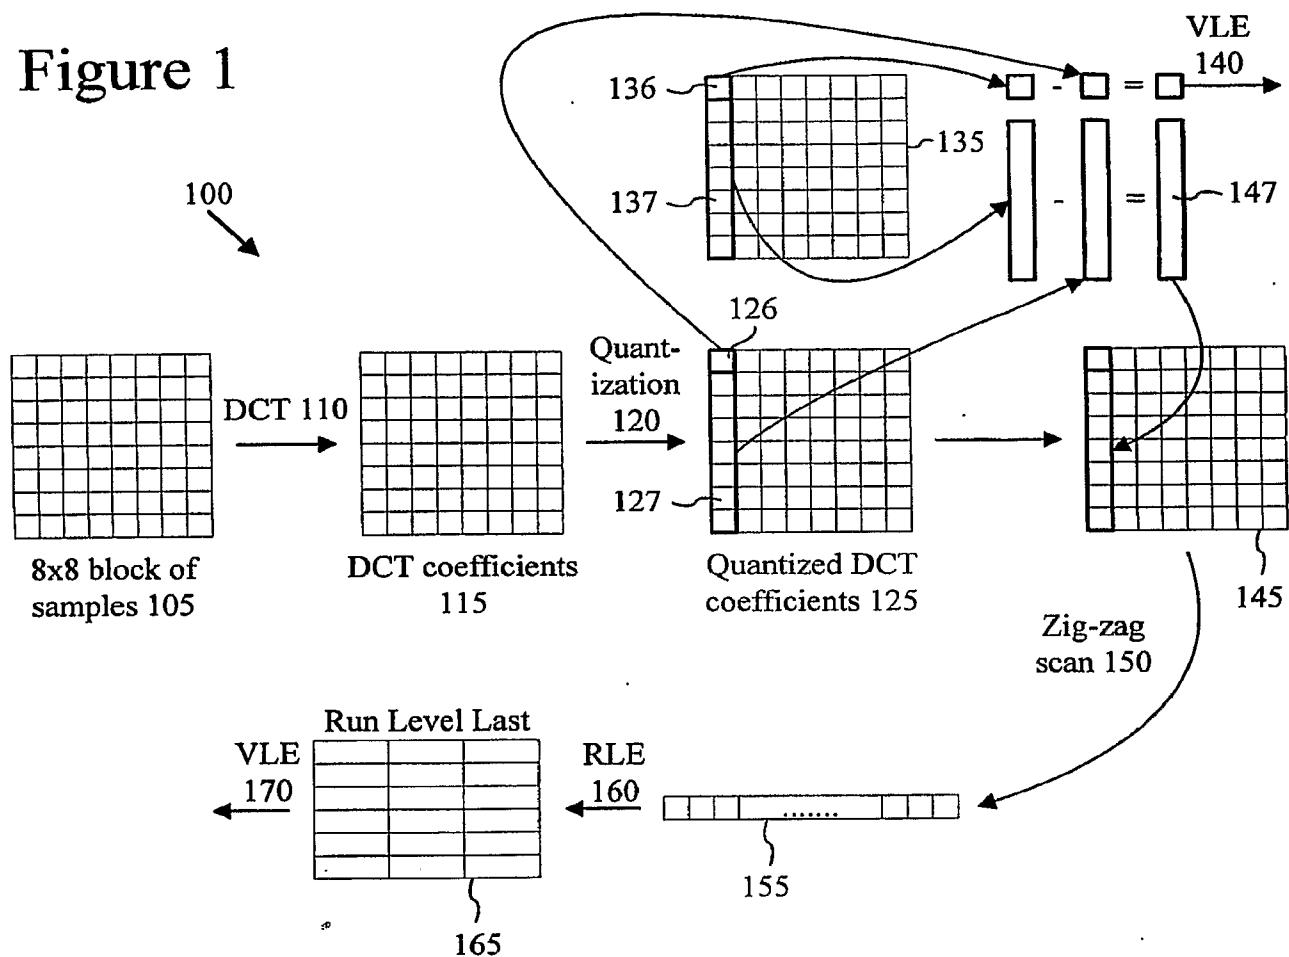
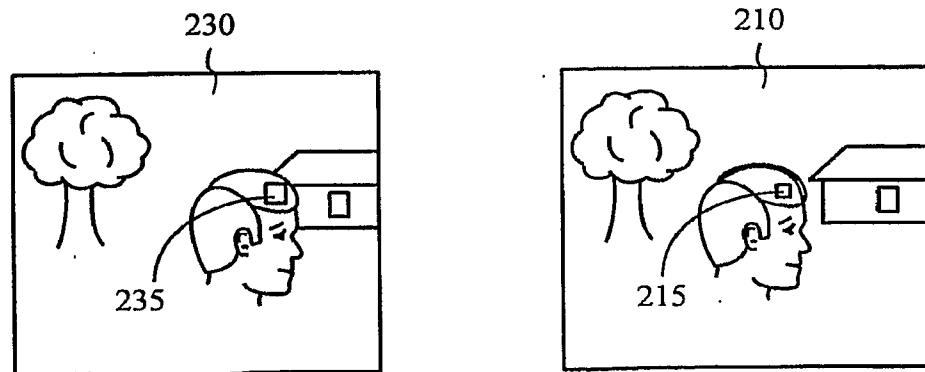
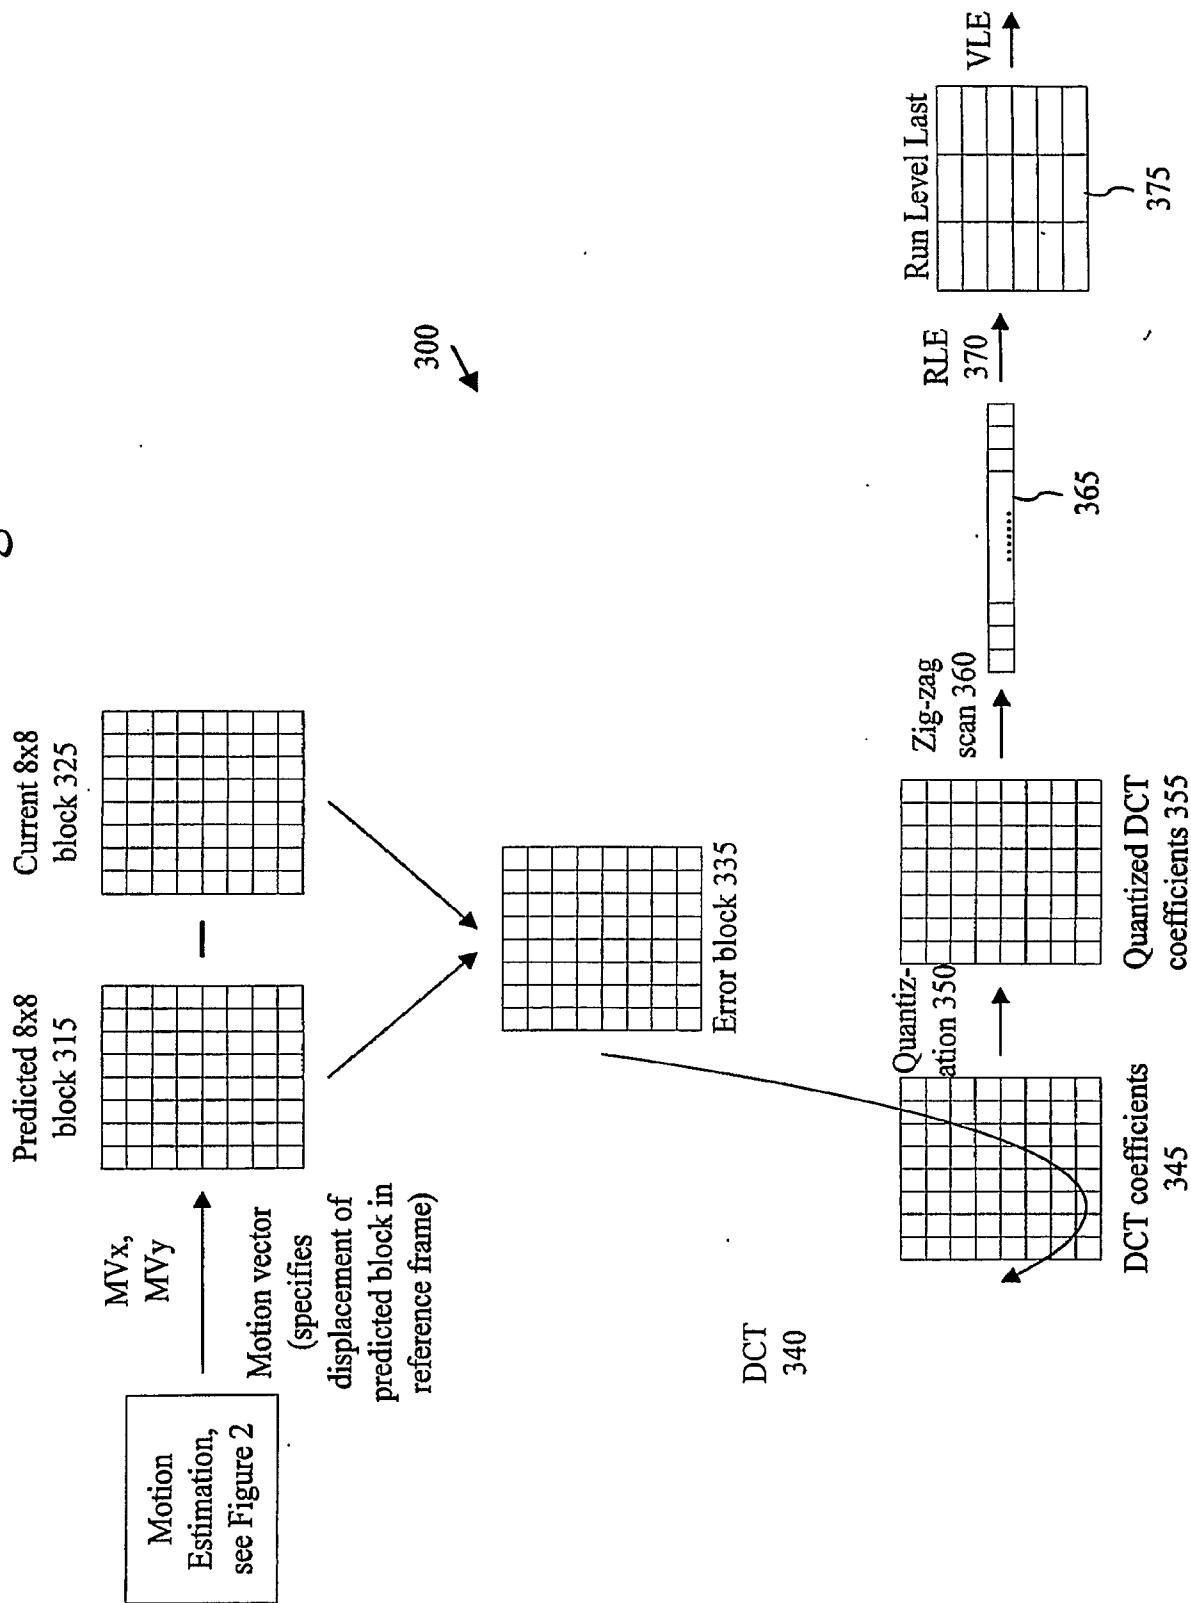
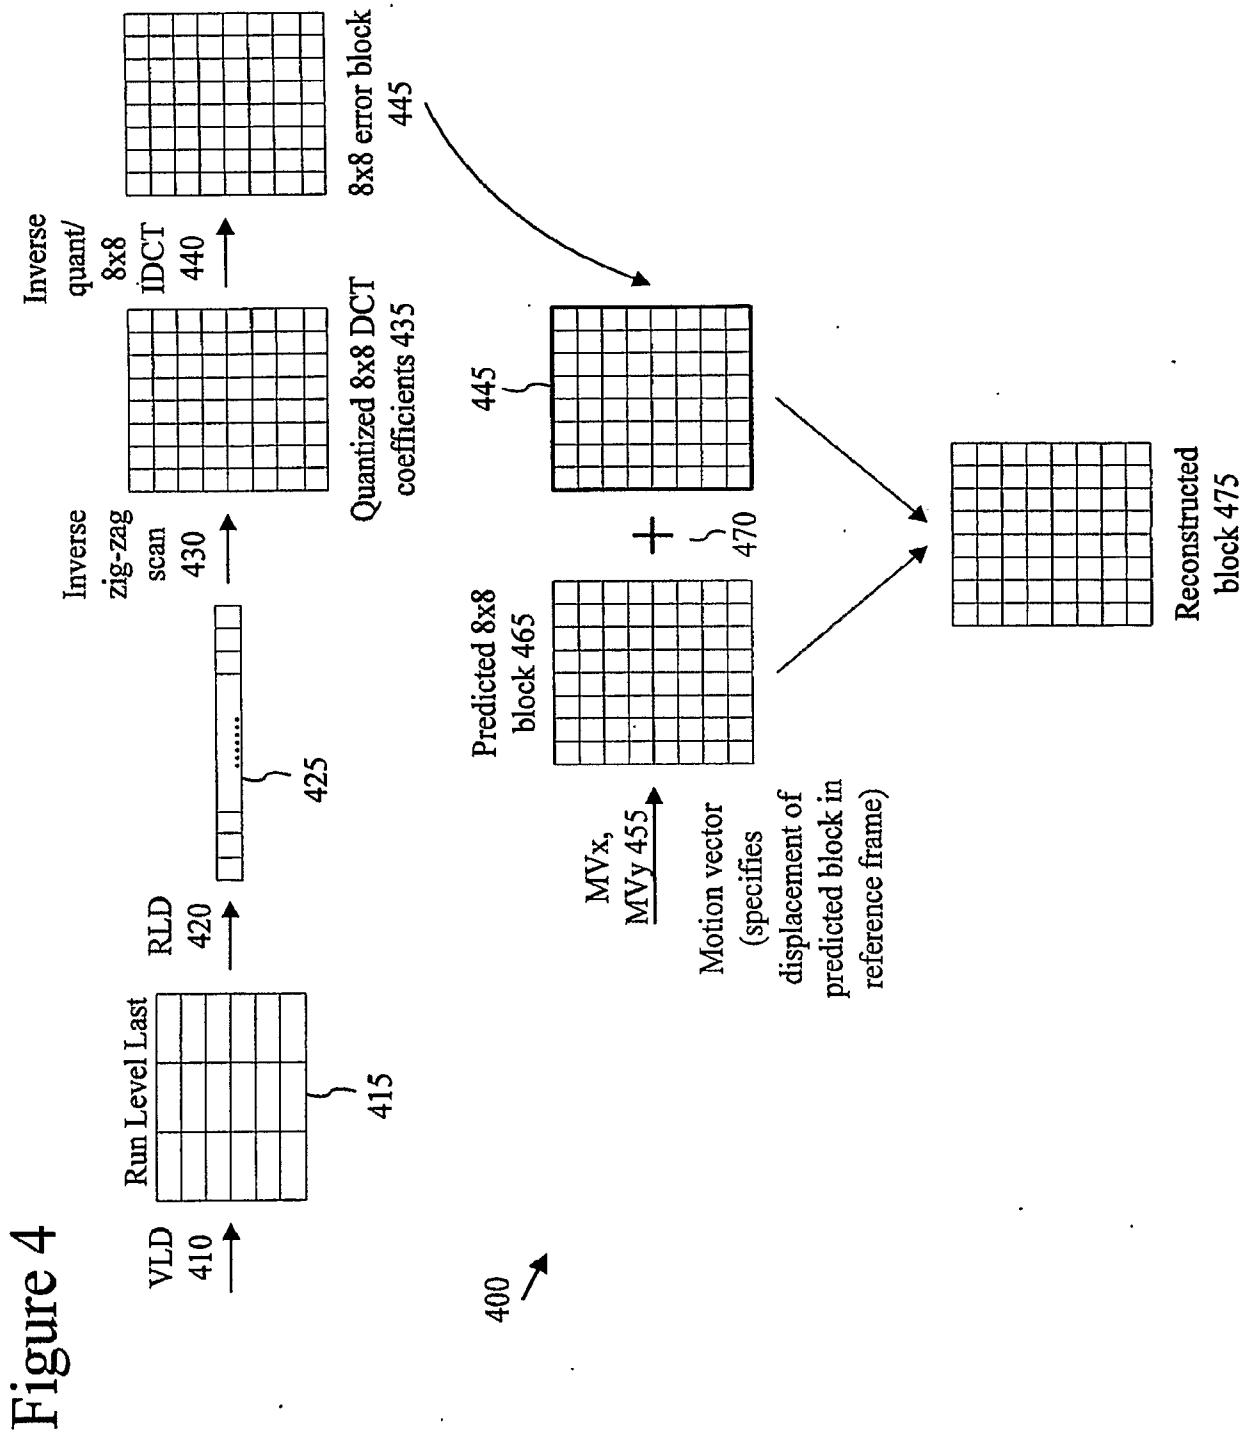
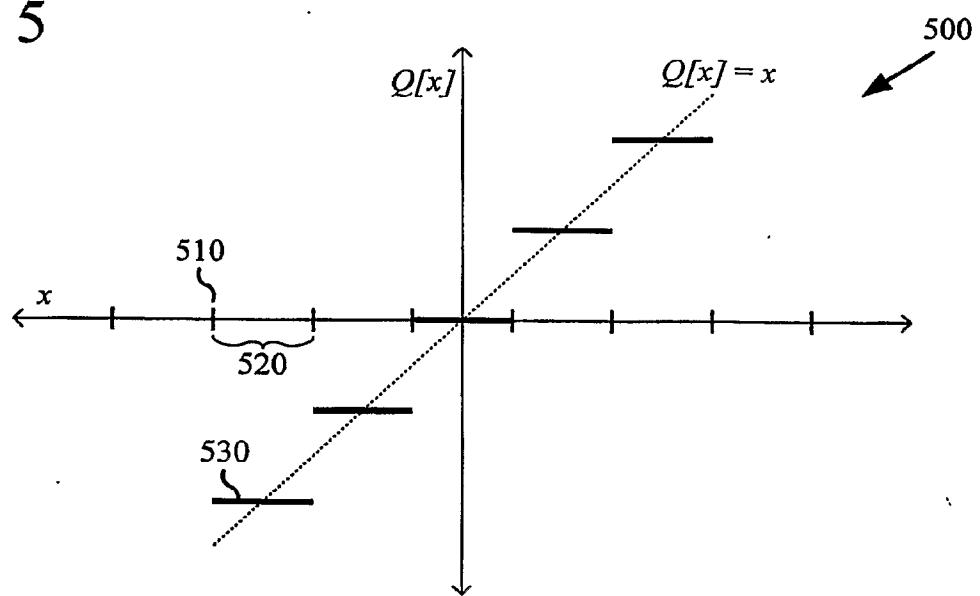
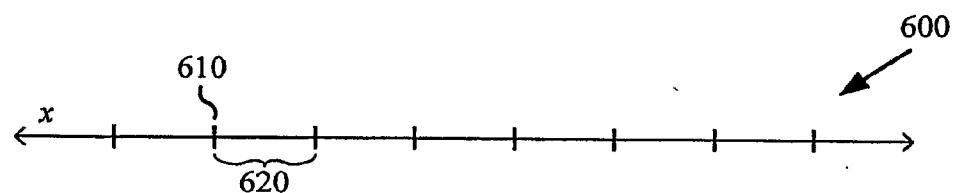
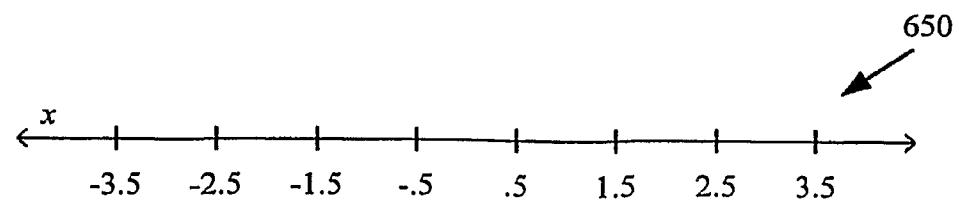
**Figure 1****Figure 2**

Figure 3





4 / 23

**Figure 5****Figure 6A****Figure 6B**

5 / 23

Figure 7

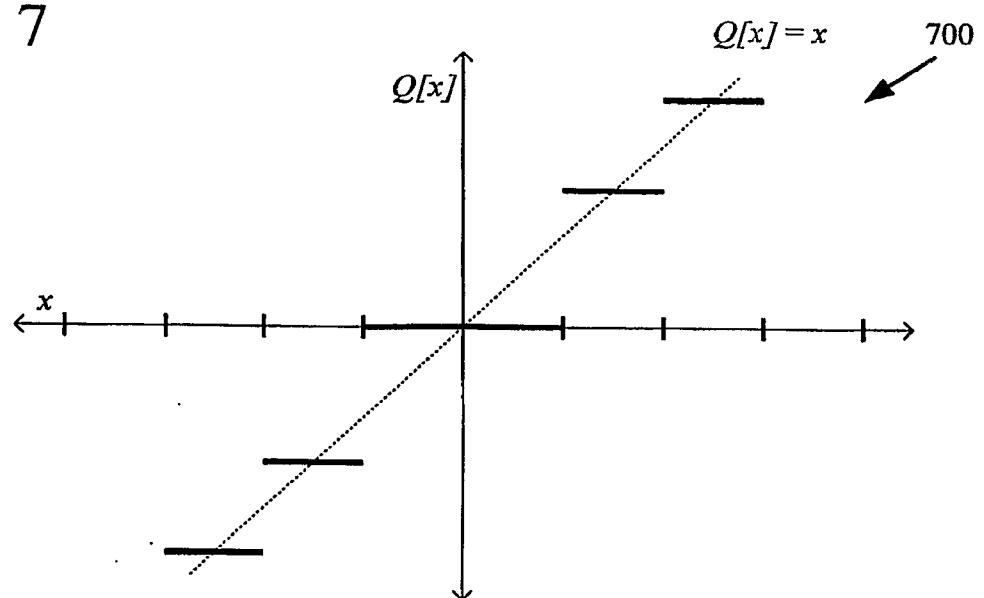


Figure 8A

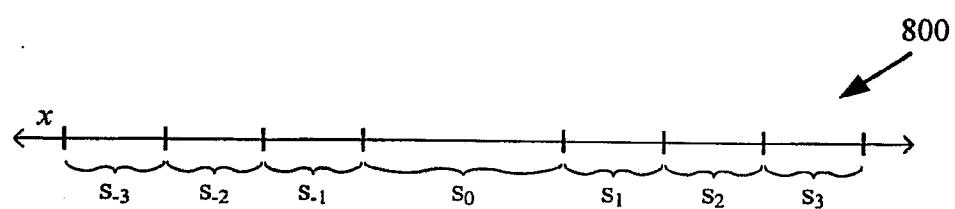
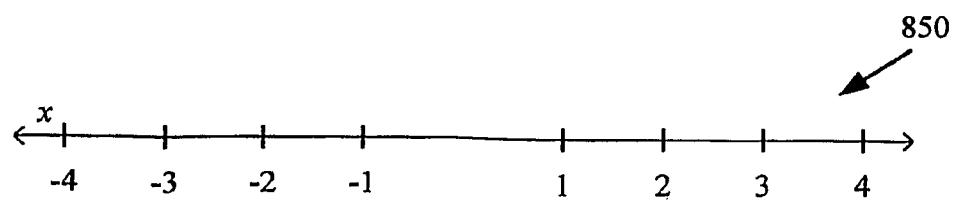
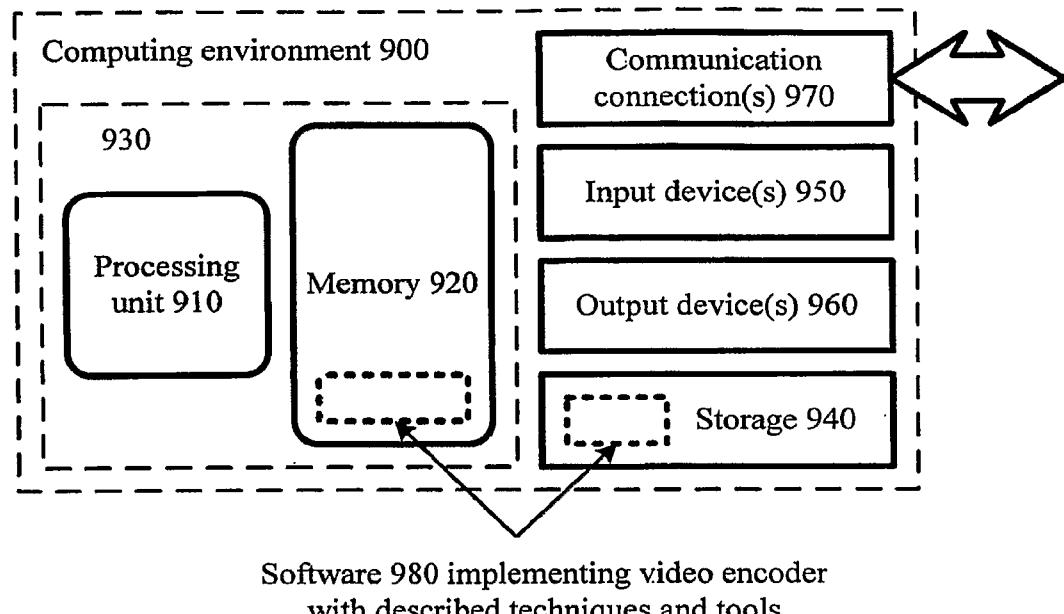
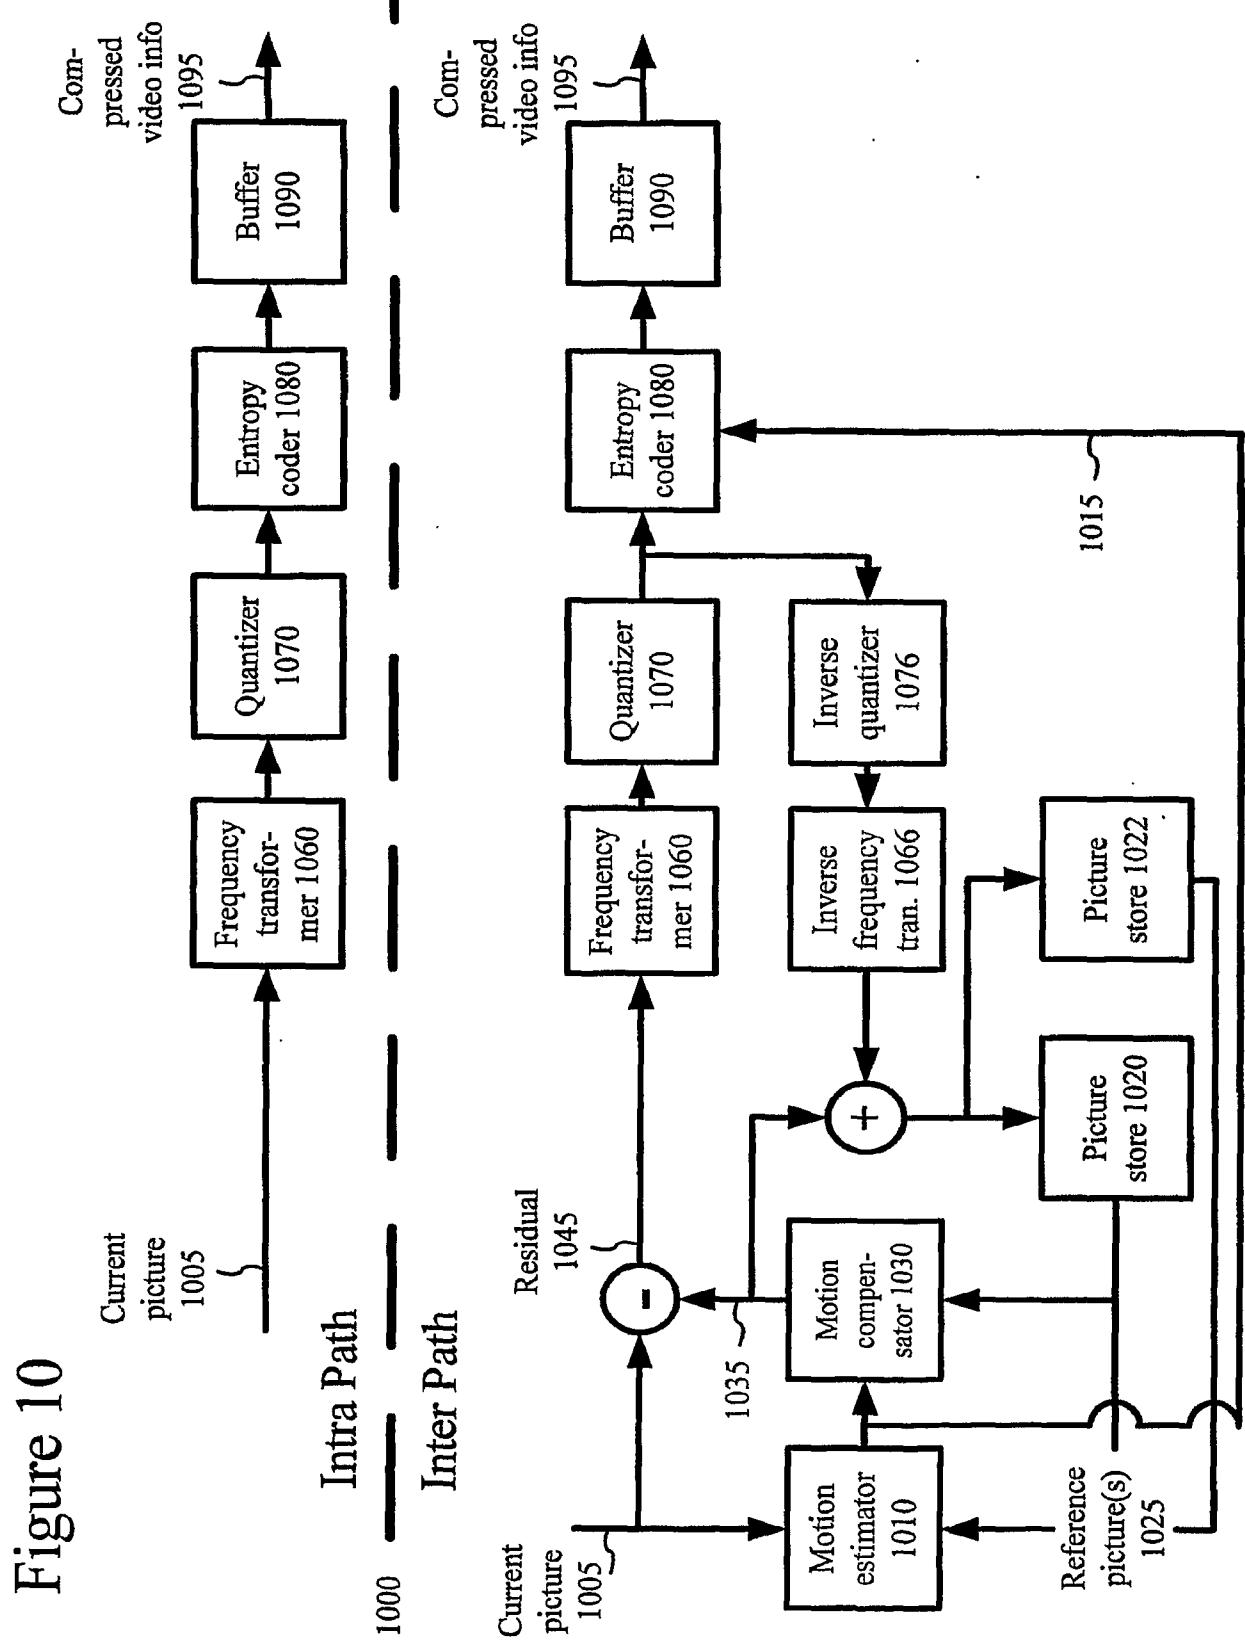


Figure 8B



**Figure 9**



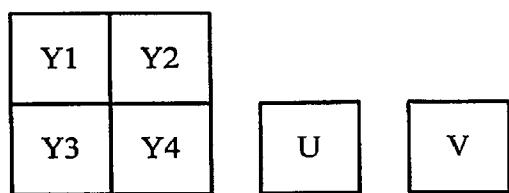
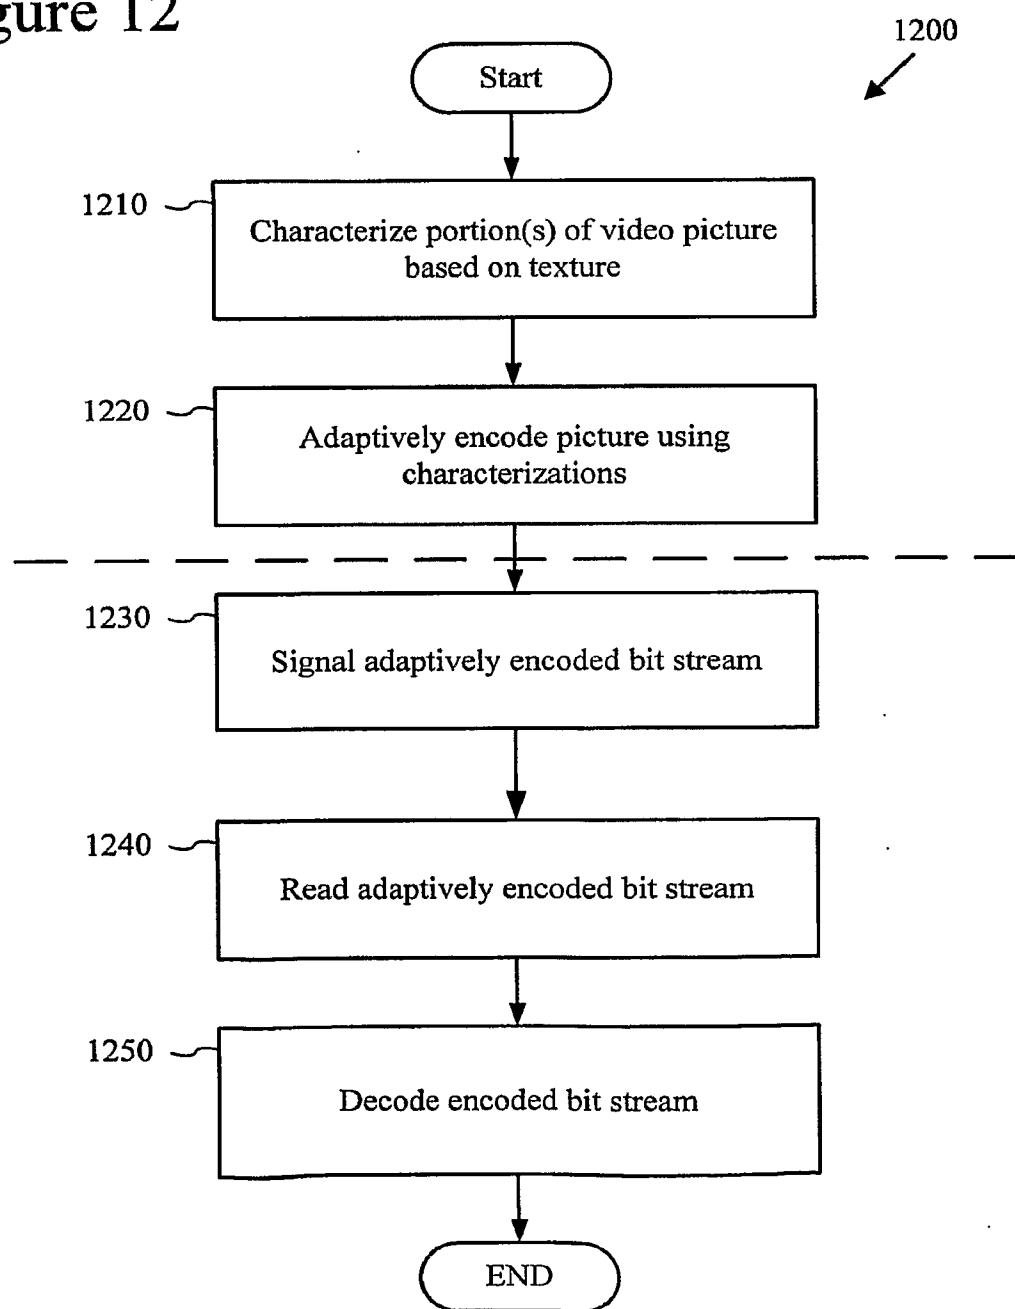
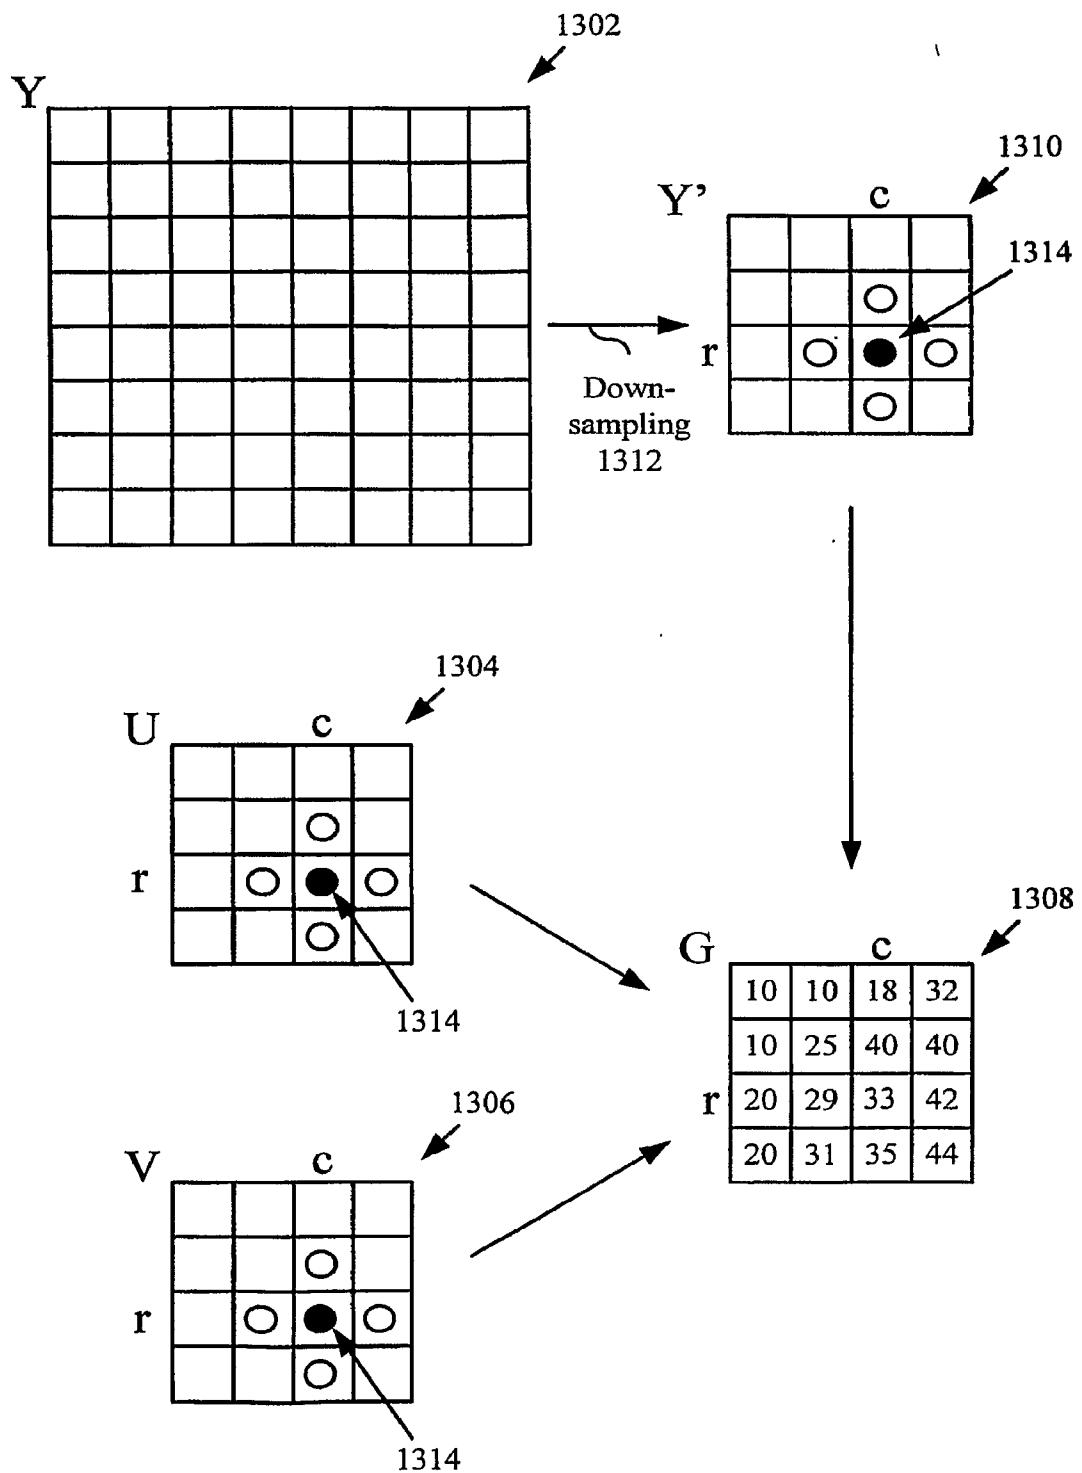
**Figure 11****Figure 12**

Figure 13



10 / 23

Figure 14

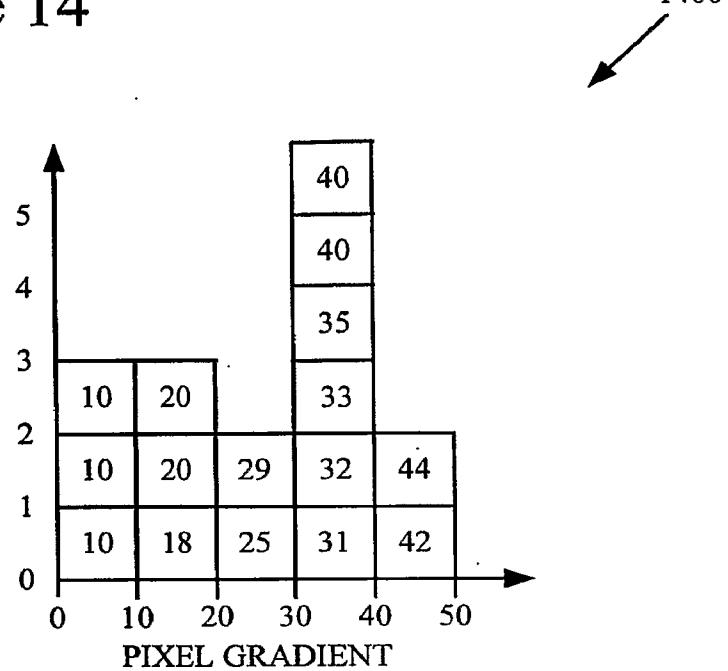


Figure 15

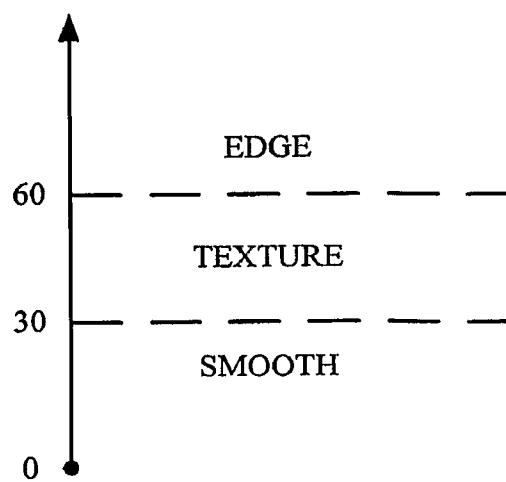


Figure 16

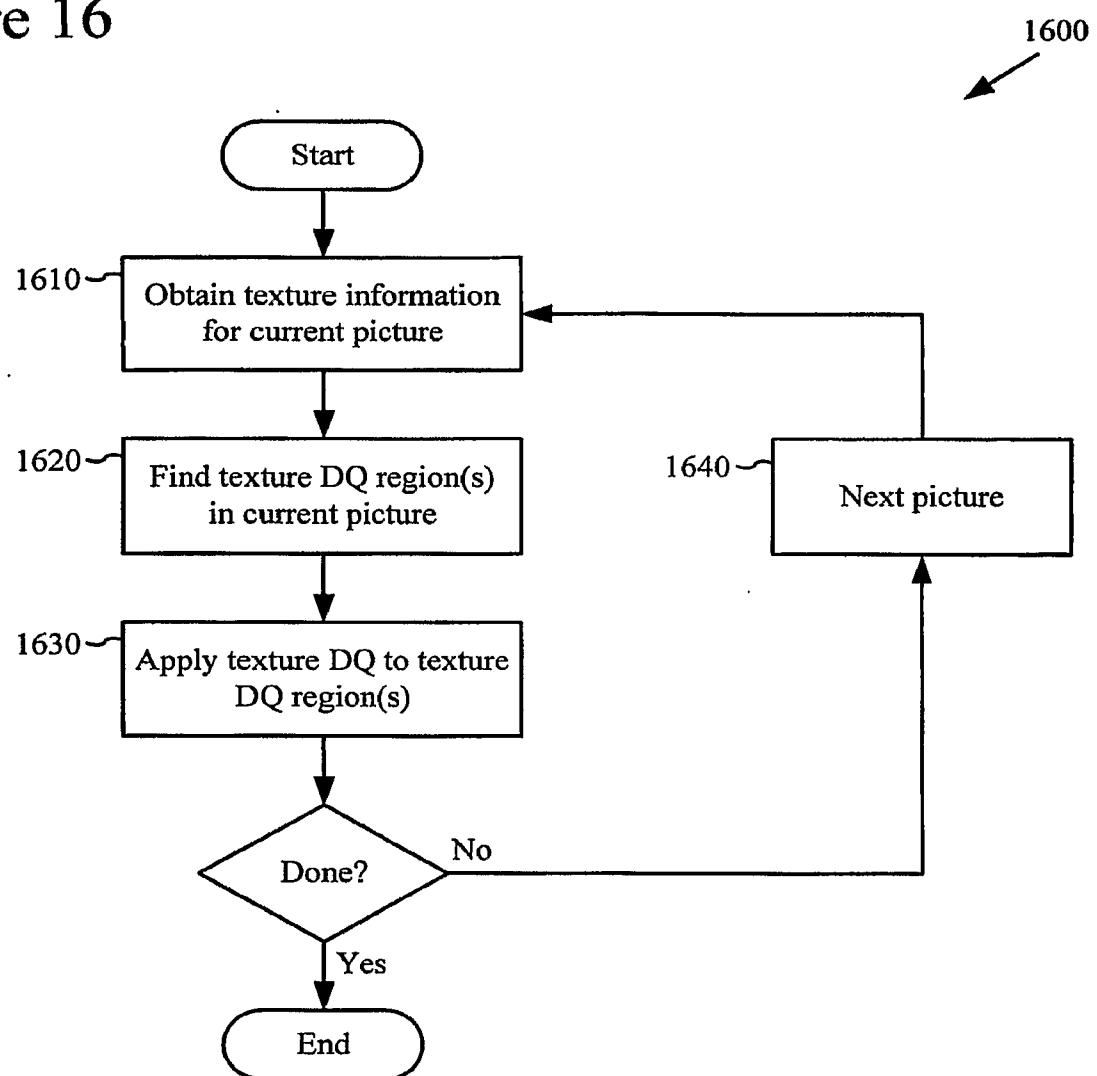


Figure 17

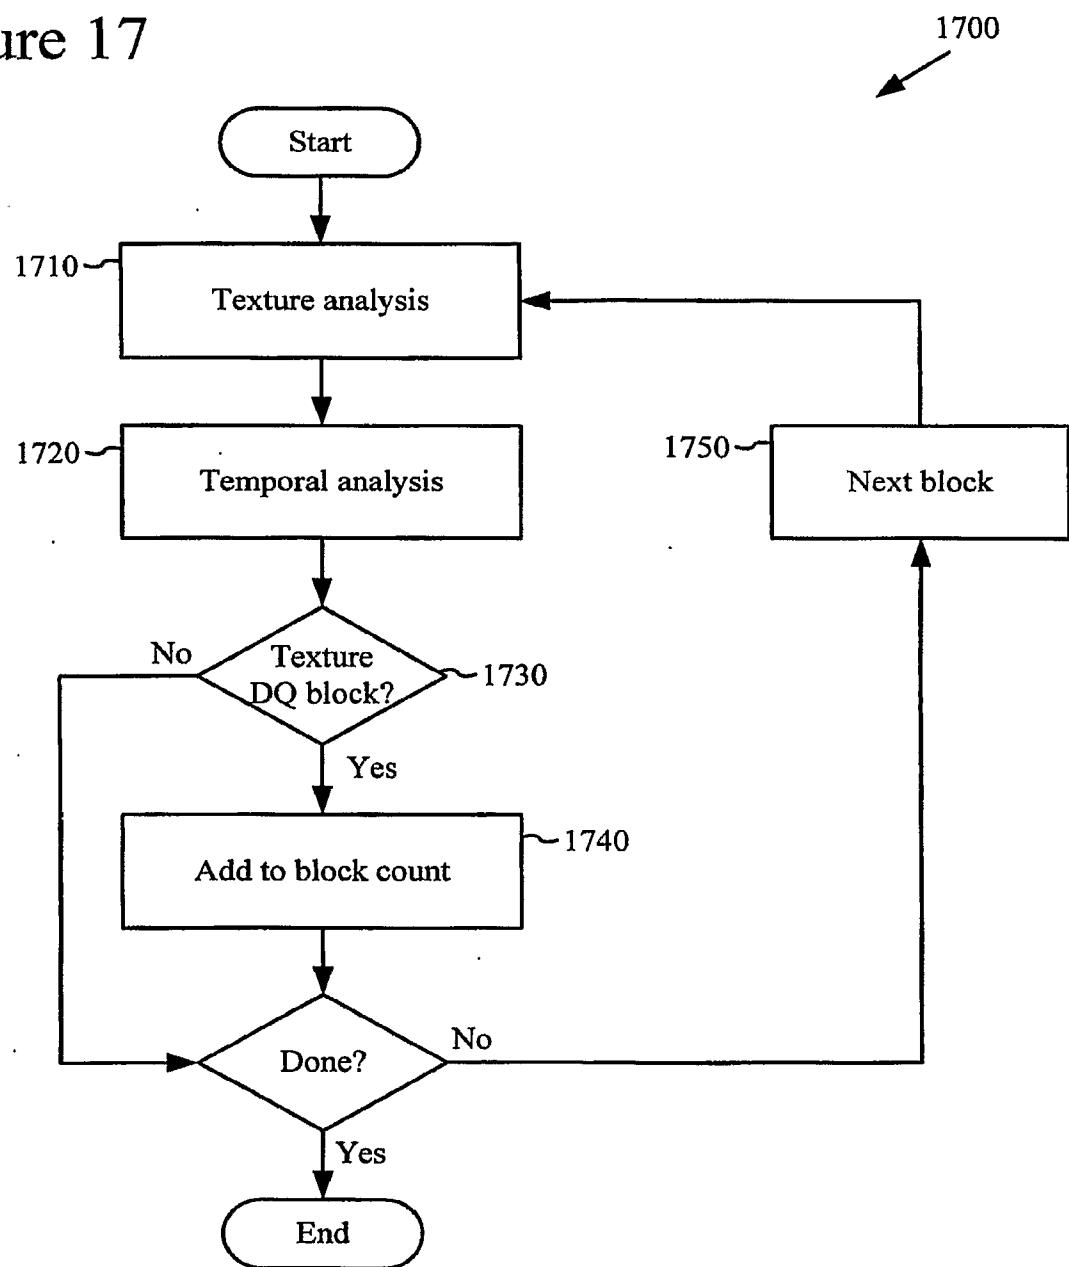


Figure 18

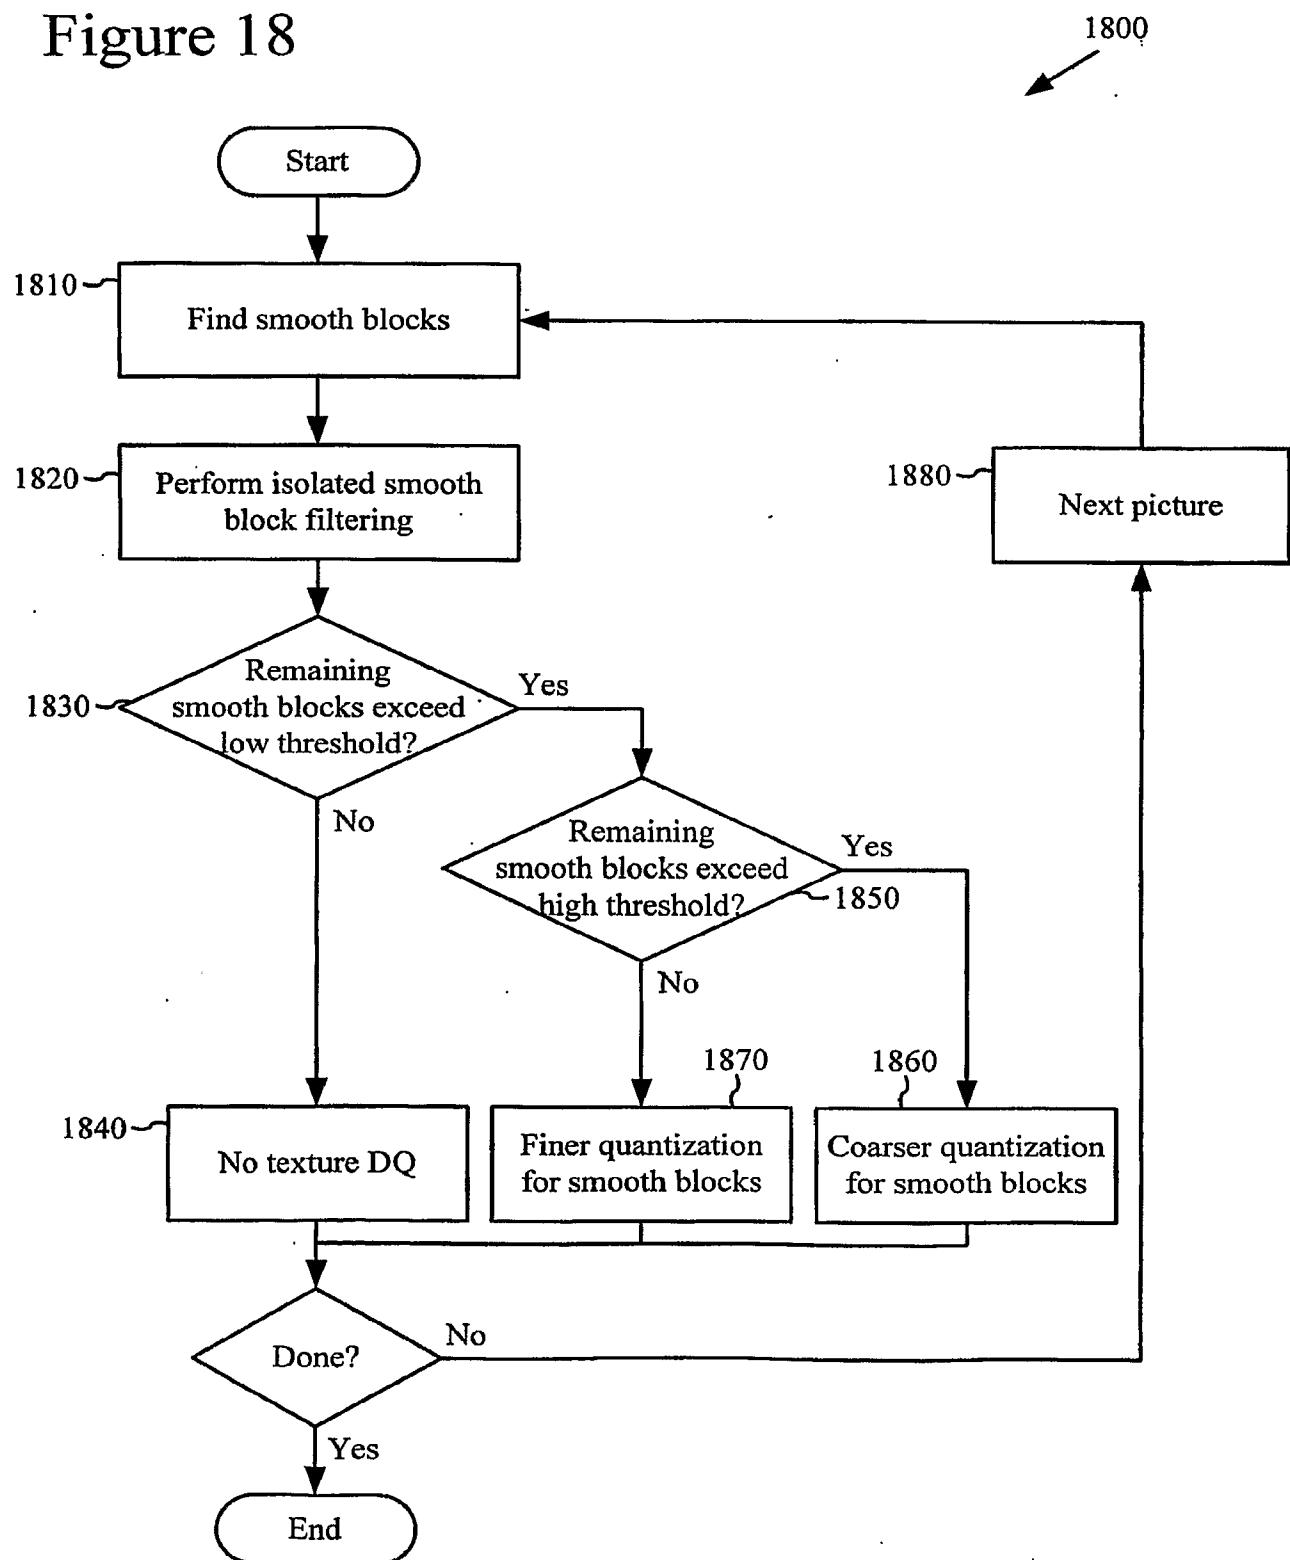
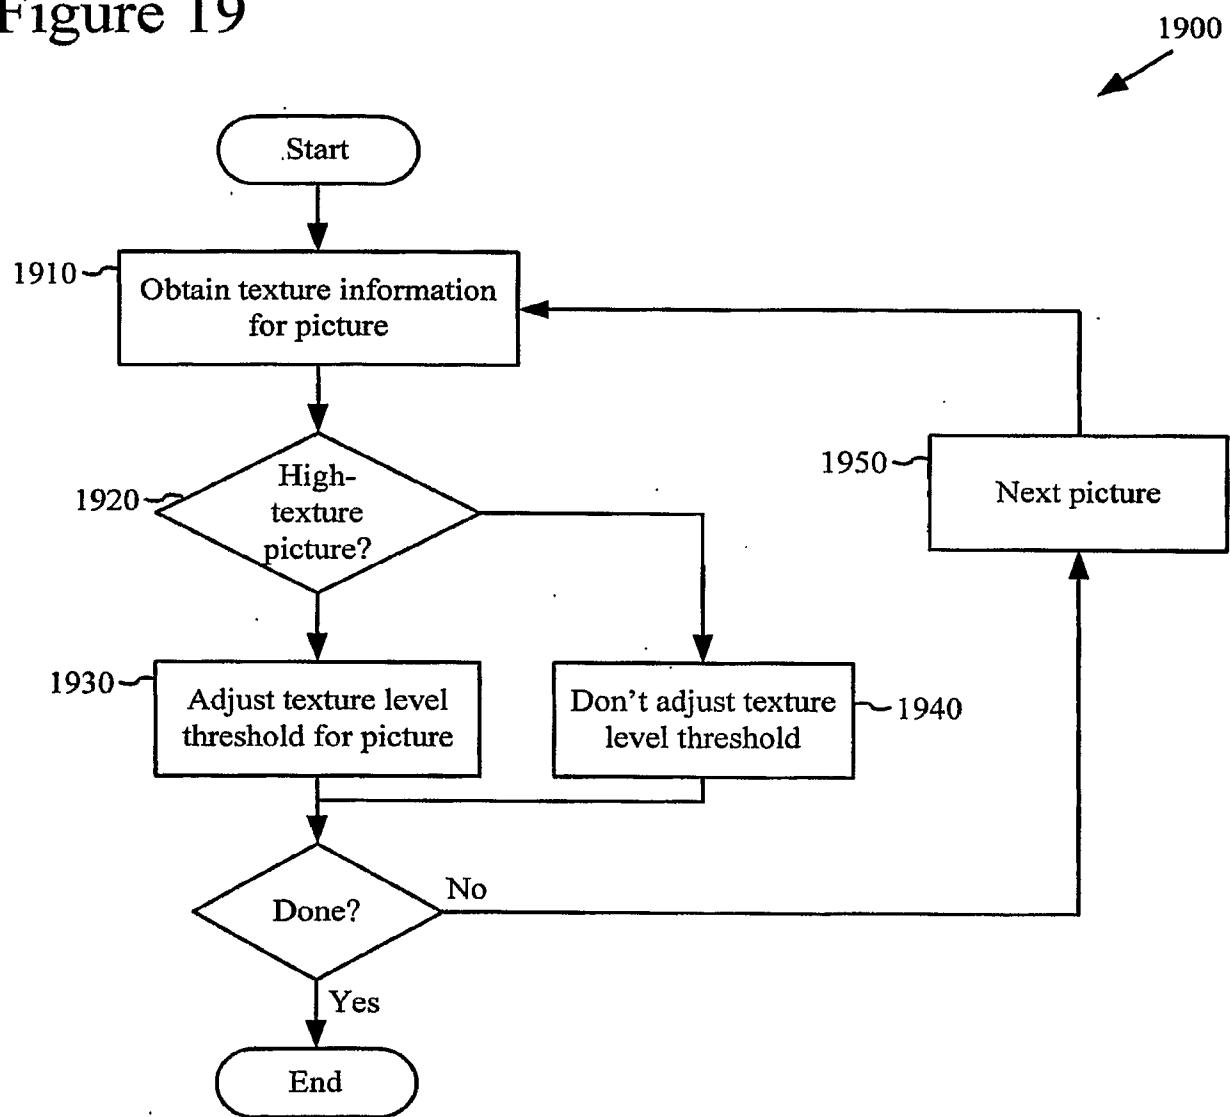


Figure 19



## Figure 20

2000

```
#define TEXTURELOWTH 30
#define TEXTUREMAX      100
#define MINFLATTH       14
#define MAXFLATTH       28

const UInt g_iFlatThTable[TEXTUREMAX-TEXTURELOWTH]={
    16, 16, 16, 16, 18, 18, 18, 18, 18, 18,
    20, 20, 20, 20, 20, 20, 20, 20, 20, 20,
    22, 22, 22, 22, 22, 22, 22, 22, 22, 22,
    24, 24, 24, 24, 24, 24, 24, 24, 24, 24,
    24, 24, 24, 24, 24, 24, 24, 24, 24, 24,
    26, 26, 26, 26, 26, 26, 26, 26, 26, 26,
    26, 26, 26, 26, 26, 26, 26, 26, 26, 26
};

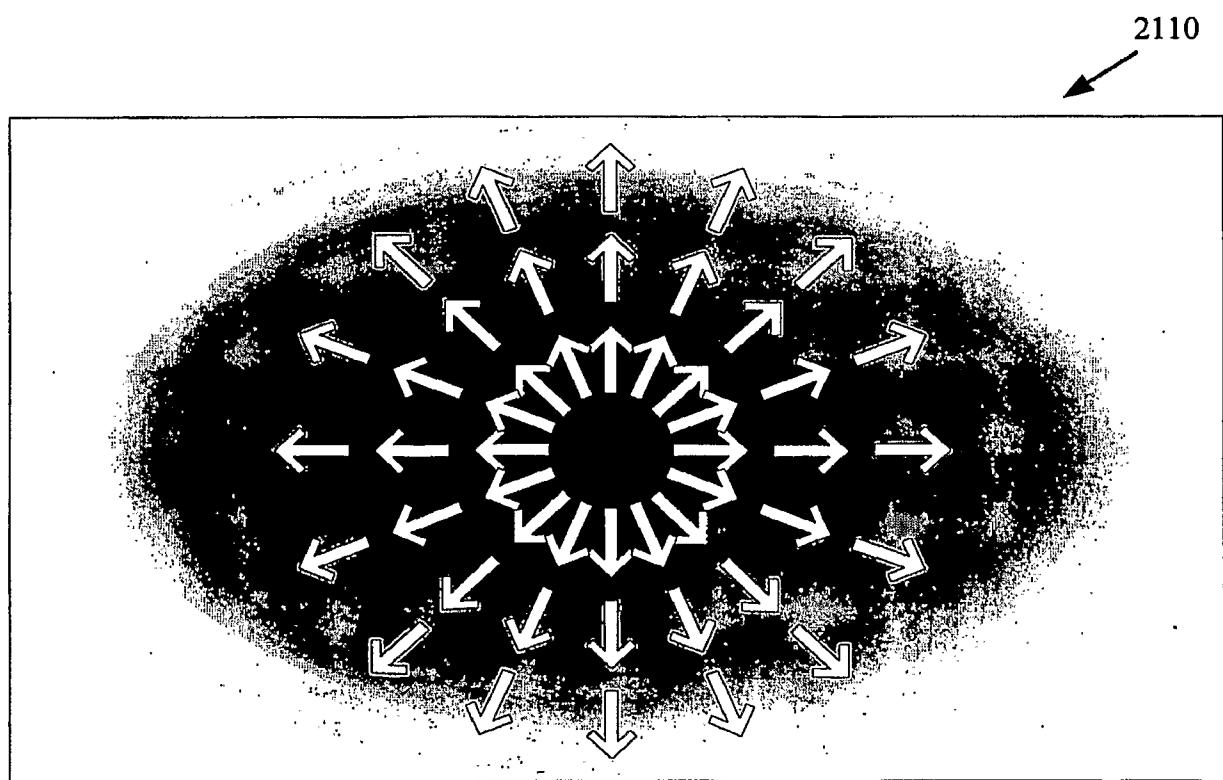
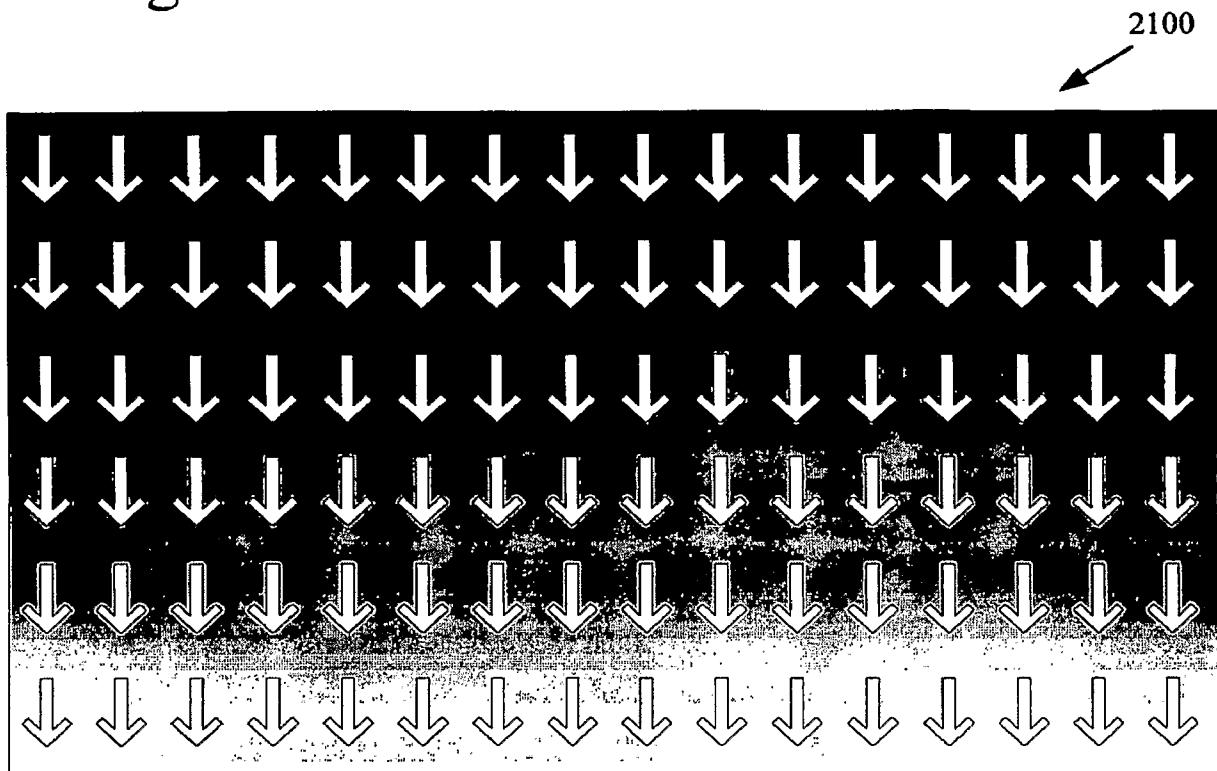
unsigned integer decideAdaptiveFlatThreshold(unsigned integer g0)
{
    Integer new_adaptive_threshold;
    if (g0 < TEXTURELOWTH) new_adaptive_threshold=MINFLATTH;
    else
        if (g0 >= TEXTUREMAX)
            new_adaptive_threshold=MAXFLATTH;
        else
            new_adaptive_threshold=g_iFlatThTable[g0-TEXTURELOWTH];

    if (not scene change key frames)
        new_adaptive_threshold = last_adaptive_threshold +
        max(-4, min(4, new_adaptive_threshold-last_adaptive_threshold));

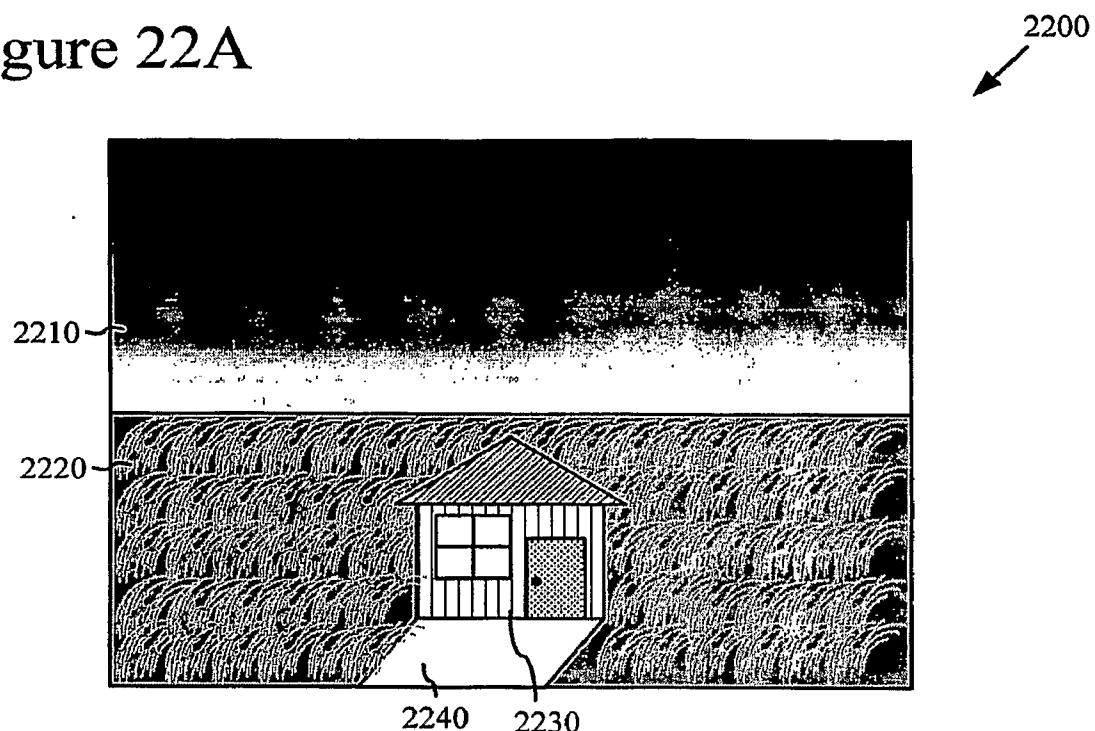
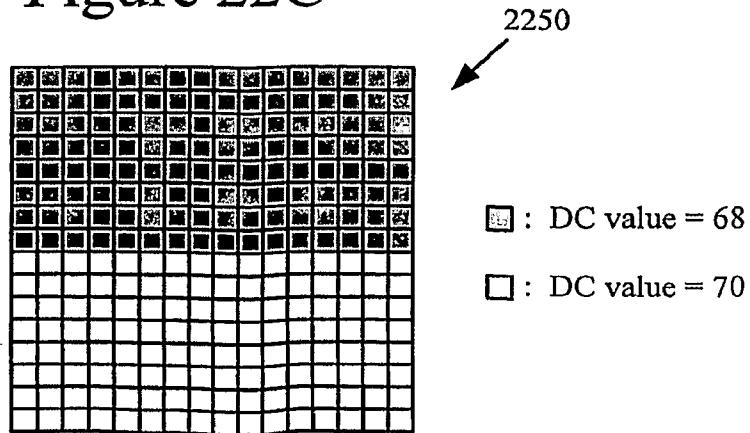
    return new_adaptive_threshold;
}
```

16 / 23

Figure 21



17 / 23

**Figure 22A****Figure 22B****Figure 22C**

18 / 23

Figure 23

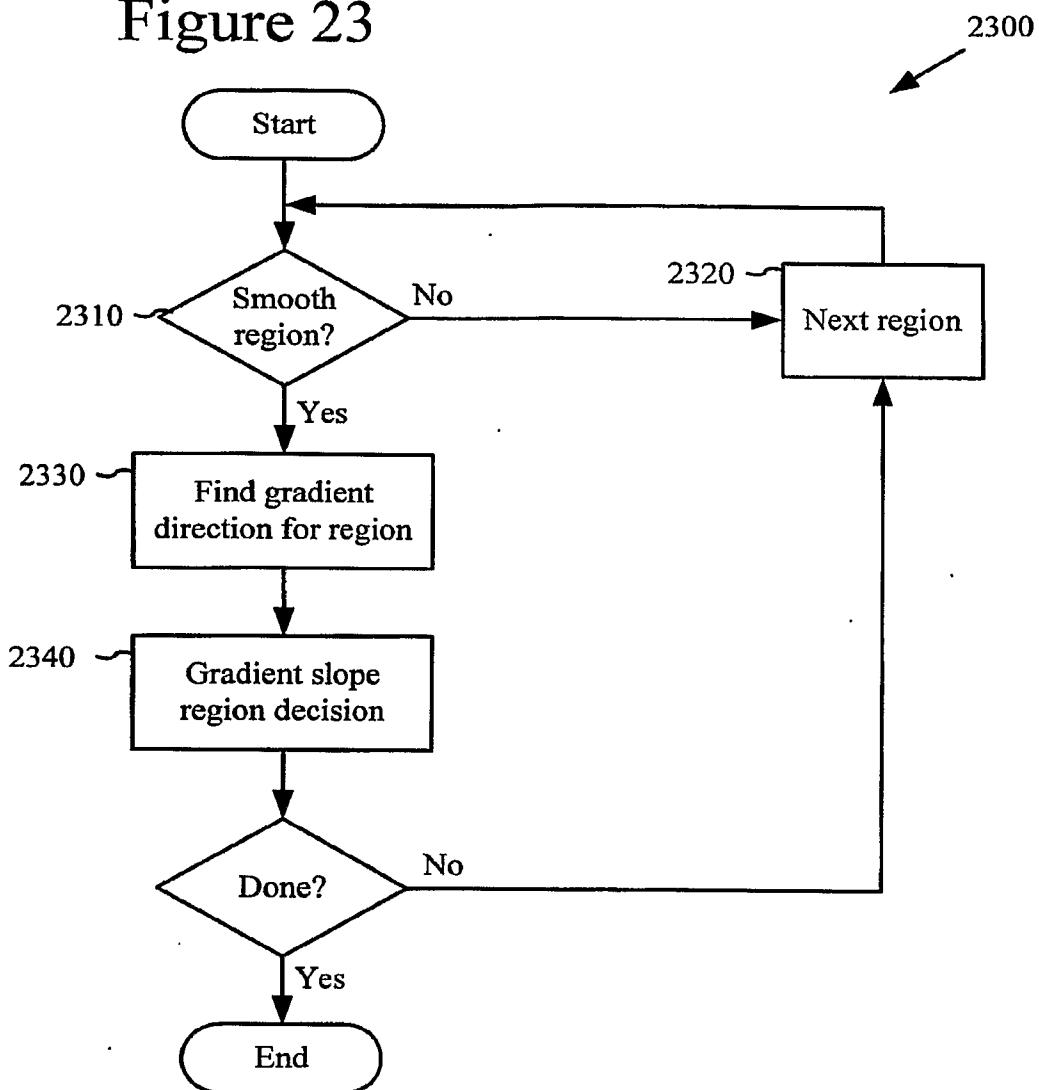
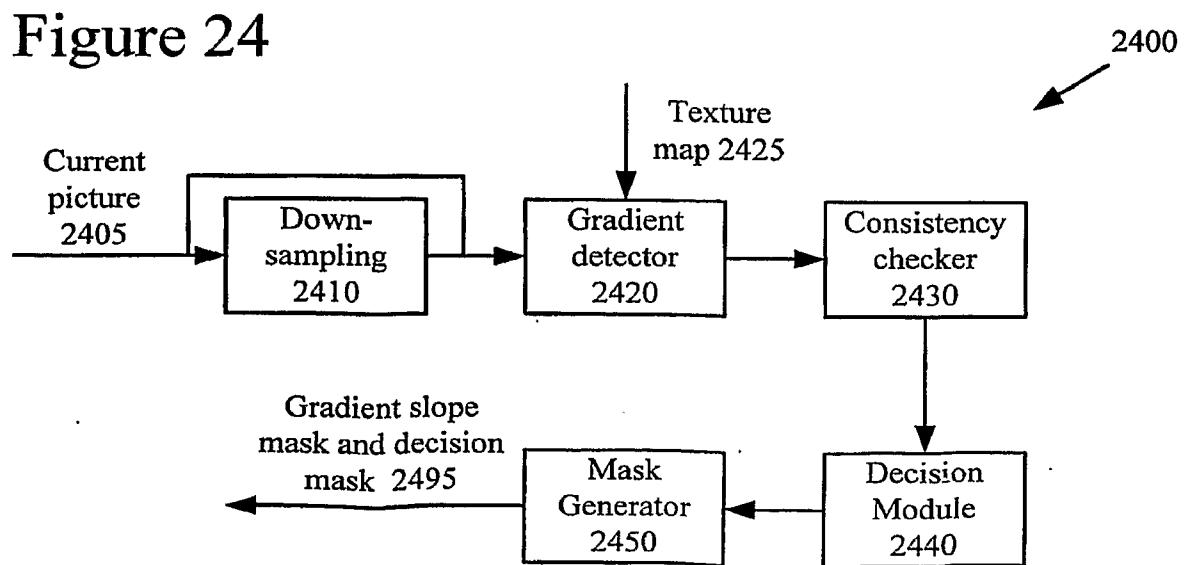


Figure 24



19 / 23

Figure 25

2500

Figure 26

$$K_H = [k_{ij}] \text{ where } 0 \leq i, j \leq 15, k_{ij} = \begin{cases} 1 & \text{if } i < 8 \\ -1 & \text{if } i \geq 8 \end{cases}$$

$$K_V = [k_{ij}] \text{ where } 0 \leq i, j \leq 15, k_{ij} = \begin{cases} 1 & \text{if } j < 8 \\ -1 & \text{if } j \geq 8 \end{cases}$$

Figure 27

2700

```

ComputeGradientDirection (IN region, OUT gradientDirection)
if (region is NOT smooth) {
    gradientDirection = -2; //texture or edge region
    return;
}

gx = sum of all elements of  $K_H \otimes$  block
gy = sum of all elements of  $K_V \otimes$  block
if (abs(gx) < 32 && abs(GradY) < 32 ) {
    gradientDirection = -1; //flat region
    return;
}
θ =  $\tan^{-1}(g_Y / g_X)$ ;
Offset = 0;
if (gx < 0) {
    Offset = 3.14159;
}
Else if (gy < 0) {
    Offset = 2 * 3.14159;
}
θ = θ + Offset;
gradientDirection = ((Int)(θ * 2598)) >> 6;
End_ComputeGradientDirection

```

20 / 23

Figure 28

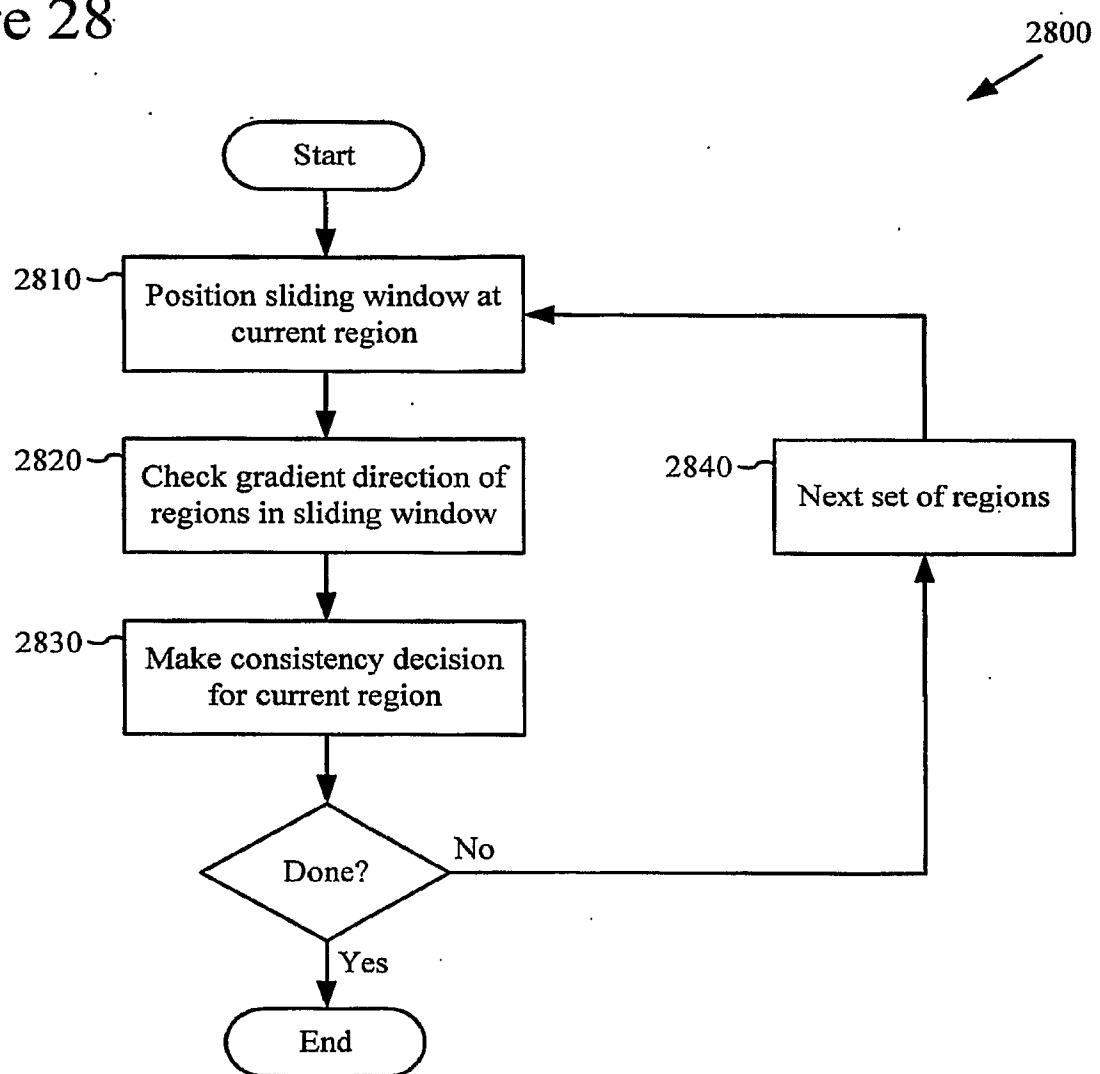


Figure 29

	0	1	2	3	4
0	0	5	10	5	0
1	0	6	24	10	0
2	0	5	10	5	0
3	0	0	0	0	0
4	0	2	0	1	0

21 / 23

Figure 30

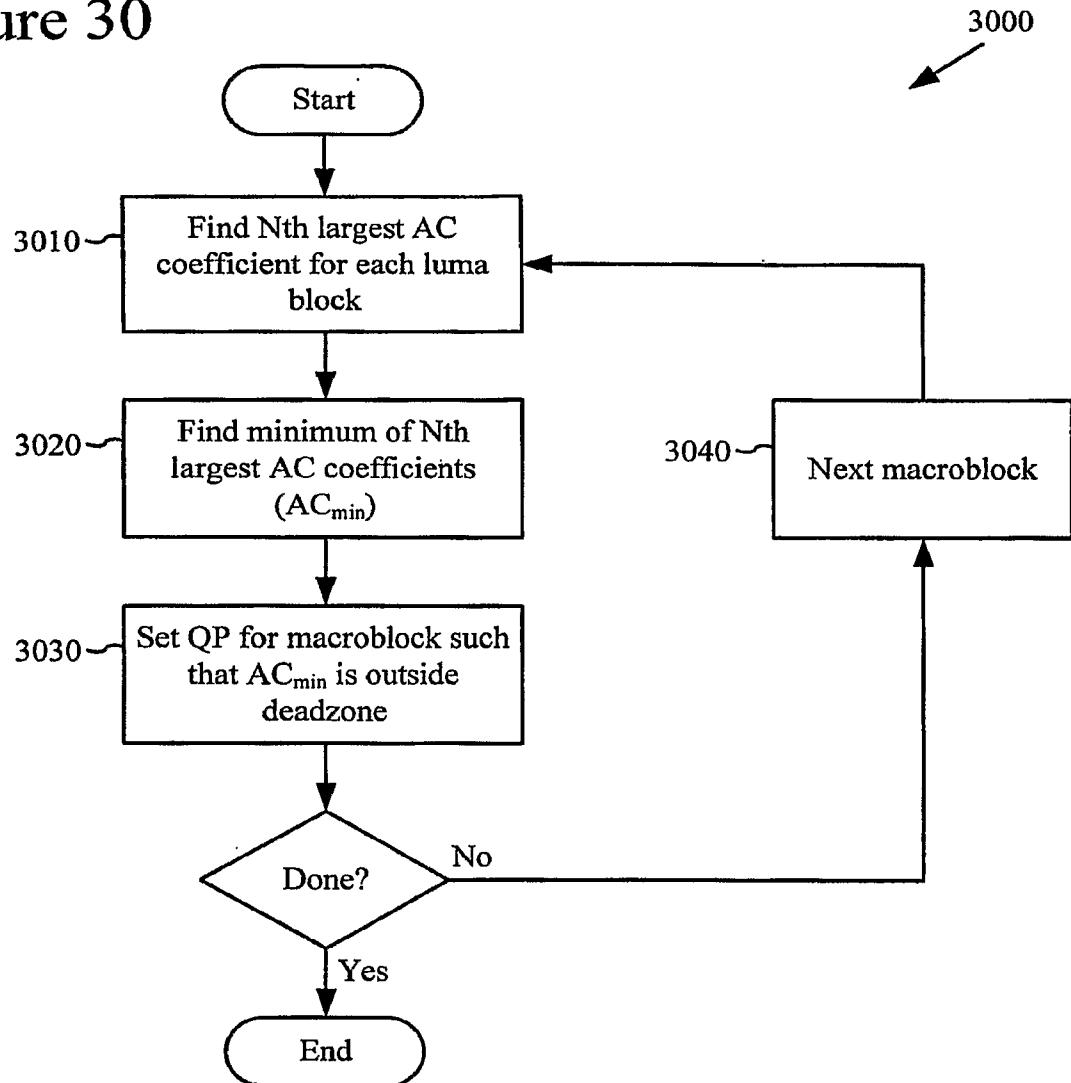


Figure 31

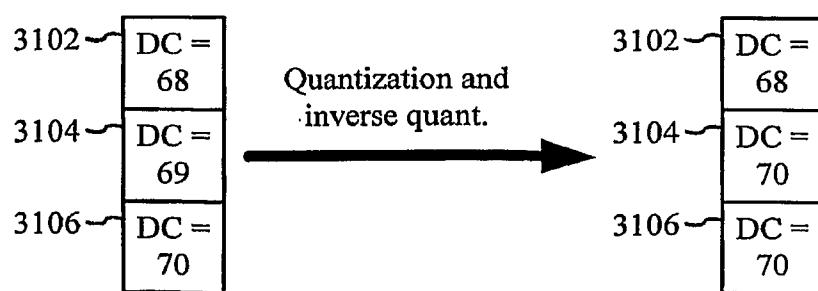


Figure 32

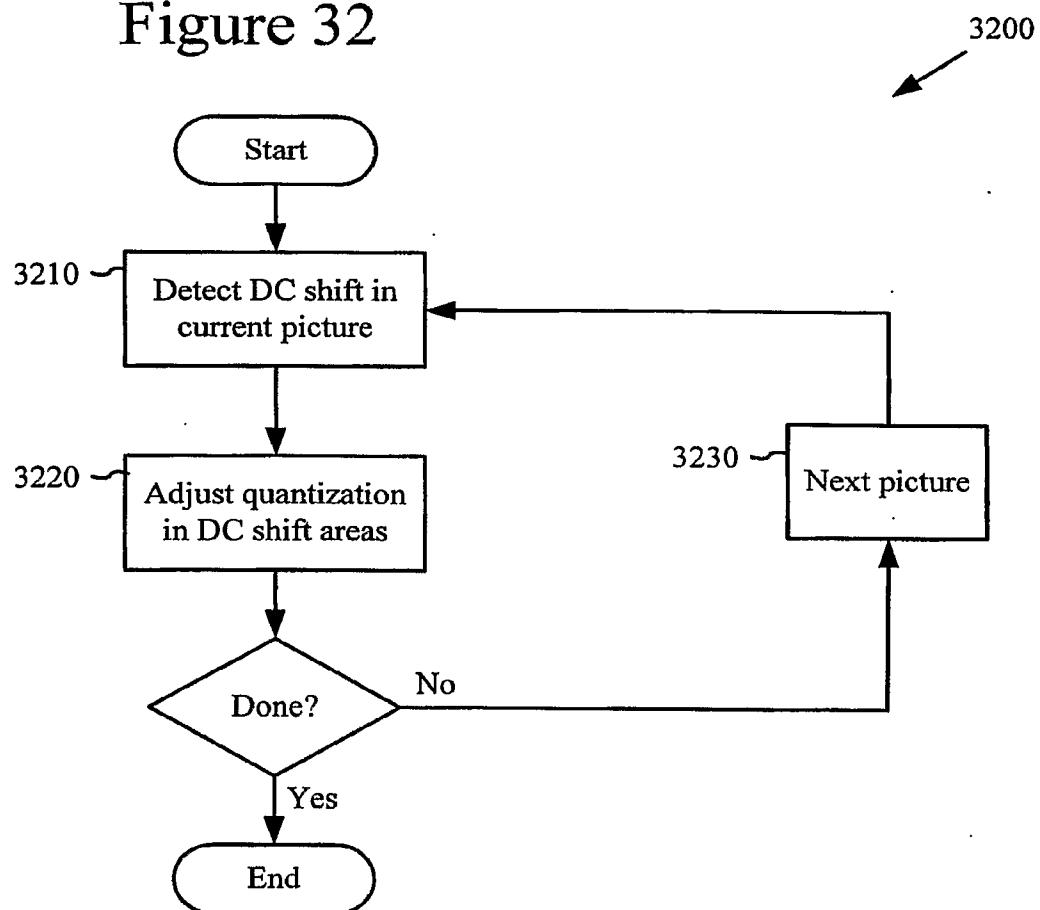
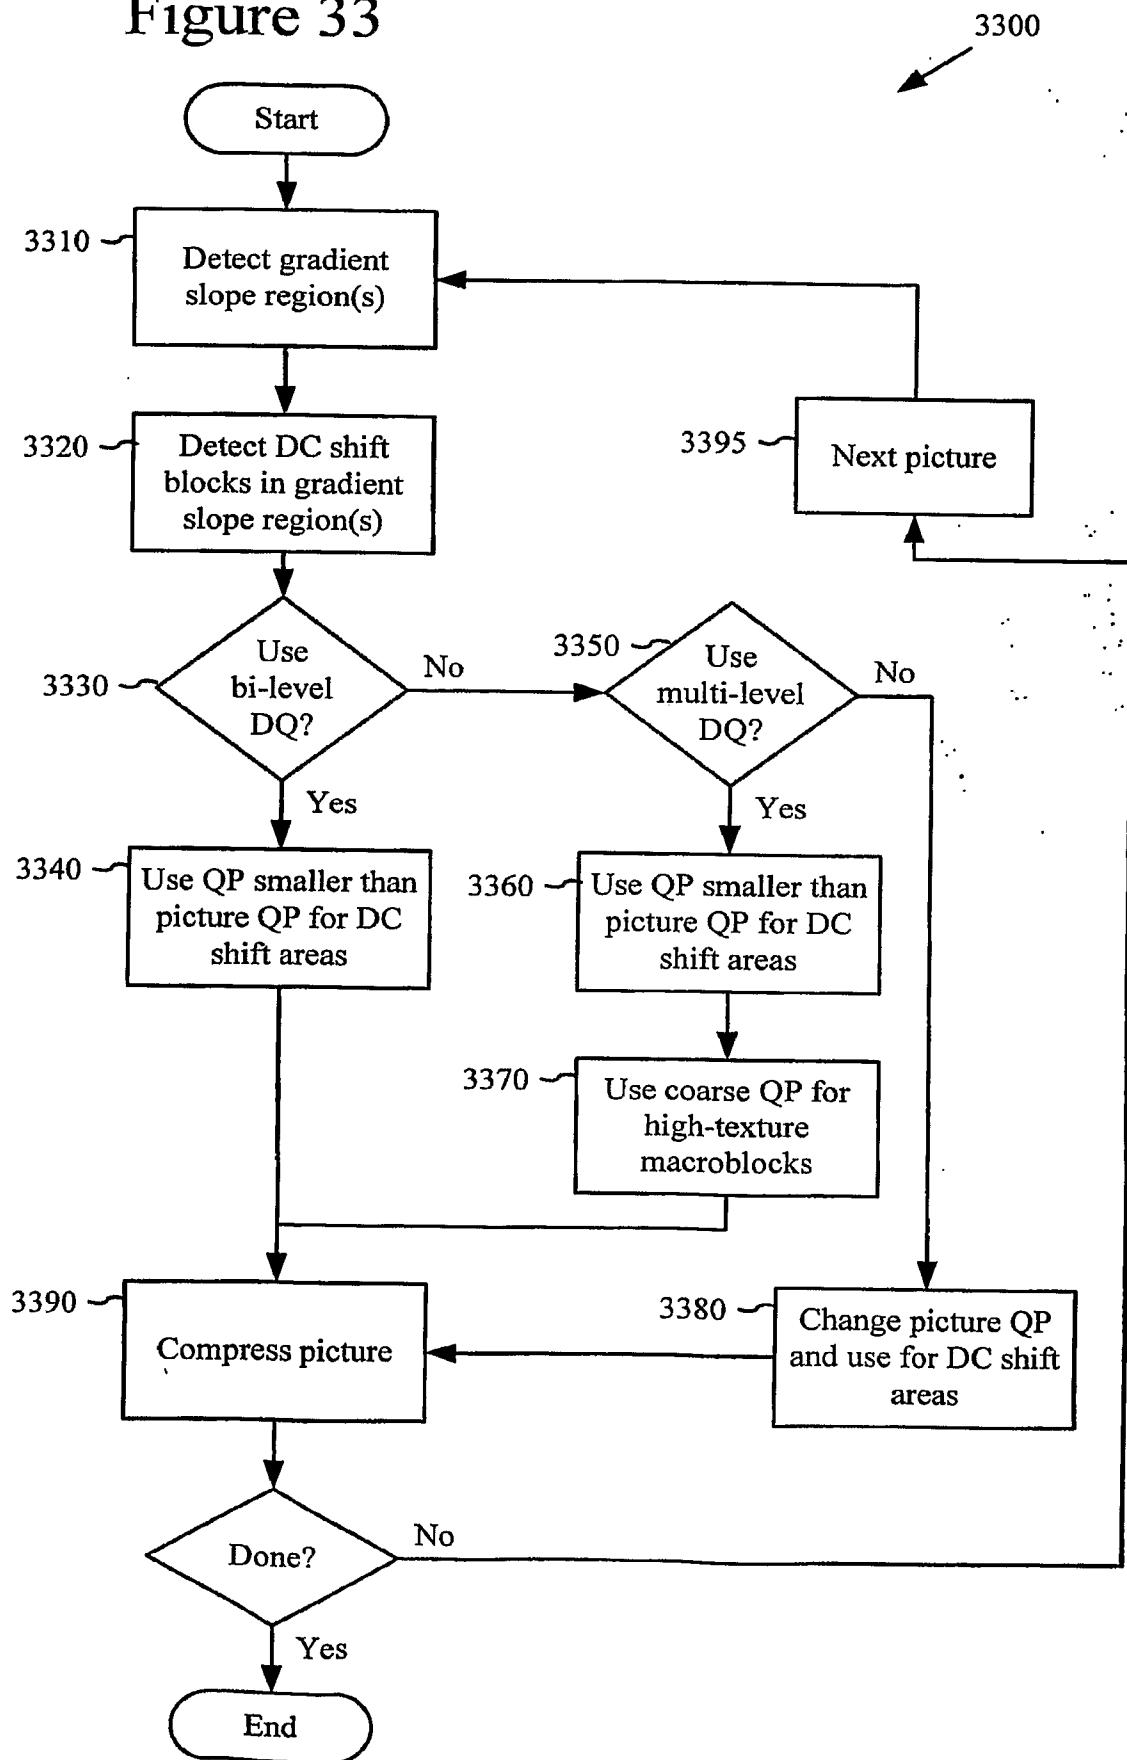


Figure 33



## INTERNATIONAL SEARCH REPORT

International application No.  
**PCT/US2007/008937****A. CLASSIFICATION OF SUBJECT MATTER*****H04N 7/24(2006.01)i***

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)

IPC 8 H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched  
Korean Utility models and applications for Utility Models: IPC as aboveElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
eKIPASS(KIPO internal): "encode, smooth region, adjust"**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 99/48300 A1 (DEAWOO ELECTRONICS CO., LTD.) 23. Sep.1999. See abstract, page 1, line 5 - page 6, line 34.	1-20
A	US 2002/0118884 A1 (HYUN DUK CHO, ALEXANDER L. MAIBORODA, PETERSBERG) 08. Oct. 2002. See abstract, page 1, [0001] - page 5, [0087].	1-20
A	US 6,463,100 B1 (SANGHEE CHO, et al.) 08. Oct. 2002. See abstract, page 1, line 1 - page 5, line 16.	1-20
A	US 2003/0095599 A1 (JIN SOO LEE, et al.) 22. May. 2003. See abstract, page 1, [0001] - page 2, [0023]	1-20

 Further documents are listed in the continuation of Box C. See patent family annex.

\* Special categories of cited documents:  
 "A" document defining the general state of the art which is not considered to be of particular relevance  
 "E" earlier application or patent but published on or after the international filing date  
 "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified)  
 "O" document referring to an oral disclosure, use, exhibition or other means  
 "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention  
 "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone  
 "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art  
 "&" document member of the same patent family

Date of the actual completion of the international search  
13 SEPTEMBER 2007 (13.09.2007)

Date of mailing of the international search report

**13 SEPTEMBER 2007 (13.09.2007)**Name and mailing address of the ISA/KR  
 Korean Intellectual Property Office  
920 Dunsan-dong, Seo-gu, Daejeon 302-701,  
Republic of Korea  
Facsimile No. 82-42-472-7140

Authorized officer

CHO, Woo Yeon

Telephone No. 82-42-481-8524



**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/US2007/008937**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
W09948300A1	23.09.1999	AU749227B2 AU7940098A1 CN1166204C CN1229324A DE69837497CO EP01076999A1 JP11298900A2 US06351492	20.06.2002 11.10.1999 08.09.2004 22.09.1999 16.05.2007 21.02.2001 29.10.1999 26.02.2002
US20020118884A1	29.08.2002	KR1020020047460 US6993197BB	22.06.2002 31.01.2006
US6463100B1	08.10.2002	KR100243430B1 KR1019990060795	01.02.2000 26.07.1999
US20030095599A1	22.05.2003	CN1288910C CN1420690 EP01313322A2 KR1020030040974	06.12.2006 28.05.2003 21.05.2003 23.05.2003