

Method and System for Improving Compressed Image Chroma Information

TECHNICAL FIELD

[0001] This invention relates to video compression, and more particularly to methods, systems, and computer programs for improving compressed image chroma information in MPEG-like video compression systems.

BACKGROUND

MPEG Background

[0002] MPEG-2 and MPEG-4 are international video compression standards defining a video syntax that provides an efficient way to represent image sequences in the form of more compact coded data. The language of the coded bits is the "syntax." For example, a few tokens can represent an entire block of samples (e.g., 64 samples for MPEG-2). Both MPEG standards also describe a decoding (reconstruction) process where the coded bits are mapped from the compact representation into an approximation of the original format of the image sequence. For example, a flag in the coded bitstream signals whether the following bits are to be preceded with a prediction algorithm prior to being decoded with a discrete cosine transform (DCT) algorithm. The algorithms comprising the decoding process are regulated by the semantics defined by these MPEG standards. This syntax can be applied to exploit common video characteristics such as spatial redundancy, temporal redundancy, uniform motion, spatial masking, etc. In effect, these MPEG standards define a programming language as well as a data format. An MPEG decoder must be able to parse and decode an incoming data stream, but so long as the data stream complies with the corresponding MPEG syntax, a wide variety of possible data structures and compression techniques can be used (although technically this deviates from the

standard since the semantics are not conformant). It is also possible to carry the needed semantics within an alternative syntax.

[0003] These MPEG standards use a variety of compression methods, including intraframe and interframe methods. In most video scenes, the background remains relatively stable while action takes place in the foreground. The background may move, but a great deal of the scene is redundant. These MPEG standards start compression by creating a reference frame called an "intra" frame or "I frame". I frames are compressed without reference to other frames and thus contain an entire frame of video information. I frames provide entry points into a data bitstream for random access, but can only be moderately compressed. Typically, the data representing I frames is placed in the bitstream every 12 to 15 frames (although it is also useful in some circumstances to use much wider spacing between I frames). Thereafter, since only a small portion of the frames that fall between the reference I frames are different from the bracketing I frames, only the image differences are captured, compressed, and stored. Two types of frames are used for such differences - predicted or P frames, and bi-directional interpolated or B frames.

[0004] P frames generally are encoded with reference to a past frame (either an I frame or a previous P frame), and, in general, are used as a reference for subsequent P frames. P frames receive a fairly high amount of compression. B frames provide the highest amount of compression but require both a past and a future reference frame in order to be encoded. Bi-directional frames are never used for reference frames in standard compression technologies.

[0005] Macroblocks are regions of image pixels. For MPEG-2, a macroblock is a 16x16 pixel grouping of four 8x8 DCT blocks, together with one motion vector for P frames, and one or two motion vectors for B frames. Macroblocks within P frames may

be individually encoded using either intra-frame or inter-frame (predicted) coding. Macroblocks within B frames may be individually encoded using intra-frame coding, forward predicted coding, backward predicted coding, or both forward and backward (i.e., bi-directionally interpolated) predicted coding. A slightly different but similar structure is used in MPEG-4 video coding.

[0006] After coding, an MPEG data bitstream comprises a sequence of I, P, and B frames. A sequence may consist of almost any pattern of I, P, and B frames (there are a few minor semantic restrictions on their placement). However, it is common in industrial practice to have a fixed pattern (e.g., IBBPBBPBBPBBPBB).

MPEG Color Space Representation

[0007] MPEG-1, MPEG-2, and MPEG-4 all utilize a Y, U, V color space for compression. There is a choice of luminance equation, but a typical conversion transformation between RGB (red-green-blue) to a YUV representation is expressed as:

$$Y = .59 G + .29 R + .12 B$$

$$U = R - Y$$

$$V = B - Y$$

The Y luminance factors for green range from 0.55 up to 0.75, depending upon the color system. The factors for red range from 0.2 to 0.3, and the factors for blue range from 0.05 to 0.15.

[0008] This transformation can be cast as a matrix transformation, which is a linear operator intended for use on linear signals. However, this simple transformation is performed in MPEG 1, 2, and 4 in the non-linear video space, yielding various artifacts and problems.

[0009] It is typical in MPEG to reduce the resolution of the U and V chroma channels to achieve higher compression. The most commonly used reduction of resolution is to use half resolution both vertically and horizontally. MPEG-2 supports full resolution chroma, as well as half resolution horizontally. However, the most commonly used MPEG-2 profiles, Main Profile at Main Level (MP @ ML) and Main Profile at High Level (MP @ HL), use half resolution horizontally and vertically. MPEG-4 versions 1 and 2 use half resolution vertically and horizontally. Note that full chroma resolution is often called 4:4:4, half chroma horizontal resolution is often called 4:2:2, and half vertical and horizontal resolution is often called 4:2:0. (It should be noted that the 4:x:x nomenclature is flawed in its meaning and derivation, but it is common practice to use it to describe the chroma resolution relationship to luminance.)

[0010] The filter which reduces the horizontal and vertical chroma resolution under the various MPEG standards is applied to non-linear video signals as transformed into the U and V color representation. When the inverse transformation is applied to recover RGB, the non-linear signals and the filters interact in such a way as to produce artifacts and problems. These problems can be generalized as "crosstalk" between the Y luminance and the U and V chroma channels, along with spatial aliasing.

[0011] Further information on linear versus non-linear representations and transformations may be found in "The Use of Logarithmic and Density Units for Pixels" by Gary Demos, presented at the October 1990 SMPTE conference, and published in the SMPTE Journal (Oct. 1991, vol. 100, no. 10). See also "An Example Representation for Image Color and Dynamic Range which is Scalable, Interoperable, and Extensible" by Gary Demos, presented at the October 1993 SMPTE conference and published in the proceedings and preprints. These papers

describe the benefits of logarithmic and linear spaces at various stages of the image compression processing pipeline, and are hereby incorporated by reference.

Chroma Sub-Sampling

[0012] The reason for reducing chroma resolution for U and V is that the human visual system is less sensitive to changes in U and V than it is to changes in luminance, Y. Since Y is mostly green, and U and V are mostly red, and blue respectively, this can also be described as a human visual sensitivity being higher for green than for red and blue. However, although U and V are treated the same in MPEG-1, MPEG-2, and MPEG-4, the human visual system is more sensitive to U (with its red component) than to V (with its blue component).

[0013] This difference in chroma sensitivity is embodied in the 1951 NTSC-2 color standard that is used for television. NTSC-2 uses a YIQ color space, where I and Q are similar to U and V (with slightly different weightings). That is, the I channel primarily represents red minus luminance and the Q channel primarily represents blue minus luminance. In NTSC-2, the luminance is given 4.5 MHz of analog bandwidth, and the I chroma channel is given 1.5 MHz of analog bandwidth. The Q channel, representing the blue-yellow axis, is given only 0.5 MHz of analog bandwidth.

[0014] Thus, the NTSC-2 television system allocates three times as much information to the I channel than it does to the Q channel, and three times as much information to the Y luminance channel than to the I channel. Therefore, the bandwidth ratio between the Y luminance channel and the Q (blue minus luminance) channel is nine. These MPEG YUV and NTSC-2 relationships are summarized in the following table:

[0015]

Ratio	YUV 4:4:4	YUV 4:2:2	YUV 4:2:0	NTSC- 2
Red, U, and I pixels to Y	1:1	2:1	4:1	3:1
Blue, V, and Q pixels to Y	1:1	2:1	4:1	9:1

Ratio of Chroma Resolution to Luminance

[0016] Clearly there is a greater difference in treatment between the luminance channel and the U and V channels under the MPEG standards than the luminance and I and Q channels in the NTSC-2 standard.

SUMMARY

[0017] The invention is directed to methods, systems, and computer programs for improving compressed image chroma information.

[0018] More particularly, in one aspect of the invention, a color video image may be improved by increasing the red resolution for an RGB representation (or the U resolution for a YUV representation) above the resolution used for blue (or V). Using lower resolution for the blue color component means less information needs to be compressed, such as in a motion compensated color video image compression system. This aspect of the invention includes a method, system, and computer program for compressing image chroma information of a color video image in a video image compression system by selecting a resolution for a red color component of the color video image that is higher than the resolution for a blue color component of the color video image.

[0019] Another aspect of the invention is a technique for reducing the level of chroma noise that results from any given value of the quantization parameter (QP) used during

compression, thereby improving image quality. This is accomplished by utilizing a lower value of QP for the U ($=R-Y$) channel than for the Y channel. Similarly, the quality of the V ($=B-Y$) channel may also be improved by utilizing a lower QP value for the V channel than for the Y channel.

[0020] Another aspect of the invention is a technique useful when higher compression is required. In this aspect, a positive QP bias is applied to the QP value for the Y channel for use with either or both of the U and V chroma channels.

[0021] Another aspect of the invention is use of a logarithmic representation to benefit image coding. Logarithmic coding, when feasible, can improve coding efficiency for YUV color space representations of images originally represented as linear RGB pixel values. At other processing steps, a conversion to and from linear representations can be beneficial.

[0022] Another aspect of the invention is a method for improving the video characteristics of a color video image in a video compression system, including: selecting a set of image channels to represent the color video image, including a luminance channel and n chroma channels, where n is at least three; and compressing the luminance channel and the n additional chroma channels to a compressed video image.

[0023] The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

[0024] FIG. 1 is a flowchart showing an illustrative method (which may be computer implemented) for increasing the resolution for U above the resolution used for V in a YUV color space representation.

[0025] FIG. 2 is a flowchart showing an illustrative method (which may be computer implemented) for applying a QP bias for chroma channels.

[0026] FIG. 3 is a flowchart showing an illustrative method (which may be computer implemented) for logarithmic coding of luminance and chroma information.

[0027] FIG. 4 is a flowchart showing an illustrative method (which may be computer implemented) for coding additional chroma channels in an image compression system.

[0028] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

Improved Color Coding Precision

[0029] As the quality of images improves with respect to the attributes of reduced noise, extended dynamic range, and extended color range, human sensitivity to color also increases. In particular, it has been observed that red in an RGB representation (or U in a YUV representation) often requires higher precision and clarity than is commonly used in video compression.

[0030] Unless blue is being used for processing (such as blue-screen special effects compositing or image analysis), human sensitivity to the blue-yellow chroma axis, as embodied by either blue or V, is adequately addressed by half resolution sampling horizontally and vertically. Thus, one quarter of the total number of pixels of an image provides sufficient quality for representing the blue or V chroma axis. However, unlike blue and V, one-half resolution coding of red and/or U is sometimes insufficient in quality with respect to large wide-dynamic range displays and projectors.

[0031] Thus, an image may be improved by increasing the red resolution for an RGB representation (or the U resolution for

a YUV representation) above the resolution used for blue (or V). Using lower resolution for the blue color component means less information needs to be compressed, such as in a motion compensated color video image compression system.

[0032] In accordance with the invention, there are three preferred methods of maintaining increased red (or U) resolution with respect to a downfiltered blue (or V) resolution:

- 1) Use full resolution for red and/or U;
- 2) Use one-half resolution on only one chroma axis, either vertically or horizontally, for red and/or U; or
- 3) Use a filtered resolution between full size and one-half, such as $2/3$ or $3/4$, on one or both chroma axes for red and/or U.

[0033] FIG. 1 is a flowchart showing an illustrative method (which may be computer implemented) utilizing higher resolution for U than the resolution used for V in a YUV color space representation (a similar method may be applied to an RGB color space representation):

[0034] Step 101: In an image compression system utilizing a YUV color space representation, downsize filter the V ($=B-Y$) channel of an input image to one-half resolution horizontally, and optionally to one-half resolution vertically.

[0035] Step 102: Downsize filter the U ($=R-Y$) channel of the image to a resolution higher than the V ($=B-Y$) channel, preferably being one of:

- a) full resolution;
- b) between one-half and full resolution horizontally, but full resolution vertically;
- c) between one-half and full resolution horizontally and vertically;
- d) between one-half and full resolution vertically, but full resolution horizontally.

[0036] Step 103: Compress the YUV image (having luminance Y and the downsize filtered U and V chroma information) using an MPEG-like compression system.

[0037] Step 104: Decompress the images into Y, U, and V channels (usually in a different computer).

[0038] Step 105: Convert the U and V channels to full resolution, using the appropriate resolution increase (i.e., the reciprocal of the downsize filter factor used in Step 101 above for V and Step 102 above for U).

[0039] Step 106: Optionally, convert the YUV picture to an RGB image for viewing, analysis, or further processing.

Differential QP Bias for Chroma

[0040] Co-pending U.S. Patent Application No. 09/798,346, entitled "High Precision Encoding and Decoding of Video Images" and assigned to the assignee of the present invention (which is hereby incorporated by reference), teaches various aspects of the use of the quantization parameter (QP) during compression. Another aspect of the present invention is a technique for reducing the level of chroma noise that results from any given value of the quantization parameter (QP) used during compression, thereby improving image quality. This is accomplished by utilizing a lower value of QP for the U (=R-Y) channel than for the Y channel. Similarly, the quality of V (=B-Y) may also be improved by utilizing a lower QP value for the V channel than for the Y channel.

[0041] A simple method of implementing a reduced chroma QP value is to subtract a constant value from the QP value used for the Y (luminance) channel. Alternatively, a separate constant value (lower than the QP value for Y) might be used for each of U and V. For example, "2" might be subtracted from the QP value for Y to yield the QP value for U, and "1" might be subtracted for the QP value for Y to yield the QP value for V. Any useful value of the amount to subtract can be used,

limited only by a minimum value of "1" for the applied QP value.

[0042] This method works for constant QP values (variable bit rate). It also works as well for variable QP values (e.g., in both constant and variable bit rate motion compensated compression systems), since the instantaneous QP value can be biased by subtracting a specified difference value from the QP value for Y to yield a QP value for each of U and V.

[0043] Further, the range of these differential chroma-biased QP values can be extended using the extended QP range function or lookup, as described in the "High Precision Encoding and Decoding of Video Images" Patent Application referenced above.

[0044] It is necessary to signal the U and V bias values from the encoder to the decoder unless a pre-arranged value is used. These can be specified once, for example, for each session, group of pictures (GOP), frame, or image region.

[0045] FIG. 2 is a flowchart showing an illustrative method (which may be computer implemented) for applying a QP bias for chroma channels:

[0046] Step 201: In an image compression system, reduce the QP value for each of the U and V chroma channels by a selected value (which may be different for each channel).

[0047] Step 202: Utilize this reduced QP value for the U and V chroma channel compressions, respectively.

[0048] Step 203: Optionally, if variable QP values are used, ensure that the reduced U and V QP value is at least "1".

[0049] Step 204: Unless a pre-set bias is to be used, signal or convey the QP value reduction amount to the decoder as often as it may change (once at a minimum).

[0050] Step 205: Decompress (usually in a different computer) the signal using the appropriate QP value for U and V (again ensuring that the reduced QP value is at least "1").

[0051] Step 206: Optionally, view the decompressed images, or use the images for additional processing or analysis.

[0052] Another aspect of the invention is a technique useful when higher compression is required. In this aspect, a positive QP bias is applied to the QP value for the Y channel for use with either or both of the U and V chroma channels (preferably checking against a QP maximum value of a compression system, if any). Separate bias can be used for each of the U and V channels. Otherwise, the steps of such an embodiment would be similar to those shown in FIG. 2.

Logarithmic Coding of Luminance and Chroma

[0053] The paper entitled "The Use of Logarithmic and Density Units for Pixels," referenced above, describes the benefits of a logarithmic representation for dynamic range. Log representations of a matching dynamic range are somewhat similar to commonly used video transfer functions. Even though similar, the logarithmic representation is more optimal in extensibility, calibration usage, and in orthogonality of color channels than are the various commonly used video representations.

[0054] Another aspect of the invention is use of a logarithmic representation to benefit image coding. It has been discovered that logarithmic coding, when feasible, can improve coding efficiency for YUV color space representations of images originally represented as linear RGB pixel values (such as at the sensor of a camera). At other processing steps, a conversion to and from linear representations can be beneficial.

[0055] As described in the "High Precision Encoding and Decoding of Video Images" Patent Application referenced above, chroma crosstalk with luminance is minimized when:

$$Ylog = \text{Log} (W_r * R + W_g * G + W_b * B)$$

$$U = \text{Log}(R) - Ylog$$

$$V = \text{Log}(B) - Ylog$$

where W_r , W_g , and W_b are the linear weightings for the red, green, and blue components of luminance, and where R , G , and B represent a linear light space. These relationships are useful in applying this aspect of the invention.

[0056] FIG. 3 is a flowchart showing an illustrative method (which may be computer implemented) for logarithmic coding of luminance and chroma information:

[0057] Step 301: In an image compression system, perform the following transformation on input (e.g., directly from a video camera) linear R , G , and B pixel values:

$$Ylog = \text{Log} (W_r * R + W_g * G + W_b * B)$$

$$U = \text{Log}(R) - Ylog$$

$$V = \text{Log}(B) - Ylog$$

where W_r , W_g , and W_b are the linear weightings for the red, green, and blue components of luminance.

[0058] Step 302: Optionally, reduce the resolution of the U and V chroma channels (as described above).

[0059] Step 303: Perform motion-compensated compression on this Y , U , and V representation of the moving image.

[0060] Step 304: Decompress the compressed images to restore Y , U , and V color components of the moving image (usually in a different computer).

[0061] Step 305: If optional Step 302 was applied, reverse the resolution reduction to restore full U and V resolution.

[0062] Step 306: Restore the linear R , G , and B pixel values using the following transformation:

$$R = \text{anti-log}(Y + U)$$

$$B = \text{anti-log}(Y + V)$$

$$G = (\text{anti-log}(Y) - W_r * R - W_b * B) / W_g$$

[0063] Step 307: Optionally, convert to other video RGB representations (alternatively, may be done in lieu of Step 306 rather than in addition to Step 306).

Additional Chroma Axes

[0064] In extended dynamic range and extended contrast range images, it may be beneficial to augment visible wavelength channels with additional channels of image information, both visible and non-visible.

[0065] The range of colors available from any given set of red, green, and blue primaries does not include all possible visible colors. The combining of proportions of red, green, and blue primary colors to create other visible colors such as yellow, orange, cyan, and brown, is a property of the human visual system known as the "metamerism".

[0066] As pointed out in the paper entitled "An Example Representation for Image Color and Dynamic Range which is Scalable, Interoperable, and Extensible", referenced above, it is possible to add additional color primaries to the three primaries of red, green, and blue. In particular, cyan, magenta, and yellow color primaries help to extend the color gamut beyond the range available from most common red, green, and blue primary values. Further, violet and ultraviolet (which brightens phosphorescent colors) can also be conveyed.

[0067] Beyond the visible colors, invisible infrared wavelengths have proven useful in penetrating clouds and haze, and in seeing in the dark. Ultraviolet wavelengths can also be useful for seeing low-amplitude visible image details, such as fingerprints and surface coatings.

[0068] Further, even in the visible wavelengths, various materials (e.g., smog and underwater algae) often reduce the amount of contrast or dynamic range of some wavelengths. This is why smog can appear brown, giving a brown tint to all objects in the distance, having reduced the blue contrast and

dynamic range. This is also why underwater photography can appear green, blue-green, or blue, since the red end of the visible spectrum is reduced in contrast and dynamic range.

[0069] The logarithmic relationships between Y, U, and V, as described above, will optimize the coding of color relationships for visible light.

[0070] In this aspect of the invention, additional chroma channels are added to the channels encoding three primary wavelengths, typically embodied by RGB or YUV representations. Further, when using a YUV color space, it is also possible to change the makeup of the Y (luminance) channel to favor the highest amplitude image signals. Thus, for example, the green visible channel might be coded using its own chroma channel, with luminance moving to other wavelength regions. This concept can be extended to where Y luminance is infrared, with red, green, and blue (and perhaps other visible and non-visible primaries) each having their own chroma channels.

[0071] In accordance with this aspect of the invention, for each new chroma channel, the following should be determined:

- 1) Should the channel be coded differentially from one or more other channels (usually from luminance, such as $U=R-Y$)?

- 2) Should the channel be given full resolution with respect to luminance, or can resolution be reduced without impairing the image quality for a given intended usage?

[0072] The determination in 1) is based upon the correlation of each coded channel with other channels. For example, ultraviolet or far-infrared wavelength images may be relatively uncorrelated to visible wavelengths, or to each other. In such a case, these channels might be coded without reference to other channels. However, any visible wavelengths are highly correlated, and thus can almost always benefit from being coded with respect to each other.

[0073] Based upon these determinations, a set of image channels can be selected, usually exceeding (or replacing and

exceeding) the three primary channels (e.g., YUV). For example, the set of selected image channels may comprise a Y' luminance channel, and n chroma channels, such as a U' first chroma channel, a V' second chroma channel, and an X' third chroma channel.

[0074] Using this example, and applying motion compensated compression, the selected value of Y' would be coded with full resolution, and the various other chroma channels (U', V', X') would be differentially or independently coded. All channels can utilize the same motion vector and macroblock motion compensation structure as would be used for conventional YUV representations, except that there would be additional channels. Each such channel would utilize an appropriate resolution with respect to Y (as determined in step 2 above). In addition, a QP bias (as described above) can be independently applied to each chroma channel, to ensure that the desired compression chroma quality is achieved.

[0075] Even when applied only to visible wavelengths, additional chroma channels can ensure not only extended color range and more accurate color, but also allow additional clarity, detail, and noise fidelity to be applied to such highly visible colors as magenta, orange, yellow, and aqua-cyan. These benefits can be particularly significant for wide-dynamic range and wide-contrast range images.

[0076] FIG. 4 is a flowchart showing an illustrative method (which may be computer implemented) for coding additional chroma channels in an image compression system:

[0077] Step 401: In an image compression system, determine an optimal luminance representation for an image, selected based upon widest dynamic range and highest resolution, including optional non-visible wavelength image signals.

[0078] Step 402: Determine n additional chroma channels to represent the image, where n is at least three.

[0079] Step 403: Optionally, for each chroma channel, determine whether it is beneficial to code differentially with respect to luminance and/or one or more other chroma channels.

[0080] Step 404: Determine the resolution desired for each chroma channel image signal from an input with respect to the luminance image signal, such resolution being equal to or less than the resolution of the luminance, and optionally apply a resolution reduction.

[0081] Step 405: Compress the $Y + n$ chroma image signals using motion compensated compression.

[0082] Step 406: Decompress the $Y + n$ chroma images (usually in a different computer).

[0083] Step 407: If resolution reduction was applied, restore the original resolutions of the chroma channels.

[0084] Step 408: Combine each chroma channel with its differential counterpart, if any, from Step 403 above.

[0085] Step 409: Optionally, perform any of the following:

a) Convert the chroma channels to a viewing space, such as RGB, or to spaces having more than three primaries, and view as a true-color image;

b) Perform the conversion of a) but view as a false-color image (such as mapping infrared to green);

c) Use the chroma channels without conversion for processing and/or analysis.

[0086] As another option, each chroma channel may have a biased QP value applied (either increasing or decreasing), relative to the QP value used for the luminance channel, to achieve a desired level of quality for each chroma channel (i.e., trading off chroma noise versus higher degree of compression).

Implementation

[0087] The invention may be implemented in hardware or software, or a combination of both (e.g., programmable logic

arrays). Unless otherwise specified, the algorithms included as part of the invention are not inherently related to any particular computer or other apparatus. In particular, various general purpose machines may be used with programs written in accordance with the teachings herein, or it may be more convenient to construct more specialized apparatus (e.g., integrated circuits) to perform particular functions. Thus, the invention may be implemented in one or more computer programs executing on one or more programmable computer systems each comprising at least one processor, at least one data storage system (including volatile and non-volatile memory and/or storage elements), at least one input device or port, and at least one output device or port. Program code is applied to input data to perform the functions described herein and generate output information. The output information is applied to one or more output devices, in known fashion.

[0088] Each such program may be implemented in any desired computer language (including machine, assembly, or high level procedural, logical, or object oriented programming languages) to communicate with a computer system. In any case, the language may be a compiled or interpreted language.

[0089] Each such computer program is preferably stored on or downloaded to a storage media or device (e.g., solid state memory or media, or magnetic or optical media) readable by a general or special purpose programmable computer, for configuring and operating the computer when the storage media or device is read by the computer system to perform the procedures described herein. The inventive system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer system to operate in a specific and predefined manner to perform the functions described herein.

[0090] A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, some of the steps described above may be order independent, and thus can be performed in an order different from that described. Accordingly, other embodiments are within the scope of the following claims.

WHAT IS CLAIMED IS:

1. A method for compressing image chroma information of a color video image in a video image compression system, including selecting a resolution for a red color component of the color video image that is higher than the resolution for a blue color component of the color video image.
2. A method for compressing image chroma information of a color video image in a video image compression system, including:
 - (a) downfiltering a blue color component of the color video image to a processed blue color component having a first resolution along at least one of the horizontal and vertical image axes of the color video image; and
 - (b) filtering a red color component of the color video image to a processed red color component having a second resolution higher than the first resolution.
3. The method of claim 2, wherein the second resolution is in the range from 0.5 to 1.0 of the full resolution of the red color component along at least one of the horizontal and vertical image axes of the color video image.
4. The method of claim 2, further including compressing at least the processed blue color and red color components to a compressed output image.
5. The method of claim 4, further including decompressing the compressed output image to obtain uncompressed processed blue and red color components.
6. The method of claim 5, further including upsize filtering the processed blue and red color components to the full resolution of the color video image.

7. The method of claims 1 or 2, wherein the video image compression system is a motion-compensated video image compression system.
8. A method for reducing chroma noise during compression of a color video image in a YUV video image compression system utilizing a quantization parameter (QP) during compression, including utilizing a first QP value for the Y color channel of a color video image, and a second QP value for at least one of the U and V color channels of the color video image, wherein the second QP value is less than the first QP value.
9. The method of claim 8, wherein the second QP value is determined by applying a bias value to the first QP value.
10. The method of claim 8, further including compressing the color video image, after application of the first and second QP values, to a compressed output image.
11. The method of claim 10, further including decompressing the compressed output image using the first and second QP values to obtain an uncompressed video image.
12. A method for achieving higher compression during compression of a color video image in a YUV video image compression system utilizing a quantization parameter (QP) during compression, including utilizing a first QP value for the Y color channel of a color video image, and a second QP value for at least one of the U and V color channels of the color video image, wherein the second QP value is greater than the first QP value.
13. The method of claim 12, wherein the second QP value is determined by applying a bias value to the first QP value.

14. The method of claim 12, further including compressing the color video image, after application of the first and second QP values, to a compressed output image.
15. The method of claim 14, further including decompressing the compressed output image using the first and second QP values to obtain an uncompressed video image.
16. A method for improving the coding efficiency for a color space representation of a video image originally represented as linear RGB pixel values in a video image compression system, including transforming the linear RGB pixel values of the video image to a logarithmic representation of luminance and chroma channel information.
17. The method of claim 16, wherein transforming includes applying the following equations to obtain a YUV logarithmic representation of the video image:
- $$Ylog = \text{Log} (W_r * R + W_g * G + W_b * B)$$
- $$U \text{ chroma channel} = \text{Log}(R) - Ylog$$
- $$V \text{ chroma channel} = \text{Log}(B) - Ylog$$
- where W_r , W_g , and W_b are linear weightings for red, green, and blue components of luminance of the video image.
18. The method of claim 17, further including reducing the resolution of the U and V chroma channels of the YUV logarithmic representation.
19. The method of claim 17, further including compressing the YUV logarithmic representation of the video image to a compressed video image.
20. The method of claim 19, further including decompressing the compressed video image to a restored YUV logarithmic representation of the video image.

21. The method of claim 20, further including transforming restored YUV logarithmic representation of the video image to linear RGB pixel values.
22. The method of claim 21, wherein transforming includes applying the following equations to obtain the linear RGB pixel values:
- $$R = \text{anti-log}(Y + U)$$
- $$B = \text{anti-log}(Y + V)$$
- $$G = (\text{anti-log}(Y) - W_r * R - W_b * B) / W_g.$$
23. A method for improving the video characteristics of a color video image in a video compression system, including:
- (a) selecting a set of image channels to represent the color video image, including a luminance channel and n chroma channels, where n is at least three; and
 - (b) compressing the luminance channel and the n additional chroma channels to a compressed video image.
24. The method of claim 23, wherein at least one chroma channel represents non-visible wavelengths.
25. The method of claim 23, wherein the luminance channel is the image channel having the highest dynamic range and resolution.
26. The method of claim 23, further including coding each chroma channel independently from each other channel.
27. The method of claim 23, further including coding each chroma channel differentially with respect to a selected other channel.
28. The method of claim 23, further including reducing the resolution of at least one chroma channel.

29. The method of claim 23, further including applying a quantization parameter (QP) value to at least one chroma channel biased with respect to a QP value applied to the luminance channel.
30. A computer program, stored on a computer-readable medium, for compressing image chroma information of a color video image in a video image compression system, the computer program comprising instructions for causing a computer to permit selection of a resolution for a red color component of the color video image that is higher than the resolution for a blue color component of the color video image.
31. A computer program, stored on a computer-readable medium, for compressing image chroma information of a color video image in a video image compression system, the computer program comprising instructions for causing a computer to:
- (a) downfilter a blue color component of the color video image to a processed blue color component having a first resolution along at least one of the horizontal and vertical image axes of the color video image; and
 - (b) filter a red color component of the color video image to a processed red color component having a second resolution higher than the first resolution.
32. The computer program of claim 31, wherein the second resolution is in the range from 0.5 to 1.0 of the full resolution of the red color component along at least one of the horizontal and vertical image axes of the color video image.
33. The computer program of claim 31, further including instructions for causing the computer to compress at least the processed blue color and red color components to a compressed output image.

34. The computer program of claim 33, further including instructions for causing a decompression computer to decompress the compressed output image to obtain uncompressed processed blue and red color components.
35. The computer program of claim 34, further including instructions for causing the decompression computer to upsize filter the processed blue and red color components to the full resolution of the color video image.
36. The computer program of claims 30 or 31, wherein the video image compression system is a motion-compensated video image compression system.
37. A computer program, stored on a computer-readable medium, for reducing chroma noise during compression of a color video image in a YUV video image compression system utilizing a quantization parameter (QP) during compression, the computer program comprising instructions for causing a computer to utilize a first QP value for the Y color channel of a color video image, and a second QP value for at least one of the U and V color channels of the color video image, wherein the second QP value is less than the first QP value.
38. The computer program of claim 37, wherein the second QP value is determined by applying a bias value to the first QP value.
39. The computer program of claim 37, further including instructions for causing the computer to compress the color video image, after application of the first and second QP values, to a compressed output image.
40. The computer program of claim 39, further including instructions for causing a decompression computer to

decompress the compressed output image using the first and second QP values to obtain an uncompressed video image.

41. A computer program, stored on a computer-readable medium, for achieving higher compression during compression of a color video image in a YUV video image compression system utilizing a quantization parameter (QP) during compression, the computer program comprising instructions for causing a computer to utilize a first QP value for the Y color channel of a color video image, and a second QP value for at least one of the U and V color channels of the color video image, wherein the second QP value is greater than the first QP value.
42. The computer program of claim 41, wherein the second QP value is determined by applying a bias value to the first QP value.
43. The computer program of claim 41, further including instructions for causing the computer to compress the color video image, after application of the first and second QP values, to a compressed output image.
44. The computer program of claim 14, further including instructions for causing a decompression computer to decompress the compressed output image using the first and second QP values to obtain an uncompressed video image.
45. A computer program, stored on a computer-readable medium, for improving the coding efficiency for a color space representation of a video image originally represented as linear RGB pixel values in a video image compression system, the computer program comprising instructions for causing a computer to transform the linear RGB pixel values of the video image to a logarithmic representation of luminance and chroma channel information.

46. The computer program of claim 45, wherein the instructions for causing the computer to transform include instructions for causing the computer to apply the following equations to obtain a YUV logarithmic representation of the video image:

$$Ylog = \text{Log} (W_r * R + W_g * G + W_b * B)$$

$$U \text{ chroma channel} = \text{Log}(R) - Ylog$$

$$V \text{ chroma channel} = \text{Log}(B) - Ylog$$

where W_r , W_g , and W_b are linear weightings for red, green, and blue components of luminance of the video image.

47. The computer program of claim 46, further including instructions for causing the computer to reduce the resolution of the U and V chroma channels of the YUV logarithmic representation.

48. The computer program of claim 46, further including instructions for causing the computer to compress the YUV logarithmic representation of the video image to a compressed video image.

49. The computer program of claim 48, further including instructions for causing a decompression computer to decompress the compressed video image to a restored YUV logarithmic representation of the video image.

50. The computer program of claim 49, further including instructions for causing the decompression computer to transform restored YUV logarithmic representation of the video image to linear RGB pixel values.

51. The computer program of claim 50, wherein the instructions for causing the computer to transform includes instructions for causing the computer to apply the following equations to obtain the linear RGB pixel values:

$$R = \text{anti-log}(Y + U)$$

$$B = \text{anti-log}(Y + V)$$

$$G = (\text{anti-log}(Y) - W_r * R - W_b * B) / W_g.$$

52. A computer program, stored on a computer-readable medium, for improving the video characteristics of a color video image in a video compression system, the computer program comprising instructions for causing a computer to:
 - (a) select a set of image channels to represent the color video image, including a luminance channel and n chroma channels, where n is at least three; and
 - (b) compress the luminance channel and the n additional chroma channels to a compressed video image.
53. The computer program of claim 52, wherein at least one chroma channel represents non-visible wavelengths.
54. The computer program of claim 52, wherein the luminance channel is the image channel having the highest dynamic range and resolution.
55. The computer program of claim 52, further including instructions for causing a computer to code each chroma channel independently from each other channel.
56. The computer program of claim 52, further including instructions for causing a computer to code each chroma channel differentially with respect to a selected other channel.
57. The computer program of claim 52, further including instructions for causing a computer to reduce the resolution of at least one chroma channel.
58. The computer program of claim 52, further including instructions for causing a computer to apply a quantization parameter (QP) value to at least one chroma channel biased with respect to a QP value applied to the luminance channel.

59. A system for compressing image chroma information of a color video image in a video image compression system, including:
- (a) means for selecting a resolution for a red color component of the color video image that is higher than the resolution for a blue color component of the color video image; and
 - (b) means for applying the selected resolution to compress the color video image.
60. A system for compressing image chroma information of a color video image in a video image compression system, including means for:
- (a) downfiltering a blue color component of the color video image to a processed blue color component having a first resolution along at least one of the horizontal and vertical image axes of the color video image; and
 - (b) filtering a red color component of the color video image to a processed red color component having a second resolution higher than the first resolution.
61. The system of claim 60, wherein the second resolution is in the range from 0.5 to 1.0 of the full resolution of the red color component along at least one of the horizontal and vertical image axes of the color video image.
62. The system of claim 60, further including means for compressing at least the processed blue color and red color components to a compressed output image.
63. The system of claim 62, further including means for decompressing the compressed output image to obtain uncompressed processed blue and red color components.

64. The system of claim 63, further including means for upsize filtering the processed blue and red color components to the full resolution of the color video image.
65. The system of claims 59 or 60, wherein the video image compression system is a motion-compensated video image compression system.
66. A system for reducing chroma noise during compression of a color video image in a YUV video image compression system utilizing a quantization parameter (QP) during compression, including
- (a) means for utilizing a first QP value for the Y color channel of a color video image, and a second QP value for at least one of the U and V color channels of the color video image, wherein the second QP value is less than the first QP value; and
 - (b) means for applying the selected QP values during compression of the color video image.
67. The system of claim 66, wherein the second QP value is determined by applying a bias value to the first QP value.
68. The system of claim 66, further including means for compressing the color video image, after application of the first and second QP values, to a compressed output image.
69. The system of claim 68, further including means for decompressing the compressed output image using the first and second QP values to obtain an uncompressed video image.
70. A system for achieving higher compression during compression of a color video image in a YUV video image compression system utilizing a quantization parameter (QP) during compression, including:

- (a) means for utilizing a first QP value for the Y color channel of a color video image, and a second QP value for at least one of the U and V color channels of the color video image, wherein the second QP value is greater than the first QP value
- (b) means for applying the selected QP values during compression of the color video image.

71. The system of claim 70, wherein the second QP value is determined by applying a bias value to the first QP value.
72. The system of claim 70, further including means for compressing the color video image, after application of the first and second QP values, to a compressed output image.
73. The system of claim 72, further including means for decompressing the compressed output image using the first and second QP values to obtain an uncompressed video image.
74. A system for improving the coding efficiency for a color space representation of a video image originally represented as linear RGB pixel values in a video image compression system, including:
- (a) means for inputting linear RGB pixel values of a video image; and
 - (b) means for transforming the linear RGB pixel values of the video image to a logarithmic representation of luminance and chroma channel information.
75. The system of claim 74, wherein transforming includes applying the following equations to obtain a YUV logarithmic representation of the video image:
- $$Ylog = \text{Log} (W_r * R + W_g * G + W_b * B)$$
- $$U \text{ chroma channel} = \text{Log}(R) - Ylog$$
- $$V \text{ chroma channel} = \text{Log}(B) - Ylog$$

where W_r , W_g , and W_b are linear weightings for red, green, and blue components of luminance of the video image.

76. The system of claim 75, further including means for reducing the resolution of the U and V chroma channels of the YUV logarithmic representation.

77. The system of claim 75, further including means for compressing the YUV logarithmic representation of the video image to a compressed video image.

78. The system of claim 77, further including means for decompressing the compressed video image to a restored YUV logarithmic representation of the video image.

79. The system of claim 78, further including means for transforming restored YUV logarithmic representation of the video image to linear RGB pixel values.

80. The system of claim 79, wherein transforming includes applying the following equations to obtain the linear RGB pixel values:

$$R = \text{anti-log}(Y + U)$$

$$B = \text{anti-log}(Y + V)$$

$$G = (\text{anti-log}(Y) - W_r * R - W_b * B) / W_g.$$

81. A system for improving the video characteristics of a color video image in a video compression system, including means for:

- (a) selecting a set of image channels to represent the color video image, including a luminance channel and n chroma channels, where n is at least three; and
- (b) compressing the luminance channel and the n additional chroma channels to a compressed video image.

82. The system of claim 81, wherein at least one chroma channel represents non-visible wavelengths.

83. The system of claim 81, wherein the luminance channel is the image channel having the highest dynamic range and resolution.
84. The system of claim 81, further including means for coding each chroma channel independently from each other channel.
85. The system of claim 81, further including means for coding each chroma channel differentially with respect to a selected other channel.
86. The system of claim 81, further including means for reducing the resolution of at least one chroma channel.
87. The system of claim 81, further including means for applying a quantization parameter (QP) value to at least one chroma channel biased with respect to a QP value applied to the luminance channel.