




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04N 7/32	A1	(11) International Publication Number: WO 99/60793 (43) International Publication Date: 25 November 1999 (25.11.99)
(21) International Application Number: PCT/US99/10605 (22) International Filing Date: 13 May 1999 (13.05.99) (30) Priority Data: 09/083,383 21 May 1998 (21.05.98) US (71) Applicant (for all designated States except US): INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US). (72) Inventors; and (75) Inventors/Applicants (for US only): ACHARYA, Tinku [IN/US]; 7292 South Roberts Road, Tempe, AZ 85283 (US). KARAM, Lina, J. [LB/US]; 4619 E. Piedmont Road, Phoenix, AZ 85044 (US). MARINO, Francescomaria [IT/US]; Apartment 125, 1019 E. Lemon Street, Tempe, AZ 85281 (US). (74) Agents: MILLIKEN, Darren, J. et al.; Blakely, Sokoloff, Taylor & Zafman LLP, 7th floor, 12400 Wilshire Boulevard, Los Angeles, CA 90025 (US).		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i>

(54) Title: THE COMPRESSION OF COLOR IMAGES BASED ON A 2-DIMENSIONAL DISCRETE WAVELET TRANSFORM YIELDING A PERCEPTUALLY LOSSLESS IMAGE

**(57) Abstract**

A coding method and apparatus that include splitting raw image data (200) into a plurality of channels (R, G1, G2, B) including color plane difference channels (R-G1, B-G2), and then compressing separately each of these channels using a two-dimensional discrete wavelet transform (210, 212, 214, 216), the compression utilizing quantization, whereby the recovery of the compressed channel data yielding a perceptually lossless image. The method and apparatus operate on images directly in their Bayer pattern form. Quantization thresholds are defined for the quantizing which may vary depending upon the channel and DWT sub-band being processed.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

THE COMPRESSION OF COLOR IMAGES BASED ON A 2-DIMENSIONAL
DISCRETE WAVELET TRANSFORM YIELDING A PERCEPTUALLY LOSSLESS
IMAGE

BACKGROUND OF THE INVENTION

1. Field of the Invention

The invention relates generally to image processing. More specifically, the invention relates to encoding and quantization for image compression.

2. Description of the Related Art

When an image is captured of a scene, environment or object, the image is represented by an array of locations known as pixels. Each pixel of an image has a value relating to one or more color planes. When an image is captured by an imaging device such as a digital camera it is often captured into a Color Filter Array (CFA) pattern known as the Bayer pattern. In the Bayer pattern, each pixel location is an intensity value relating to only one of the three primary rendering colors (Red (R), Green (G) and Blue (B)). The Bayer pattern arranges pixels as follows:

G R G R ...
B G B G ...
G R G R ...
B G B G ...

Since there are twice as many G related pixels as either of B or R pixels, the G or Green color plane may be considered as two separate color planes G1 (G pixels on the same row as R pixels) and G2 (G pixels on the same row as B pixels). Thus, a Bayer pattern "raw" image can be considered to contain four independent color planes. To obtain a full resolution color image (e.g., for rendering), each pixel location should have all three R, G and B components, not just one. To achieve this a process known as color interpolation is employed where

missing color components for a pixel are estimated based on neighboring pixels.

After an image is captured and perhaps color interpolated then the image is often "compressed" or reduced in terms of the total number of bits that would be needed to store or transmit the image. Such image compression is ordinarily applied after color interpolation, but it would be advantageous in certain instances to perform suitable compression before color interpolation while an image is still in the Bayer pattern raw image format. Image compression plays a key role in multimedia applications such as video conferencing, digital imaging and video streaming over a network. Image compression schemes for such applications should be designed to reduce the bit-rate of storage or transmission of the image while maintaining acceptable image quality for the specific application.

Image compression techniques can be classified as either "lossy" or "lossless". With lossless compression, the original image prior to compression can be exactly recovered when the compressed image is decompressed. Consequently, lossless techniques, whose compression ratios depend upon the entropy of an image, do not ordinarily achieve high compression ratios and, since they preserve a high percentage of original image information, may also be computationally expensive. By contrast, lossy compression schemes provide only an approximation of the original image. Thus, with lossy compression, greater compression ratios can be achieved but often with loss in image quality compared to lossless techniques. One such lossy technique is a transform-based coding known as JPEG (Joint Photographic Experts Group) which transforms pixels of an input image using the well-known Discrete Cosine Transform (DCT). The resulting transformed pixel values are quantized or mapped to smaller set of values in order to achieve compression. The quality of a compressed

image that is decompressed will depend greatly on how the quantization of the transformed pixels are performed. The compression ratio (the size of the original raw image compared to the compressed image) will also be affected by the quantization, but can be enhanced by the binary encoding of the data after quantization.

Further, high compression ratio algorithms such as JPEG suffer from deficiencies such as "blocking artifacts". For these algorithms, an image is divided into blocks of pixels such as 8x8, or 16x16 blocks. These blocks are processed independently of each other and thus, between blocks, there is an observable discontinuity in luminance or color which constitutes a "blocking artifact".

These and other image compression schemes which achieve high compression ratios and sometimes, also acceptable decompressed image quality, operate on images better when the images are in "luminance-chrominance" format. Unlike Bayer pattern or color interpolated full RGB image "spaces" (i.e. formats) which represent a pixel's color as a selected mixture of primary colors (such as red, green and blue), luminance-chrominance format images define each pixel in terms of hue and saturation levels. Since imaging devices such as digital cameras ordinarily capture images in Bayer pattern format, an image must first be color interpolated into full resolution [RGB] and then have its "color space" converted into a luminance-chrominance format such as YCrCb before luminance-chrominance techniques can be applied. Such color interpolation and color space conversion is often cost-prohibitive as well as time-consuming and thus, not desirable.

Figure 1 shows one such traditional approach. An original image 100, captured for instance from a device such as a digital camera, is ordinarily in a raw image format such as the Bayer pattern. As such, each pixel will not have full color representation. Thus, the image is passed, either

entirely at once or, block by block, to a color interpolation pixel procedure 110. The color interpolation procedure generates a full-color pixels from the image 100, each pixel of which has full color resolution (e.g., R, G, and B components). The full color image is then color space converted (block 120) from RGB to YUV or other appropriate space. Such conversion may improve the compression ratio that is achievable. Once converted, the image is then passed to a primary compression procedure (block 130). Such compression may include a variety of procedures, such as JPEG or Fourier analysis, etc., but often has, as a component, a procedure known as quantization. An image is quantized by mapping a range of values representing the image pixels to a smaller range of values. After compression, the compressed image values can be encoded (block 140) such that they are suitable for transmission or storage.

This traditional approach suffers several drawbacks. First, the entire procedure is computationally complex particularly in the color interpolation and color space conversion. The color space conversion alone requires (for RGB to YCrCb space, for example) nine multiplications and six additions for each and every pixel. Often, such complicated techniques are unable to be effectively implemented in small, cost-conscious devices such as digital cameras.

If images are to be compressed on a digital camera or other imaging device, the compression techniques described above would be ineffective or prohibitive. Thus, there is a need for an image quantization and compression technique which is computationally inexpensive so as to reduce the cost of the digital cameras on which they are employed. To avoid the need for first performing color interpolation, a quantization and compression method should be developed that can be applied directly to Bayer pattern raw image data that is generated by portable imaging devices and that can exploit the correlation

between different color planes to achieve high compression ratios. Additionally, there is a need for enhancing the speed of quantization and compression so that the image capture and store in local memory or transfer from the imaging device can be performed more in real-time while still preserving image quality.

SUMMARY OF THE INVENTION

What is disclosed is a method that includes splitting raw image data into a plurality of channels including color plane difference channels, which exploits the correlation of color planes that compose the raw image data, and then compressing separately each of these channels using a two-dimensional discrete wavelet transform, the compression utilizing quantization, the decompression of the compressed channel data yielding a perceptually lossless image.

BRIEF DESCRIPTION OF THE DRAWINGS

The objects, features and advantages of the method and apparatus for the present invention will be apparent from the following description in which:

Figure 1 shows a traditional approach to image compression.

Figure 2 illustrates image compression data flow according to one embodiment of the invention.

Figure 3 illustrates recovery of a compressed and encoded image according to one embodiment of the invention.

Figure 4 shows the results of iteratively applying a 2-dimensional DWT to an image.

Figure 5 is a table of sample quantization threshold values for given sub-bands and channels.

Figure 6 is a block diagram of one embodiment of the invention.

Figure 7 is a block diagram of an image processing apparatus according to an embodiment of the invention.

Figure 8 is a system diagram of one embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

Referring to the figures, exemplary embodiments of the invention will now be described. The exemplary embodiments are provided to illustrate aspects of the invention and should not be construed as limiting the scope of the invention. The exemplary embodiments are primarily described with reference to block diagrams or flowcharts. As to the flowcharts, each block within the flowcharts represents both a method step and an apparatus element for performing the method step. Depending upon the implementation, the corresponding apparatus element may be configured in hardware, software, firmware or combinations thereof.

Figure 2 illustrates image compression data flow according to one embodiment of the invention.

It is desirable in digital applications such as still or motion imaging that an original image such as one captured by a digital camera be compressed in size as much as possible while maintaining a certain level of quality prior to its being transferred for decompression and displayed. Ideally, the compression technique chosen can also be applied to any kind of data transfer mechanism. The disclosed compression technique, that is the subject of one or more embodiments of the invention, has been specifically developed to adaptively utilize the response of the human visual system to color and light to maintain image quality.

As mentioned earlier, a raw image that is captured by a digital camera or other similar device will typically be represented in a Bayer pattern. The sensor array 200 is a set of pixel locations or "senses" that provide for each location,

an intensity value of the light incident upon the sensors from the environment/scene being imaged. In a Bayer pattern, each pixel location of an image in sensor array (hereinafter "original image") 200 will have an association with a color plane—Red(R), Green(G), or Blue(B). Since the Bayer pattern has two associated values for every R and B, the Green color plane may be considered as two planes G1 and G2. The G1 associated pixels lie in the Bayer pattern on the same row in original image 200 as R associated pixels, while G2 associated pixels lie on the same row as B associated pixels.

According to one embodiment of the invention, the correlation between an R associated pixel and its G1 associated neighboring pixel as well as the correlation between a B associated pixel and its neighboring G2 pixel are both exploited advantageously. The pixels in the Bayer pattern are capable of being subjected to compression directly without the need for color interpolation and/or color space conversion in this embodiment of the invention. The G1 and G2 associated pixels are passed directly to compression (blocks 212 and 216). The R and B pixels are treated less directly. The R pixel value is subtracted by its west neighboring G1 pixel value (block 205). This difference (R-G1) is passed to compression (block 210). Likewise, each B associated pixel is subtracted from its east neighboring G2 associated pixel (block 206). This difference (B-G2) is then passed to compression (block 216).

According to one embodiment of the invention, the difference channels, R-G1 and B-G2 are created in order to take advantage of strong correlation between color planes. These "channels" along with the G1 and G2 channels are each passed to appropriate compression stages. The pure color channels are decorrelated in one embodiment by using subtraction, but other methods of decorrelation may also be utilized. Since green is the most perceptible color (of the

three—R, G and B) to the human vision system, the Green planes G1 and G2 are preserved as channels and utilized as reference channels for decorrelating Red and Blue.

Each of the four channels, R-G1, G1, G2, and B-G2 are passed to compression blocks 210, 212, 214, and 216, respectively. In each compression block 210, 212, 214 and 216, according to an embodiment of the invention, two processes occur. The first process is a 2-Dimensional Discrete Wavelet Transform (2-D DWT). The DWT is more useful in image compression than Fourier or other periodic-based transforms since it describes abrupt changes, discontinuities, and thus, edge features of images more accurately and efficiently. The 2-D DWT generates "sub-bands" of the image as shown and described below with respect to **Figure 4**. After the DWT is performed, a second process known as quantization is performed.

Quantization is the procedure of mapping a set of n possible values to a set of m possible, where $m < n$. By quantizing, the total number of possible data values for the DWT image data set is reduced. The mapping is achieved according to some mathematical formula $y = f(x)$, where x is the DWT data value and y is the quantized data value. With such a formula, the number of total bits needed to represent the image is diminished. While this introduces some error, there are several methods in the art which can be employed to reduce the error. After the transformed image data is quantized, it is then encoded. Encoding 130 arranges (packs) the quantized data so that it has a convenient representation. The compressed and encoded image data may then be stored onto a media, transmitted from one system to another or distributed over a communication pathway such as a network. Further, the compressed and encoded image data need not be collected and transferred as a single frame, but can be streamed, encoded value by encoded value, out to its destination.

Depending upon the precise parameters used for DWT transformation, quantization and encoding, the compression ratio, which is the size of the original image divided by the size of the compressed image, will vary. This embodiment of the invention provides for an enhanced compression process that can serve to advantageously increase decompressed image quality, decrease the complexity of decompress and optimize the compression ratio.

Given that other image compression techniques may also utilize the DWT, the quality of the decompressed image will depend in great part upon the quantization utilized. One important aspect of the invention is a perceptually lossless quantization approach, the results of which are a lossy compression that is perceived by the human vision system to be lossless when decompressed. Further, the quantization approach adopted in this embodiment of the invention is capable of fast and easy computation giving more real-time performance to the hardware on which the compression is implemented. By exploiting the property of the DWT to create sub-bands of the image, an adaptive quantization procedure is provided in one embodiment of the invention that is responsive to sub-band properties and color channel properties.

For each channel, (R-G1, G1, G2 and (B-G2), a quantization threshold value is defined for each image sub-bands generated by the 2-D DWT process. Each such threshold $Q(s,c)$, where "s" represents the sub-band and "c", the channel, is used in quantizing the DWT result values in that channel "c" and sub-band "s". Thus, for values X_{sc} (or DWT coefficients used in obtaining those values), the quantized value, Y_{sc} , is simply

$$Y_{sc} = \text{round}\left(\frac{X_{sc}}{Q(s,c)}\right)$$

where the function round (k) rounds up or down the value k to the nearest integer. Thus, in given sub-band and channel, the

quantization is a scalar and uniform quantization formula, and, therefore, capable of fast and efficient hardware implementation. In one embodiment of the invention, the quantization thresholds modify the DWT coefficients themselves, thus eliminating the need for separate quantization (see **Figure 6** and associated description). Further, the use of a threshold adapted particularly depending on the channel and sub-band greatly enhances the quality of the recovered image upon decompression over quantization techniques that are uniform or arbitrary with respect to colors (channels) and edge perceptibility (a function of DWT sub-band resolving). The absolute error introduced during quantization is equal or less than $Q(s,c)/2$. In one embodiment of the invention, experimentally derived values for $Q(s,c)$ are determined such that the error leads to no perceptual loss of image quality.

Figure 3 illustrates recovery of a compressed and encoded image according to one embodiment of the invention.

The decoding blocks, inverse quantization blocks and inverse DWT blocks comprise a process which attempts to recover the original image 200 from the compressed and encoded image data 240 (see **Figure 2**). The decoded and decompressed image obtained will not be an exact pixel-for-pixel reconstruction of the original image, since the compression is "lossy". However, by utilizing DWT properties and perceptually lossless quantization techniques that are the subject of various embodiments of the invention, the loss can be made imperceivable to human vision, and thus the quality of the decompressed image is increased over other lossy techniques. Further, the ease of the inverse DWT procedure when compared with other inverse techniques makes it suitable for fast and easy implementation.

The compressed and encoded image data 240 may be efficiently stored channel by channel and sub-band by sub-band

(see **Figure 4**). Thus, the compressed and encoded channels (R-G1), G1, G2 and (B-G2) may be separately decoded and decompressed. First, the data belong to each channel is decoded (for instance, Zero Run length, Huffman decoding, etc.) (blocks 310, 312, 314 and 316). Each channel and sub-band of data, may have been encoded using techniques different from those of other sub-bands, and channels, thus, will need to be decoded taking any differences in encoding technique into account. Each channel of decoded data is then decompressed (blocks 320, 322, 324 and 326). As with the compression blocks shown in **Figure 2**, the decompression consists of two procedures — dequantizing the decoded data and then performing an inverse DWT (IDWT).

The dequantization block will simply multiply the decoded data (which is the quantized DWT coefficients) values by the quantization threshold $Q(s,c)$ for the given sub-band and appropriate channel. After dequantization, an inverse DWT is performed for each channel and sub-band's data. Once the IDWT is completed, an approximation of the original image 200 may be obtained pixel by pixel. By adding back G1 to the (R-G1) recovered value (block 325) and G2 to the (B-G2) recovered value (block 326), each Bayer pattern pixel value R, G1, G2 and B from the original sensor array 200 may be approximately recovered. The recovered R, recovered G1, recovered B2 and recovered B values may or may not be identical with the values of original image 200 but will show visually lossless or perceptually lossless properties due to the exploitation of color channel correlation. Thus, the recovered image will be of high quality. According to another embodiment of the invention, the dequantization process may be merged with the inverse DWT by modifying the inverse DWT coefficients by the appropriate quantization thresholds.

The decompression can be implemented as hardware, software or from one or a combination thereof and can be

separate physically from the apparatus performing the function of the encoding compression process. The basic data flow for lossy compression schemes consist of compression and decompression and often will also include an intermediate transfer from the compression block to the desired destination which has access to decompression capability.

Figure 4 shows the results of iteratively applying a 2-dimensional (2-D) DWT to an image.

As described in co-pending U.S. patent application entitled, *An Integrated Systolic Architecture for Decomposition and Reconstruction of Signals Using Wavelet Transforms*, Serial No. 08/767,976 (hereinafter "Patent Application '976"), application of 2-D DWT upon an image space will result in the creation of four "sub-bands." For instance, **Figure 4** shows that an image *S* is decomposed by the 2-D DWT into four sub-bands *S1*, *S2*, *S3* and *S4*. Of these, the most critical sub-band is *S1*. The sub-band *S1* is also referred to as the "LL" sub-band, based on the double low-pass filtering used to generate it. *S1* (LL) is essentially in scaled approximation of the original image *S*, and contains the most salient image information. The sub-bands *S2*, *S3* and *S4* contain edge information and when the input image is noisy, also a considerable amount of that noise. The sub-bands *S2*, *S3* and *S4* are also referred to as HL, LH and HH sub-bands, respectively, due to the various low-pass and high-pass filtering used to generate them. Since the sub-bands *S2*, *S3* and *S4* are perceptually less significant than the *S1* sub-band, these sub-bands may be more "roughly" quantized (i.e., assigned a higher threshold *Q*) so that the values therein are compressed greater. The *S1* sub-band may not even need to be quantized directly, since this sub-band is utilized in generating higher level DWTs. As mentioned earlier, the original image, according to one embodiment of the invention, is subjected to a 2-D DWT channel by channel. The four

channels utilized in one embodiment of the invention include (R-G1), G1, G2 and (B-G2). The data composing each of these channels may be considered an "image" in their own right upon whom a 2-D DWT is performed. The four sub-bands S1, S2, S3 and S4 constitute a level 1 DWT. Thus, the subscript 1 in **Figure 4** below the designations LL_1 , HL_1 , LH_1 and HH_1 indicate that these sub-bands belong to level 1.

The level 1 sub-bands S1, S2, S3 and S4 result from applying the 2-D DWT once to the image S. If the 2-D DWT is applied again, to the sub-band result S1, a two level 2-D DWT is said to have been performed. The level 2 2-D DWT would result in the generation of four new sub-bands S11, S12, S13 and S14 which are sub-bands generated from the sub-band S1 from the level 1 2-D DWT. These sub-bands S11, S12, S13 and S14 have the designation LL_2 , HL_2 , LH_2 and HH_2 , respectively, since they are level 2 DWT sub-bands. Again, the LL_2 sub-band S11 contains the most salient features from S1 while the S12, S13, and S14 sub-bands contain edge and possibly noise information from the sub-band S1. The 2-D DWT may be thus applied many times to LL sub-band of each level to obtain more and more levels of DWT resolution and thus, image sub-bands. According to one embodiment of the invention, only a level 1 2-D DWT procedure is considered. If more levels of 2-D DWT processing occurs, each of the newly created sub-bands would be assigned a Q or quantization threshold for each channel present therein. The determination of the $Q(s,c)$ value for a given sub-band "s" and channel "c" has been arrived at empirically for a 9-7 bi-orthogonal spline DWT filter. The results of this study are tabulated in **Figure 5**.

Figure 5 is a table of sample quantization threshold values for given sub-bands and channels.

The quantization thresholds $Q(s,c)$ may be determined/selected in a wide variety of ways. In one study conducted in conceiving the various embodiments of the

invention, empirical data regarding the perception of a group of observers to a set of DWT compressed images was collected. In these experiments, the thresholds $Q(s,c)$ were increased until artifacting due to quantization was observable. A wide variety of images, each with differing characteristic, were selected for the study. The table of **Figure 5** illustrates the results of the study and are statistically assured to be applicable to any image provided that the 9-7 bi-orthogonal spline filter, which is well-known in the art, is used as a basis for the DWT. When a different technique such as DCT or different basis for DWT, is utilized, new quantization thresholds may need to be selected or determined since each filter renders differently the same image. In one embodiment of the invention, such values may be preloaded into a ROM or in another embodiment, these values may be written to a rewritable memory so that they can be modified.

Referring to **Figure 4**, the sub-bands labeled S1, S2, S3 and S4 all belong to a level 1 DWT. The sub-band S4 has quantization thresholds approximate 5 times greater than the S2 and S3 sub-bands. Thus, the information (data values) in the S4 sub-band are quantized to a greater degree, and hence, more compressed. The error implicit in a higher quantization threshold and, consequently, fewer mapped values is tolerable since the S4 sub-band contains the least relevant visually perceptible image details, such as diagonal edges and noise. As mentioned above, S1 contains most of the salient and visually crucial information of the original image S. For a k level DWT, the lowest k-1 LL sub-bands are preserved (i.e., Q of 1) and thus not quantized since these sub-bands are themselves resolved further into higher level LL, LH, HL and HH sub-bands. The LL_k or highest level DWT sub-band is quantized ($Q>1$) since there is no higher level resolution of LL_k within which the quantization would have been accounted for. The sub-bands S2 and S3 for all levels have quantization

thresholds that lie in-between S1 and S4 sub-bands of the same level.

With regard to channels, quantization thresholds were determined for R, G, B, and then (R-G1), and (B-G2). The G values hold true for both G1 and G2. In general, Blue can be quantized more roughly (with higher thresholds) than Green (G1, G2) and Red. However, in embodiment of the invention, the channels (R-G1) and (B-G2) rather pure R and B are considered in the compression process. These "difference" or decorrelated channels have much higher quantization thresholds than the pure color channels R, G, and B. This is due to the fact that the edge information of the image is accounted for in G1 and G2 planes (channels). When these G1 and G2 values are subtracted, from R and B plane values, respectively, the resulting difference preserves the chrominance component in R and B planes not found in G1 and G2. Thus, the difference channels (R-G1) and (B-G2) optimized the contribution of R and B planes to the overall image and its perceived quality. Observations have shown that the S4 sub-band in the difference channels (R-G1) and (B-G2) contain no image information that is perceptually different from the information contained in G1 and G2 channels, and thus, zero values are assigned to the entire sub-band (a Q value of 00). The sub-band S4 does not need to be stored for the difference channels since there is no perceivable information, therein, according to one embodiment of the invention. The higher the DWT level, the more precision or resolution is obtained of the input sub-band LL of the previous level. Though a 3-level DWT is shown in **Figure 4**, any number of DWT levels may be generated according to design needs.

Figure 6 is a block diagram of one embodiment of the invention.

According to an embodiment of the invention described with respect to **Figures 2** and **3**, compression and

decompression processes are divided into two stages -- one for quantization (and dequantization) and another separate stage for the DWT (and inverse DWT). However, according to another embodiment of the invention, the quantization (dequantization) stages can be coalesced into the DWT (and inverse DWT). The DWT and inverse DWT outputs are generated by a set of cascading filters (see Patent Application '976) whose coefficients are the coefficients of the DWT (and inverse DWT) function. These coefficients are multiplied by input pixel values and the products selective added to generate the DWT outputs. If the quantization thresholds are combined algebraically with the DWT coefficients, quantization will be achieved during the DWT computation itself.

In the embodiment of **Figure 6**, the 2-D DWT is implemented by repeating the one-dimensional (1-D) DWT twice. This approach is possible due to the separability of the DWT filters (see Patent Application '976). It may be possible to implement 2-dimensional DWT or other 2-dimensional transforms using a single bi-dimensional filter given that such a filter is feasible. By contrast, as outlined in Patent Application '976, the one-dimensional filter approach performs the DWT in a row-wise fashion and then in a column-wise fashion to the row-wise DWT result. For instance, consider a channel "c" 600 as shown in **Figure 6**. This represents, in one embodiment of the invention, the pixel data from a particular color/difference channel, G1, G2, (R-G1) or (B-G2), but may also, in another embodiment, represent an entire image or image portion. First, a row-wise DWT is performed by the 1-D DWT module 610. A set of control signals #1 regulates the operation of this row-wise DWT and can supply coefficients depending on the level of the DWT (see below). The module 610 generates and "L" band and "H" band from the channel "c".

Once the row-wise DWT is performed, the resulting "L" and "H" bands are transposed by a matrix transposer circuit 620.

Such matrix transposer circuits are well known in the art. The matrix transposer circuit 620 provides column-by-column the result from matrix transposer circuit 620 as input to the second 1-D DWT module 630. The second 1-D DWT module 630 is regulated and provided coefficients, if necessary, by means of a set of control signal #2. The result of performing a 1-D DWT column-wise on the 1-D DWT, row-wise data transposed from matrix transposer circuit 620 is shown as the 2-D DWT result data 640. Each pass through the row-wise, 1-D DWT, transpose and column-wise 1-D DWT is equivalent to performing a 2-D DWT. The result data 640 is composed sub-bands LL, HL, LH and HH and comprises one level of the DWT, as referred to and described in **Figure 4**.

The process described above is to generate the results from one level of the DWT. If more than one level, such as three levels of DWT resolution is desired, then a counter may be utilized and loaded with the value 3. Each instance that a 2-D DWT cycle is completed, the count is decremented. A decision block 650 may check this counter to determine if another DWT level is needed. If another level is desired, the "LL" - sub-band is piped back into the 2-D DWT process to generate another set of sub-bands therefrom. **Figure 4**, for instance, shows a 3-level 2-D DWT result. At each level, the sub-band LL_k , where k is the level, is used as the input and then decomposed into four further sub-bands by use of the 2-D DWT. This procedure repeats until the last desired level of DWT resolution is reached. Also, when each level of the DWT is complete, the sub-bands HL, LH and HH are sent to an encoder 660 which performs binary encoding such as Huffman or Run Length Encoding upon the data. The encoded data is then stored as a portion of a compressed channel c'670. At each level before the last level of DWT resolution, the LL sub-band is not encoded since it is being fed back to the 2-D DWT process to generate further sub-bands. At the last level of

the DWT, the LL sub-band is sent to be encoded by encoding 660. The output of encoder 660 is stored and assembled and will constitute a compressed channel c'670 when complete. In the above-described manner, each of the channels R-G1, G1, G2 and B-G2 processed into compressed channels.

To achieve quantization during the DWT filtering (performed by DWT modules 610 and 630), the filtering coefficients must be modified by the quantization coefficients $Q(s,c)$, where s is the sub-band and c , the channel. The modification of the DWT coefficients varies according to the nature of the filter and the sub-band being processed and is summarized below:

ROW-WISE 1-D DWT:

- low-pass filtering over the LL_{k-1} sub-band (or over the source image for $k=1$) (generation of the sub-band L): each weight (coefficient) l_i of the filter is scaled by the factor:

$$\frac{1}{\sqrt{Q(LL_k, c)}}$$

- high-pass filtering over the LL_{k-1} sub-band (or over the source image) (generation of the sub-band H): each weight (coefficient) h_i of the filter is scaled by the factor:

$$\frac{\sqrt{Q(LL_k, c)}}{Q(HL_k, c)};$$

COLUMN-WISE 1-D DWT:

- low-pass filtering over the L and H sub-bands (generation of the sub-bands LL and LH): each weight (coefficient) l_i of the filter is scaled by the factor:

$$\frac{1}{\sqrt{Q(LL_k, c)}}$$

- high-pass filtering over the L sub-band (generation of the sub-band HL): each weight (coefficient) h_i of the filter is scaled by the factor: $\frac{\sqrt{Q(LL_k, c)}}{Q(HL_k, c)}$

- high-pass filtering over the H sub-band (generation of the sub-band HH): each weight (coefficient) h_i of the filter is scaled by the factor: $\frac{(HL_k, c)}{Q(HH_k, c)\sqrt{Q(LL_k, c)}}$;

where $Q(HL_k, c)$, $Q(HH_k, c)$ and $Q(LL_k, c)$ are respectively the perceptually lossless thresholds of the sub-bands HL, HH and LL in the k^{th} level for the channel c . The above conditions directly achieve the quantization since $Q(HL_k, c) = Q(LH_k, c)$. In fact, after the row-wise and the column-wise filtering the four sub-bands, at any level, result in being scaled (i.e., quantized), respectively by the factors:

- LL sub-band: $\frac{1}{\sqrt{Q(LL_k, c)}} \frac{1}{\sqrt{Q(LL_k, c)}} = \frac{1}{Q(LL_k, c)}$;
- LH sub-band: $\frac{1}{\sqrt{Q(LL_k, c)}} \frac{\sqrt{Q(LL_k, c)}}{Q(HL_k, c)} \frac{1}{\sqrt{Q(LL_k, c)}} = \frac{1}{Q(HL_k, c)}$;
- HL sub-band: $\frac{\sqrt{Q(LL_k, c)}}{Q(HL_k, c)} \frac{1}{\sqrt{Q(LL_k, c)}} \frac{\sqrt{Q(LL_k, c)}}{Q(HL_k, c)} = \frac{1}{Q(HL_k, c)}$;
- HH sub-band: $\frac{\sqrt{Q(LL_k, c)}}{Q(HL_k, c)} \frac{Q(HL_k, c)}{Q(HH_k, c)\sqrt{Q(LL_k, c)}} = \frac{1}{Q(HH_k, c)}$;

Figure 7 is a block diagram of an image processing apparatus according to an embodiment of the invention.

Figure 7 is a block diagram of internal image processing and compression components of an imaging device incorporating at least one embodiment of the invention. In the exemplary circuit of **Figure 7**, a sensor 700 generates pixel components which are color/intensity values from some scene/environment. The n-bit pixel values generated by sensor 700 are sent to a capture interface 710. Sensor 700 in the context relating to the invention will typically sense one of either R, G, or B components from one "sense" of an area or location. Thus, the intensity value of each pixel is associated with only one of three color planes and may form together a Bayer pattern such as that shown above. Capture interface 710 resolves the image generated by the sensor and assigns intensity values to the

individual pixels. The set of all such pixels for the entire image is in a Bayer pattern in accordance with at least one of the embodiments of the invention.

It is typical in any sensor device that some of the pixel cells in the sensor plane may not respond to the lighting condition in the scene/environment properly. As a result, the pixel values generated from these cell may be defective. These pixels are called "dead pixels." The "pixel substitution" unit 715 replaces each dead pixel by the immediate previously valid pixel in the row.

A RAM table 716 consists of the row and column indices of the dead pixels, which are supplied by the sensor. This RAM table 716 helps to identify the location of dead pixels in relation to the captured image. Companding module 725 is circuit designed to convert each original pixel of n -bit (typically $n=10$) intensity captured from the sensor to an m -bit intensity value, where $m < n$ (typically, $m=8$). Companding module 725 is not needed if the sensor 700 and capture interface 710 provide an m -bit per-pixel value.

According to at least one embodiment of the invention, as described above, sets of m -bit pixel value(s) may be directly compressed without resorting to color interpolation and/or color space conversion. Channel generator 727 is coupled to companding module 725 and can receive therefrom m -bit pixel data values which may be arranged according to the Bayer pattern. Each m -bit value is used by channel generator to generate the four channels (R-G1), G1, G2 and (B-G2). For instance, if pixels are captured row-by-row, a first row would yield R and G1 pixel values and thus outputs only at channels (R-G1) and G1. The next captured row would yield G2 and (B-G2) channels. The channel generator 727 sends two channels during one row, and the other remaining two channels during the next row. These channels are then input to a compressor/quantizer 728. A RAM table 729 can be used to

store DWT coefficients and/or quantization thresholds for each channel/sub-band as desired in executing the compression techniques described above. Further, add and multiply units, shifters, and control signaling can be provided in compressor/quantizer 728 to carry out the necessary DWT computation (see Patent Application '976). Compressor/quantizer 728 can be designed to provide high-pass and low-pass DWT outputs for each channel and sub-band. These compressed channel outputs, which represent the compressed image data, are then binary encoded by an encoder 730. Encoder 730 may use run-length, Huffman or other suitable coding to pack the compressed data for storage into storage array(s) 740.

Each of the RAM tables 716, 726 and 729 can directly communicate with bus 760 so that their data can be loaded and then later, if desired, modified. Further, those RAM tables and other RAM tables may be used to store intermediate result data as needed. Though the individual components (selectors, shifters, registers, add, multiply units and control/address signals) of modules 727, 728 and 730 have not been detailed, one skilled in the art will readily be able to implement such a device, given the details set forth for various embodiments of the invention.

Figure 8 is a system diagram of one embodiment of the invention.

Illustrated is a computer system 810, which may be any general or special purpose computing or data processing machine such as a PC (personal computer), coupled to a camera 830. Camera 830 may be a digital camera, digital video camera, or any image capture device or imaging system, or combination thereof and is utilized to capture an image of a scene 840. Essentially, captured images are processed by an image processing circuit 832 so that they can be efficiently stored in an image memory unit 834, which may be a ROM, RAM or

other storage device such as a fixed disk. The image contained within image memory unit 834 that is destined for computer system 810 can be according to one embodiment of the invention, stored directly as a compressed image. In most digital cameras that can perform still imaging, images are stored first and downloaded later. This allows the camera 830 to capture the next object/scene quickly without additional delay. The invention in its various embodiments, particularly in providing a compressed image that is directly converted from the captured 8-bit Bayer pattern, reduces the computation requirements of the camera 830 and the associated costs, allowing for a more inexpensive camera.

The image processing circuit 832 carries out the compression, quantization and encoding, directly from the Bayer pattern sense (with other intermediate steps such as pixel substitution or companding, see **Figure 7** and associated description) of camera 830 in this embodiment of the invention. When a compressed and encoded image is downloaded to computer system 810, it may be rendered to some output device such as a printer (not shown) or to a monitor device 820. If, according to one embodiment of the invention, the image is in Bayer pattern format after being decompressed, it may need to be converted to an RGB full color resolution format prior to rendering. Image decompression may be achieved using a processor 812 such as the Pentium® (a product of Intel Corporation) and a memory 811, such as RAM, which is used to store/load instruction addresses and result data and is a well-known operation in the art of colorimetry.

In an alternate embodiment, the compression process described above may be achieved in a software application running on computer system 810 rather than directly in camera 830. In such an embodiment, the image processing circuit may advantageously store only the Bayer pattern image. The application(s) used to perform the integrated color

interpolation and color space conversion after download from camera 830 may be from an executable compiled from source code written in a language such as C++. The instructions of that executable file, which correspond with instructions necessary to scale the image, may be stored to a disk 818 or memory 811. Further, such application software may be distributed on a network or a computer-readable medium for use with other systems. It would be readily apparent to one of ordinary skill in the art to program a computing machine to perform perceptually lossless quantized DWT compression an image if the methodology described above is followed.

When an image, such as an image of a scene 840, is captured by camera 830, it is sent to the image processing circuit 832. Image processing circuit 832 consists of ICs and other components which execute, among other functions, the DWT based perceptually lossless compression of an image. The image memory unit 834 will store the compressed channel data. Once all pixels are processed and stored or transferred to the computer system 810 for rendering the camera 830 is free to capture the next image. When the user or application desires/requests a download of images, the imaged stored in the image memory unit, whether stored as XYZ space images or as Bayer pattern images, are transferred from image memory unit 834 to the I/O port 817. I/O port 817 uses the bus-bridge hierarchy shown (I/O bus 815 to bridge 814 to system bus 813) to temporarily store the XYZ color space image data into memory 811 or, optionally, disk 818. Computer system 810 has a system bus 813 which facilitates information transfer to/from the processor 812 and memory 811 and a bridge 814 which couples to an I/O bus 815. I/O bus 815 connects various I/O devices such as a display adapter 816, disk 818 and an I/O port 817, such as a serial port. Many such combinations of I/O devices, buses and bridges can be utilized with the

invention and the combination shown is merely illustrative of one such possible combination.

In one embodiment of the invention, the compressed images can be decompressed/recovered to a perceptually lossless version on computer system 810 by suitable application software (or hardware), which may utilize processor 812 for its execution. A full resolution RGB image may be created by color interpolation data and then be rendered visually using a display adapter 816 into a perceptually lossless image 850. Since color interpolation and color space conversion are readily facilitated on-camera in one embodiment of the invention, it may be possible to implement a communication port in camera 830 that allows the image data to be transported directly to the other devices.

In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are accordingly to be regarded as illustrative rather than restrictive.

The exemplary embodiments described herein are provided merely to illustrate the principles of the invention and should not be construed as limiting the scope of the invention. Rather, the principles of the invention may be applied to a wide range of systems to achieve the advantages described herein and to achieve other advantages or to satisfy other objectives as well.

CLAIMS:

What is claimed is:

1. A method comprising the steps of:
splitting raw image data into a plurality of channels including color plane difference channels, said color plane difference channels exploiting the correlation of color planes composing said raw image data; and
compressing separately each of said channels using a two-dimensional discrete wavelet transform, said compression utilizing quantization, wherein the decompression of said compressed channel data yielding a perceptually lossless image.
2. A method according to claim 1 wherein the step of splitting includes the steps of:
arranging first color plane into a first color channel and second color channel;
generating a first difference channel that differences said first color channel from values associated with a second color plane; and
generating a second difference channel that differences said second color channel from values associated from a third color plane, said first color plane having twice the number of associated values as values associated with either said second or third color plane.
3. A method according to claim 2 wherein said first color plane is Green.
4. A method according to claim 2 wherein said second color plane is Red.

5. A method according to claim 2 wherein said third color plane is Blue.

6. A method according to claim 2 wherein said raw image data is arranged in a Bayer pattern.

7. A method according to claim 6 wherein said values of the first color channel are located on a first row in the Bayer pattern and values of the second color channel are located on a second row in the Bayer pattern, said second row immediately succeeding said first row.

8. A method according to claim 1 wherein the step of compressing includes the steps of:

performing a two-dimensional discrete wavelet transform (2-D DWT) on each channel, generating thereby for each channel a set of sub-bands, including LL sub-band containing salient channel information, the performing constituting a level of the 2-D DWT;

if further resolution is desired, performing a 2-D DWT on the LL sub-band of each channel generated in the preceding level, generating thereby four new sub-bands and constituting a further level of the 2-D DWT; and

quantizing said sub-bands of the 2-D DWT at each level, a separate quantization threshold defined for each said sub-band in each said channel at each level.

9. A method according to claim 1 wherein the step of compressing includes the steps of:

performing a two-dimensional discrete wavelet transform (2-D DWT) on each channel, generating thereby for each channel a set of sub-bands including an LL sub-band containing salient image information, said transform including filtering coefficients modified by quantization thresholds, said

thresholds defined for each channel in each sub-band, the performing constituting a level of the 2-D DWT; and

if further resolution is desired, performing a 2-D DWT on each channel of the LL sub-band generated in the preceding level, generating thereby four new sub-bands and constituting a further level of the 2-D DWT.

10. A method according to claim 8 wherein the step of performing a 2-D DWT includes the steps of:

performing a one-dimensional discrete wavelet transform (1-D DWT) row-wise to generate an "L" band and "H" band therefrom;

matrix transposing said "L" band and "H" band; and

performing a 1-D DWT column-wise on said "L" and "H" bands to generate therefor said LL, an LH, an HL and an HH sub-band.

11. A method according to claim 8 wherein the step of quantizing includes the step of:

dividing each of the values in each sub-band in each channel by the corresponding quantization threshold; and

rounding to the nearest integer the result of said division, said rounding yielding a quantized of the value version from said sub-band and said channel.

12. A method according to claim 9 wherein the modification of filtering coefficients varies according to the high-pass or low-pass nature of the filtering coefficients.

13. An apparatus for processing an image comprising:

a channel generator configured to generate both color difference channels and color channels from said image;

a compressor coupled to said channel generator, said compressor compressing each channel generated by said channel

generator, said compressed channels wherein decompression of said compressed channels yield a perceptually lossless version of said image.

14. An apparatus according to claim 13 wherein said compressor comprises:

a two-dimensional discrete wavelet transform module (2-D DWT) for each channel, said 2-D DWT module configured to a set of sub-band generate a set of sub-band outputs for each channel, including a low-low (LL) sub-band.

15. An apparatus according to claim 14 wherein said compressor further comprising:

a quantizer coupled to each said 2-D DWT module, said quantizer mapping the set of possible output values from said 2-D DWT module to subset of values therefrom, said mapping varying for each sub-band in each channel.

16. An apparatus according to claim 14 wherein each said 2-D DWT module utilizes quantization threshold modified filtering coefficients, said quantization threshold modification varying according to the particular sub-band, channel and nature of the filtering coefficients.

17. An apparatus according to claim 14 wherein said 2-D DWT module comprises:

a first one-dimension discrete wavelet transform (1-D DWT) module configured to perform a row-wise DWT, to generate thereby an "L" and an "H" bands;

a matrix transposer coupled to said first 1-D DWT module to receive said "L" and "H" bands, said transposer configured to transpose rows for columns in the channel data; and

a second 1-D DWT module coupled to receive said transposed "L" and "H" band data, and operating upon said data

to perform a column-wise DWT, said second 1-D DWT module configured to generate thereby a complete DWT level comprising said LL sub-band and an LH, HL, and HH sub-bands.

18. A method according to claim 1 wherein the step of compressing includes the steps of:

performing a two-dimensional discrete wavelet transform (2-D DWT) on each channel, generating thereby for each channel a set of sub-bands, the performing constituting a level of the 2-D DWT;

if further resolution is desired, performing a 2-D DWT on at least one of said sub-bands of each channel generated in the preceding level, generating thereby a new set of sub-bands and constituting a further level of the 2-D DWT; and

quantizing said sub-bands of the 2-D DWT at each level, a separate quantization threshold defined for each said sub-band in each said channel at each level.

19. A system comprising:

an image processor, said processor configured to split raw image data into a plurality of channels including color plane difference channels, said difference channels exploiting correlation between color planes that compose said raw image data, said processor compressing said channels using quantization and discrete wavelet transforms, the decompression of said compressed channel data yielding a perceptually lossless image; and

an image memory coupled to said processor, said memory configured to store said compression channel data.

20. A system according to claim 19 comprising:

a computer system coupled to said image memory to receive said compressed channel data and configured to decompress said

compressed channel data and render a perceptually lossless image therefrom.




FIG. 1

2/7




FIG. 2





FIG. 3

3/7



Sub-Band	Level 1 2D-DWT				Level 2 2D-DWT			Level 3 2D-DWT			
	Channel	S1	S2, S3	S4	S11	S12, S13	S14	S111	S112, S113	S114	
	R	1	22	100	1	15	51	15	15	30	
	G (i.e., G1 or G2)	1	20	90	1	15	40	15	14	26	
	B	1	24	130	1	18	65	15	17	39	
	R-G1	1	70	∞	1	42	90	22	202	50	
	B-G2	1	105	∞	1	65	130	52	53	106	

FIG. 5

5/7




FIG. 6

6/7





FIG. 7

7/7



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US99/10605

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :H04N 7/32

US CL :348/398

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)

U.S. : 348/397-398, 437-438; 382/240, 244; H04N 7/32

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,602,589 A (VISHWANATH ET AL) 11 February 1997, figs. 3, 6-7, 9, 11, col. 2, lines 37-48, col. 6, lines 1-67.	1, 13, 19-20
Y,P	US 5,850,482 A (MEANY ET AL) 15 December 1998, figs. 2, 7, col. 9, lines 31-67, col. 2, lines 1-67.	1, 13, 19-20
Y,P	US 5,880,856 A (FERRIERE) 09 March 1999, figs. 4-7, col. 4, line 1-43, col. 5, begin line 42 to col. 8, lines 1-24.	1, 13, 19-20

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

* Special categories of cited documents:	*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
A document defining the general state of the art which is not considered to be of particular relevance	*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
B earlier document published on or after the international filing date	*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
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*Z* document member of the same patent family
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	

Date of the actual completion of the international search

16 JULY 1999

Date of mailing of the international search report

09 SEP 1999

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

VU LE

Telephone No. (703) 308-6613