MULTI-SCALE PROBABILISTIC DITHERING FOR SUPPRESSING BANDING ARTIFACTS IN DIGITAL IMAGES

Sitaram Bhagavathy, Joan Llach and *Jie fu Zhai*

Thomson Corporate Research
2 Independence way, Princeton, NJ 08540, USA.
{sitaram.bhagavathy, joan.llach, jiefu.zhai1}@thomson.net

ABSTRACT

A method is proposed for reducing the visibility of "banding artifacts," i.e. false contours resulting from color quantization in digital images. The method performs a multi-scale analysis on the neighborhood of each pixel, determines the presence and scale of banding artifacts, and probabilistically dithers (perturbs) the color of the pixel. The overall effect is to "break down" the false contours making them less visible. The proposed method may be used to reduce banding artifacts at the same bit depth as the input image or at higher bit depths. The banding detection mechanism ensures that artifact-free regions remain unaffected during the process.

Index Terms— Banding artifact removal, false contour removal, debanding, decontouring

1. INTRODUCTION

Depending on the application, digital images are represented at various bit depths, e.g. 8 bits per pixel (bpp), 10 bpp, and so on. A higher bit depth allows more colors to be represented and therefore a higher visual quality. Often however, the colors in digital images are quantized in order to reduce the bit depth, e.g. when a 10 bpp image needs to be shown on an 8 bpp display. In areas with smooth gradients, color quantization may produce "bands" each of which is mostly constant in color, with a small color difference between adjacent bands. Boundaries between such bands may be visible as false contours, also referred to as "banding artifacts."

There are two main approaches for reducing visible false contours as a result of bit depth reduction in digital images. The first one affects the images before or during the quantization process. This may involve adding noise to the image prior to quantization [1], diffusing the quantization error among neighboring pixels [2, 3], or a feedback-based quantization strategy [4]. The second approach affects the image *after* quantization. This becomes necessary in order to reduce the visibility of banding artifacts that have already appeared as a result of bit depth reduction. The proposed method belongs to the latter category of *post-quantization* methods.

Post-quantization methods take an input image at a bit depth N containing false contours and output an image at a bit depth M ($\geq N$) wherein false contours are visibly reduced. The two major problems to be tackled by post-quantization methods are the detection of regions containing banding artifacts, and the reduction of banding artifacts. Lee et al. [5] proposed a two-stage false contour detection algorithm that first eliminates smooth regions and then separates false contours from edges and texture using directional contrast features. The false contours are reduced by applying 1-D directional smoothing filters whose directions are orthogonal to that determined by the directional contrast features.

Ahn et al. [6] detect flat regions, or regions of low frequency content, since these have a higher likelihood of containing banding artifacts. A random shuffle is then applied to each pixel in the flat regions. This process exchanges the color of the pixel with that of a random pixel in its neighborhood. The downside of the shuffle is that small details in almost flat regions (e.g. stars in a sky region) may be spread out or lost. This is a possibility since regions that are mostly flat except for small details may be misdetected as flat regions.

Daly et al. [7] propose a "decontouring" method that predicts where false contours will occur by low pass filtering the input image (resulting in image A at bit depth $P > N$) and then quantizing the filtered image (resulting in image B at bit depth N). The difference between image A and image B (at bit depth P), containing the predicted contours, is then subtracted from the input image. This results in an output image (at bit depth $M = P$) with fewer false contours. Daly et al. try to remove false contours without adding noise or dither to the image. This is useful for noise-free images, such as images generated by computer graphics and line art with gradients. Daly et al. point out that *white* noise is ineffective in masking false contours at admissible noise levels. In this paper, we show that a low amplitude *signal-dependent* dithering is quite effective in masking false contours.

Smoothing filters can be used for reducing false contours only when the bit depth of the output is greater than that of the input (i.e. $M > N$), so that the output image can represent the intermediate colors created by the smoothing process. Therefore, the methods of Lee et al. [5] and Daly et al. [7] are not applicable when the output image has to be of the same bit depth as the input image (which is a requirement in some applications). Furthermore, the prior methods do not address the issue of the scale of banding artifacts. The scale of banding refers to the proximity between adjacent false contours. This issue is important since false contours which are closely bunched together will need to be handled differently from those which are far apart. In [5] and [7], closer contours will necessitate smoothing filters with smaller supports, and in [6], closer contours will necessitate a smaller neighborhood for the random shuffler. However, the scale of banding is not determined in any of the above methods.

The method we propose performs a multi-scale analysis on the neighborhood of each pixel, and determines the presence and scale of banding artifacts around that pixel. Thereafter, the color of each pixel is probabilistically dithered based on the distribution of colors in its neighborhood (of appropriate scale). The overall effect is to "break down" the false contours making them less visible. The proposed method has the following advantages:

- Effective suppression of banding artifacts of different scales;
- Banding artifacts can be reduced at the same bit depth as the input image or at higher bit depths;
- Low amplitude signal-dependent dithering reduces the visibility of false contours with a relatively low level of noise addition; and
- Preserves fine detail in the image, including small details occurring in almost flat regions.

2. BANDING ARTIFACT REDUCTION

The method comprises two main steps: multi-scale banding detection and banding reduction by probabilistic dithering. We shall now describe these steps in detail.

2.1. Multi-scale banding detection

With regard to reducing banding artifacts, it is in general desirable to alter only regions where such artifacts occur. Furthermore, banding artifacts may occur at various scales and estimating the scale helps in effectively reducing such artifacts. We propose a method for detecting the presence and scale of banding in the neighborhood of each pixel. The method involves a multi-scale analysis of the color distribution in the neighborhood of a pixel.

The input to the method is an image I. The symbol, $I(x, y)$, denotes the value or color of the pixel with coordinates (x, y). At each pixel location, we first estimate the most likely scale of banding around it and then decide whether or not significant banding is present at that scale. This process is as follows. Let us consider a number of scales, $s = 1, 2, \dots, S$, each of which corresponds to a neighborhood, $N_s(x, y)$, around the pixel (x, y). As the scale index s increases, so does the size of the neighborhood N_s.

At each scale s, a confidence score, $c(s)$, for the likelihood of banding is computed as follows.

$$c(s) = p(0, s) \times \text{MAX}\left[\frac{p(-1, s)}{p(0, s) + p(-1, s)}, \frac{p(1, s)}{p(0, s) + p(1, s)}\right], \quad (1)$$

where $p(k, s)$ refers to the fraction (or probability) of pixels in $N_s(x, y)$ having the value $I(x, y)+k$. The term $\text{MAX}[a, b]$ refers to the greater of a and b. The second term in the RHS of (1) gives a measure of the likelihood of banding being present. Since banding is the result of quantizing a smooth gradient, the pixels in adjacent bands differ by a value of 1. If $N_s(x, y)$ overlaps two or more bands, the second term in the RHS of (1) has a high value due to a relatively high number of pixels differing by 1 from the center pixel value. The first term in the RHS of (1) measures how significant the banding effect is. If the fraction of affected pixels, $p(0, s)$, is low, then it is likely to be visually insignificant.

It is desirable to avoid affecting regions with low likelihood of banding artifacts. To this end, before we choose the best scale of banding, we eliminate scales which have a low likelihood of representing bands. This is done by applying some criteria on the probabilities $p(k, s)$ where $k \in \{-1, 0, 1\}$. In our implementation, we *detect* the presence of banding at pixel (x, y) at scale s if

$$p(0, s) > T \text{ and } [p(-1, s) > T \text{ or } p(1, s) > T], \quad (2)$$

where T is a preset threshold. If none of the scales, s, obey the criteria in (2), we assume that the pixel (x, y) is not part of a banding artifact. If banding is detected at one or more scales at (x, y), we choose the scale, s^*, with the highest confidence score among all scales at which banding is detected, i.e. $c^* = c(s^*)$, as the most likely scale of banding. Probabilistic dithering is then applied to reduce the visibility of the artifact.

2.2. Banding reduction by probabilistic dithering

At each pixel where banding has been detected, a probabilistic dithering is applied in order to break down the false contours and reduce the visibility of banding. Our approach is superior to adding white noise since, from a distance, white noise is averaged out by the human visual system resulting in the reappearance of the banding pattern. The proposed method exploits the local properties of banding artifacts in order to effectively mask its appearance.

Let J be the output of the banding reduction process, i.e. the "debanded" image. Let (x, y) be a pixel location where banding of scale s^* has been detected. The output value, $J(x, y)$, is obtained by dithering or perturbing the input pixel value, $I(x,y)$, according to certain probabilities. We compute the expected mean value in the local neighborhood, $N_{s^*}(x,y)$, as follows:

$$m = p'(-1,s^*)(z-1) + p'(0,s^*)z + p'(1,s^*)(z+1), \quad (3)$$

where $p'(k,s^*) = p(k,s^*)/[p(-1,s^*)+p(0,s^*)+p(1,s^*)]$, and $z = I(x,y)$. The probabilities, $p(k,s^*)$, are defined in the previous section. Ideally, the output value, $J(x,y)$, should be equal to m. However, since a pixel can only take discrete values, a dithering strategy is devised as follows in order to make the neighborhood mean value approach m.

Let us define a *lower value* as $\lfloor m \rfloor$ and an *upper value* as $\lfloor m \rfloor + 1$, where $\lfloor m \rfloor$ refers to the largest integer not greater than m. Let us then define a probability $q = m - \lfloor m \rfloor$. In order to obtain a neighborhood mean close to m, the output value, $J(x,y)$ is assigned the upper value with probability q and the lower value with probability $(1-q)$. In practice, a uniform random number, $r \in [0,1]$, may be generated and the following rule may be used.

$$J(x,y) = \begin{cases} \lfloor m \rfloor + 1 & \text{if } r < q \\ \lfloor m \rfloor & \text{if } r \geq q \end{cases} \quad (4)$$

After applying probabilistic dithering at all pixel locations where banding is detected, the debanded image, J, may then be sent to the next stage, e.g. encoder, display. Note that the upper value is clipped at the maximum allowed value of $2^{b_I} - 1$ (where b_I is the bit depth of I) and the lower value is clipped at zero.

In the above described banding detection and reduction methods, the input and output may be grayscale images or one component of color images. In general, color images are represented by multiple components, e.g. YUV, RGB. Bit depth reduction of such images involves separately quantizing each component of the image. Conversely, in order to reduce banding artifacts, we may apply the proposed method separately to each component of the quantized image.

3. BIT DEPTH EXTENSION

Bit depth extension is the reverse process of bit depth reduction, wherein the bit depth of an image is increased, e.g. from 8 bpp to 10 bpp. The additional bits may be used effectively to further suppress the visibility of existing banding artifacts. The method proposed in the previous section may be modified to further reduce the visibility of banding artifacts during bit depth extension. The banding detection mechanism remains the same and is performed at the input bit depth. Let the bit depth of the input image, I, be b_I, and that of the output image, J, be b_J ($> b_I$). The smallest non-zero difference of 1 between pixels in the input now translates to a difference of $d_{\min} = 2^{b_J - b_I}$ in the output.

The probabilistic dithering method of Sec. 2.2 is modified thus. Consider the pixel at (x,y). If no banding is detected at (x,y), we set the output value $J(x,y) = d_{\min} I(x,y)$. If banding of scale s^* is detected at (x,y), the expected mean value in the local neighborhood, $N_{s^*}(x,y)$, is computed by modifying (3) as follows:

$$m = d_{\min}\left[p'(-1,s^*)(z-1) + p'(0,s^*)z + p'(1,s^*)(z+1)\right]. \quad (5)$$

Based on the above value of m, the upper and lower values, and the probability q are computed in the same way as described in the previous section. The dithering step is also identical to that described earlier. Since $b_J > b_I$, the upper and lower values are closer in color space at bit depth b_J than at bit depth b_I. Thus, at higher bit depths, the dithered result appears less noisy while still effectively reducing the visibility of banding artifacts.

4. EXPERIMENTAL RESULTS

In this section, we compare the proposed method of banding artifact reduction with the "decontouring" method proposed in Daly et al. [7], which we consider to be a prominent method in the recent literature. Fig. 1(a) shows a 6 bits-per-pixel (bpp) grayscale image of size 256×256 with a set of intensity bands of varying width or scale. Each band differs from its adjacent bands by one intensity level, thus simulating banding artifacts arising from a quantization process. The two white dots represent small details which should ideally be preserved during the banding artifact reduction process.

Fig. 1(b) shows the 8 bpp result after applying the method of Daly et al. Note that this method always outputs an image at a higher bit depth than that of the input image. The quality of the result depends on the scale of the low pass filter used. We experimented with filters of different scales and chose a 20×20 averaging filter, since it seems to be the most effective scale for this example. From Fig. 1(b), it can be observed that (i) although the method is effective in reducing existing false contours, it could introduce new ones; and (ii) the low pass filter could blur small details (the white dots).

Fig. 1(c) shows the result of applying our proposed method. The dithering is performed at the same bit depth as the input image, i.e. 6 bpp. In our method, banding artifacts can be reduced at the same bit depth as the input image or at higher bit depths. The former cannot be done using the method of Daly et al. From Fig. 1(c), it can be seen that the proposed method is effective in breaking down the contours between bands making the artifacts less visible. The multi-scale banding detection mechanism is effective in tackling bands of different sizes. It also ensures that artifact-free regions, including small details (the white dots), mostly remain unaffected during the dithering process. Note that, in our implementation, $N_s(x,y)$ (in Sec. 2.1) is a square neighborhood[1]

[1] At the borders of the image, $N_s(x,y)$ may not lie entirely inside the image. Ways of handling this problem include truncating the neighborhood to only the portion that lies inside the image, mirroring a band of border pixels, and ignoring a band of border pixels so that the neighborhood always stays inside the image.

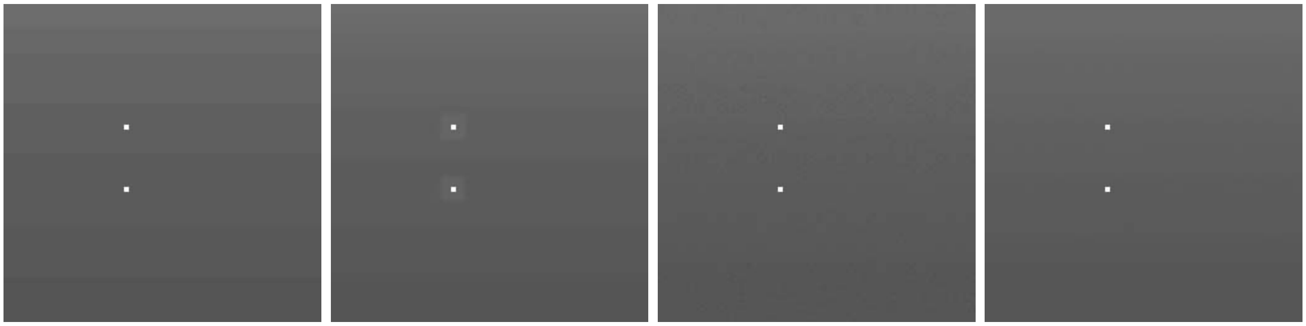

(a) (b) (c) (d)

Fig. 1. (a) Input 6 bpp image with bands and small details (white dots); (b) Result of the "decontouring" method of Daly et al. [7] at 8 bpp; (c) Result of the proposed method at 6 bpp; and (d) Result of the proposed method at 8 bpp. (Note that these images are best viewed in the electronic pdf version.)

centered at (x, y), and $T = 0.2$ in (2).

In applications where the bit depth of the output can be higher than that of the input, the bit depth extension method (Sec. 3) may be applied. This helps us to suppress banding artifacts while introducing less perceivable noise than at lower bit depths. Fig. 1(d) shows the result after extending the bit depth from 6 bpp to 8 bpp. The result at 8 bpp has less perceivable noise than the result at 6 bpp (see Fig. 1(c)).

Since banding artifacts arise from quantization of smooth gradients, the banding reduction process should result in the reintroduction of smooth gradients. We now compare the proposed method with that of Daly et al. in terms of the smoothness of the output intensity gradients. Fig. 2(a) plots the intensity average of each row of the 6 bpp input and 8 bpp output image after applying the proposed method. The two intermediate peaks in the curves are due to the white dots in Fig. 1(a). Note that the staircase profile of the input image is smoothed out effectively by applying the proposed method. In comparison, Fig. 2(b) shows the row averages for the method of Daly et al. The resulting image also has a staircase profile, although the "stairs" are of smaller scale than in the input.

5. CONCLUSION

In this paper, we have proposed a method for reducing the visibility of banding artifacts, i.e. false contours arising from color quantization in digital images. The method comprises two steps, multi-scale banding detection and probabilistic dithering. The former enables the effective handling of bands of various scales. It also ensures that artifact-free regions and small details remain unaffected during the dithering process. The probabilistic dithering step breaks down false contours, reducing the visibility of banding artifacts. This step can be performed at the same bit depth as the input or at higher bit depths. We have provided experimental results demonstrating the effectiveness of the proposed method and its advantages over a prominent previous method.

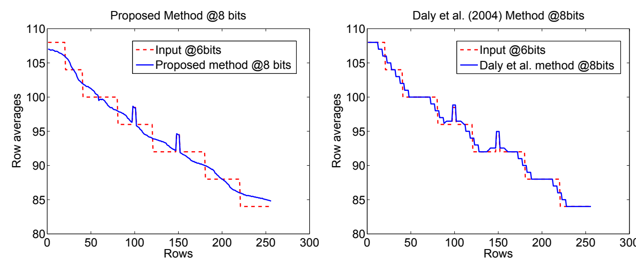

(a) (b)

Fig. 2. Row averages of intensities for (a) the proposed method; and (b) the method of Daly et al. [7].

References

[1] L. G. Roberts, "Picture coding using pseudo-random noise," *IRE Transactions on Information Theory*, vol. IT-8, pp. 145–154, Feb 1962.

[2] R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial grayscale," in *Proc. the Society for Information Display*, 1976, vol. 17, pp. 75–77.

[3] Victor Ostromoukhov, "A simple and efficient error-diffusion algorithm," in *Proc. SIGGRAPH 2001, in ACM Computer Graphics, Annual Conference Series*, 2001, pp. 567–572.

[4] G. Joy and Z. Xiang, "Reducing false contours in quantized color images," *Computers and Graphics*, vol. 20, no. 2, pp. 231–242, 1996.

[5] Ji Won Lee, Bo Ra Lim, Rae-Hong Park, Jae-Seung Kim, and Wonseok Ahn, "Two-stage false contour detection using directional contrast and its application to adaptive false contour reduction," *IEEE Transactions on Consumer Electronics*, vol. 52, no. 1, pp. 179–188, Feb 2006.

[6] Wonseok Ahn and Jae-Seung Kim, "Flat-region detection and false contour removal in the digital TV display," in *Proc. IEEE International Conference on Multimedia and Expo (ICME)*, 2005, pp. 1338–1341.

[7] S. J. Daly and X. Feng, "Decontouring: prevention and removal of false contour artifacts," in *Proc. SPIE International Society for Optical Engineering*, 2004, vol. 5292, pp. 130–149.