Coding Colour Quantized Images by Local Colour Quantization

G. Qiu
School of Computing & Mathematics, University of Derby
Keddleston Road Campus, Derby, United Kingdom

Abstract

Colour quantization is an established method for the reproduction of high visual quality colour imagery on systems with limited frame buffer resources. The ever increasing demand for fast and image/graphic intensive applications for cheap desktop platforms has called for greater savings in display and storage resources and simple, easy to implement architectures. Because human visual information system processes visual information over specific frequency ranges, it is possible to compress colour quantized image further by exploiting spatial visual redundancies. This paper presents a technique which reduces the frame buffer requirements of colour quantized images and at the same time maintains the visual quality of the images. The method has a simple hardware architecture and can be easily implemented using existing technology. Computer simulation results are presented to illustrate the effectiveness of the method.

1. Introduction

Most computer colour displays are based on a frame buffer architecture, in which the images is stored in a video memory, from which controllers constantly refresh the display screen [1]. Usually, digital images of natural scenes are captured as *true colour* images, where the colour of each pixel is represented by tristimuli of an appropriate colour space (e.g. RGB colour space) and each primary colour co-ordinate is quantized to 8-bit resolution. Because the actual number of colours present in most natural images is much smaller than the numbers of colours that can be represented by true colour resolution, most of today's colour CRT displays use an 8-bit frame buffer in combination with an RGB colour palette. Typically, 256 colours properly chosen for an image are sufficient to form a reproduction that is visually almost identical to the original.

Compared with 24-bit frame buffer architecture, a 3:1 reduction in buffer size is achieved using 8-bit frame buffer and a colour palette. The growing number of multimedia applications has created an increased demand for more efficient use of frame buffers to free up more space to support graphics and video applications. Compression of palettized colour images has attracted increasing interest [6, 7].

In traditional colour quantization, a colour palette is designed in 3-D colour space using some sort of clustering techniques such as vector quantization [14]. The representative colours are chosen from the ensemble of pixel colours present in the image according to some statistically optimal criteria. This approach can be regarded as global colour quantization. This colour selection process exploits statistical colour redundancies, however, spatial redundancy of the image is not exploited.

From human vision perspective, visual information is processed over specific frequency ranges, normally measured in cycle per viewing angle (c/deg). The overall passband of the receptive field of the normal human eye is in the range of 1 - 10 c/deg, with maximum response in the range between 3 - 6 c/deg; outside that range, the response fall off quickly [3]. In a typical viewing geometry of computer screen, only a single cycle within a small image block (e.g. 4 x 4 block) is perceivable [4]. This suggests that from human perception's point of view, no more than two colours are necessary to represent a small area of the image. Based on this reasoning, we designed our colour quantization scheme to compress palettized colour images.

The organization of the paper is as follows: In section 2, traditional colour quantization techniques are briefly described. Section 3 explains the procedure of local colour quantization of colour quantized images. Section 4 presents computer simulation results and concluding remarks are given in section 5.

2. Colour Quantization

Colour quantization of colour images has been researched extensively over the past two decades. We find it convenient to discuss colour quantization in the context of frame buffer image display, although colour quantization can also be used as a means for image data compression. In the context of image display, an image is first colour quantized, and a colour table, i.e., the palette, which contains the representative colours of the image is constructed. Each pixel of the image is represented by an index number, which identifies the colour entry in the colour palette for that pixel. To display the pixel, the colour of the entry is reproduced. A schematic diagram of a palettized image display system is shown in Fig. 1.

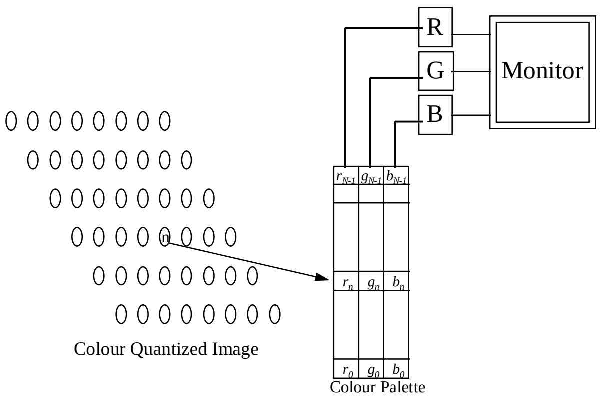

Figure 1 A schematic diagram of palettized image display system

Considerable research work has been done towards finding the colour palette [2, 9, 10]. These traditional colour quantization techniques design the colour palette in 3-D colour space using various forms of vector quantization or data clustering techniques [14]. These methods design the palette purely based on point level information and do not exploit spatial redundancies that exist in natural images. A more aggressive approach that forms low dimensional spatial vectors for colour palette design was recently introduced to achieve higher compression [5]. A different approach to achieve further data reduction is to apply compress/coding techniques to colour quantized images [6, 7]. We have investigated a simple and yet very effective method to compress colour quantized images.

3. Local Colour Quantization of Colour Quantized Images

The rationale for developing the current method is the psychophysics of human visual system. From human vision perspective, visual information is processed over specific frequency ranges, normally measured in cycle per viewing angle (c/deg). The overall passband of the receptive field of the normal human eye is in the range of 1 - 10 c/deg, with maximum response in the range between 3 - 6 c/deg; outside that range, the response fall off quickly [3]. In a typical viewing geometry of computer screen, only a single cycle within a small image block (e.g. 4 x 4 block) is perceivable [4]. This suggests that from human perception's point of view, no more than two colours are necessary to represent a small area of the image.

Exploiting this pychophysical evidence of human visual system for image data compression is not new. The earliest work in this area may be traced back to the block truncation coding (BTC) technique [11, 12]. Using BTC as a source coding technique for full colour image compression has also been investigated [13]. We shall show the principle of BTC can be used as an effective and simple method to compress colour quantized images. In addition to achieve reduction in data, and therefore frame buffer size, the present method also has a very simple and elegant architecture and can be easily implemented using special hardware.

The palettized image is divided into small blocks and each block is quantized into two colours using a simple clustering algorithm. A K-means clustering algorithm [14] is used to find the two representative colours for each small block (Notice the clustering is performed on the colours of the block, not the colour indices. Details will be explained shortly). The original palette is used to determine the colour indices of the two locally quantized representative colours. A bitmap is used to identify the two different colours. The scheme is illustrated in Fig. 2.

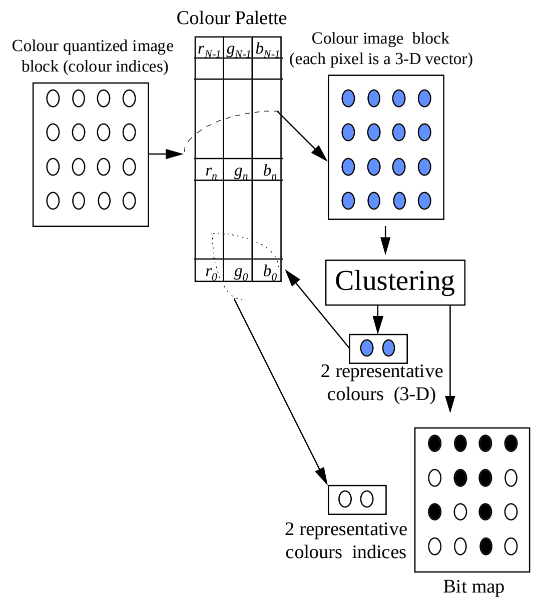

Figure 2 Local colour quantization of colour quantized image

Data clustering techniques, such as the K-means algorithm are normally computationally intensive and sub-optimal. The clustering task we have at hand is relatively simple because there are only a small number of low dimensional (3-D) samples (16 for 4 x 4 blocks), and only two clusters are to be produced. There are many different techniques to make the algorithm converge quickly [14], in our simulation, we have used the following strategy. First, the mean vector of the entire block is calculated. Second, two cluster centers are initialized by applying a very small amount of random noise to the mean vector. Third, K-means iteration process is applied until converge. It was found that it is only necessary to go through all the sample twice to obtain the final two representative colours. Because the samples are spatially close to each other, many of them might have the same or very similar colours after quantization, this may explain the fast convergence speed.

Assuming a 256 colour palette and a 4 x 4 block, each block is encoded by a 16-bit bitmap and two 8-bit colour indices. The bit rate is 2 bits per pixel (12:1 compression over true colour image). Although this bit rate is relative high when compared to source coding techniques such as JPEG [15], the present method achieves a further 4:1 compression over the original colour quantized (palettized) image. The merits of such an approach should be recognized by its simplicity, good visual image quality and moderately high data/buffer reduction rate.

This coding scheme can be easily implemented in a simple frame buffer architecture as shown in Fig. 3. To display a 4M x 4N colour image, the frame buffer uses one 4M x 4N x 1 bit memory to store the bitmap and two M x N x 8 bit memories to store the two colour indices of each 4 x 4 block. The bits from the bitmap memory control the switch which connects the correct colour index to the palette.

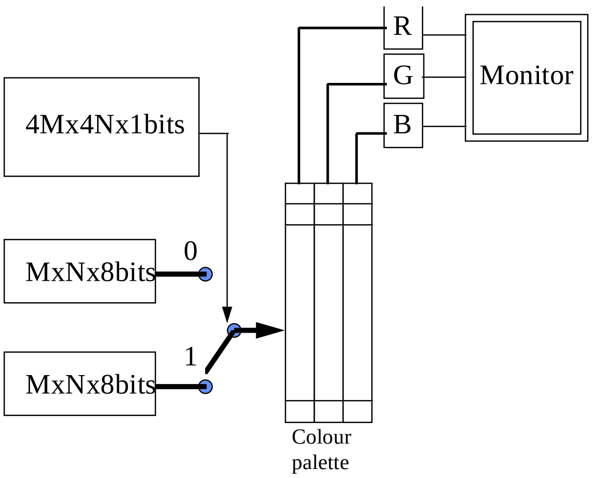

Figure 3 A simple video buffer architecture for displaying local colour quantized palettized colour image

4. Simulation Results

Computer simulations have been performed to evaluate the effectiveness of the technique. We subject tested the visual difference between palettized images (256 colours) and their locally colour quantized counterparts by pair comparison method. Two images were display simultaneously on an 18 inches monitor, we asked the subjects to view the screen at a distance of about 18 inches. We found that using a block size of 4 x 4 to locally colour quantize the images, no visually difference was detectable. For most of the natural images, the subjects could not detect any visual difference for a block size as large as 6 x 6 (corresponding to less than 1.5 bit per pixel).

(a)

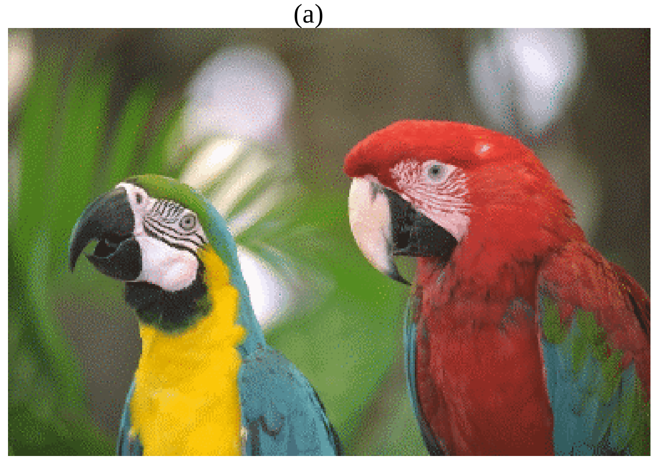

(b)

Figure 4 (a) Palettized Parrots image, 8 bits /pixel. (b) Local colour quantized image of (a), 2 bits / pixel

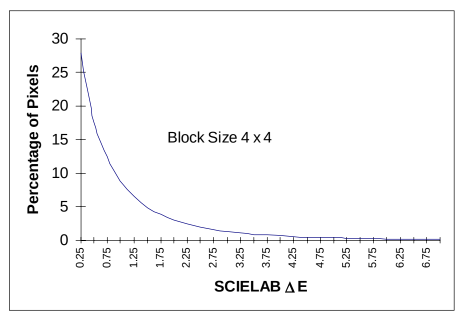

Figure 5, S-CIELAB ΔE distribution between colour quantized and local colour quantized Parrots image, average $\Delta E = 1.2$

The ultimate aim of the work is to achieve data reduction while maintaining the visual quality of the image. In this sense, subject testing is the best way to judge the effectiveness of the method. However, subject testing is susceptible to the influence of many uncertainty factors, in other word, the results can be unreliable. Recently, an objective colour image metric, the S-CIELAB system [8] was introduced to measure the visual difference between two colour images. Fig. 5 shows the S-CIELAB ΔE between a palettized image and its compressed version of the Parrots image (shown in Fig. 4) using a block size of 4 x 4. Fig. 6 shows the S-CIELAB ΔE for a block size of 6 x 6 and Fig. 7 shows the same for a block size of 8 x 8.

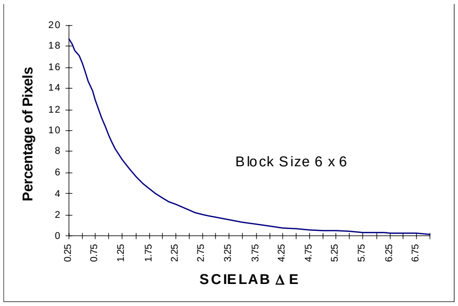

Figure 6, S-CIELAB ΔE distribution between colour quantized and local colour quantized Parrots image, average ΔE=1.6

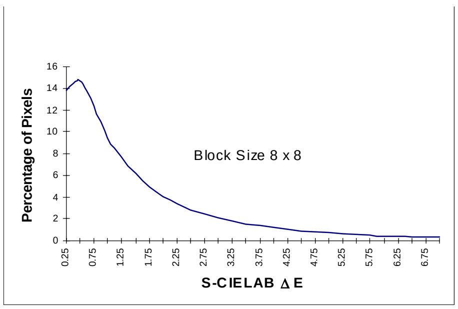

Figure 7, S-CIELAB ΔE distribution between colour quantized and local colour quantized Parrots image, average ΔE=1.8

5. Concluding Remarks

In this paper, we have introduced a technique for palettized colour image coding. The technique can reduce frame buffer requirements while maintaining the visual quality of the image. It should be noted that while the amount of compression achieved by this technique is still quite modest (typically 12:1 compression which is significantly better than most conventional colour quantization schemes), the merit of this work should be considered as achieving reduction in frame buffer size and maintaining the simplicity of hardware architecture.

Acknowledgment

I would like to thank Gerald Schaefer of the Colour and Imaging Institute, University of Derby for calculating the S-CIELAB ΔE shown in Figures 5 - 7.

References

1. D. Travis, Effective Color Displays: Theory and Practice. New York: Academic, 1991
2. M. Gervautz and W. Purgathofer, "A simple method for colour quantization: Octree quantization", Graphics Gems, New York: Academic, 1990
3. L. Kaufman, Perception, Oxford University Press, 1979
4. D. Chen and A. Bovik, "Visual pattern image coding", IEEE Trans. Communications, vol. 38, pp. 2137 - 2146, 1990
5. X. Wu, "YIQ vector quantization in a new color palette architecture", IEEE Trans. on Image Processing, vol. 5, pp. 321-329, 1996
6. A. Zaccarin and Bede Liu, "A new approach for coding color quantized images", IEEE Trans. on C&S on Video Technology, pp.442-453, 1993
7. M. Gormish, "Compression of palettized images by color", Proc. IEEE International Conference on Image Processing, pp. I274 -277, 1995
8. X. Zhang and B. A. Wandell, "A spatial extension of CIELAB for digital color image reproduction", SID 96 Digest, pp. 731 -735, 1996
9. P. Heckbert, "Color image quantization for frame buffer display", Computer Graphics, vol. 16, pp. 297-304, 1982
10. R. Balasubramanian and J. Allebach, "A new approach to palette selection for color images", Journal of Imaging Technology, vol. 17, pp. 284 - 290, 1991
11. E. J. Delp and O. R. Mitchell, "Image compression using block truncation coding", IEEE Trans. Communications, vol. COM-27, pp. 1335-1342, 1979
12. G. Qiu, M. R. Varley and T. J. Terrell, "Improved Block Truncation Coding using Hopfield Neural Network", IEE Electronics Letters, Vol. 27, No. 21, pp. 1924-1926, 1991
13. Y. Wu and D. C. Coll, "Single bit-map block truncation coding of color images", IEEE Journal on Selected Areas in Communications, vol. 10, pp. 952 - 959, 1992
14. A. Gersho and R. M. Gray, Vector Quantization and Signal Compression, Boston: Kluwer, 1992
15. W. B. Pennebaker and J. L. Mitchell, JPEG Still Image Compression Standard, New York: Van Nostrand Reinhold, 1992