Colour Plane Synchronisation in Colour Error Diffusion

Donald G. Bailey
Physics Department, Massey University
Palmerston North, New Zealand
D.G.Bailey@massey.ac.nz

Abstract

On a bi-level colour device, only eight distinct colours (including black and white) are able to be represented. One common method of mapping image data to such a device is to use error diffusion techniques. These techniques diffuse the error from representing a particular pixel to adjacent unprocessed pixels, with the effect that locally the average pixel value closely approximates that of the original image. However, when applied to colour images, the pattern of dots within the three colour planes lose synchronisation, resulting in coloured patterns within regions of low saturation. This paper demonstrates that this problem may be overcome by modulating the threshold level by image intensity. The results presented show that a threshold deviation on 15% is sufficient to give the desired linear relationship between desynchronisation and saturation, and that the resynchronisation occurs within one or two pixels of an edge between a saturated and unsaturated region.

1. Introduction

A common problem in computer graphics is that of representing an image with a limited range of colours or grey levels, whether onto a display or printing the image. Within the computer, the image may be consist of many shades of grey, or colours. The output device, whether a display or printer, is usually more limited in the range of greys or colours available. This is especially the case if the output device is bi-level, that is it can only represent black and white, or for a colour device, black, red, green, blue, yellow, magenta, cyan, and white. The problem then becomes one of generating a reasonable likeness of the image with what is available.

Error diffusion [1] is a technique that is commonly used to map from the image pixels to the output representation. The idea is to threshold each pixel at a pre-defined level or levels, and pass any error from the representation to neighbouring unprocessed pixels. This error effectively biases the threshold level of the neighbouring pixels, resulting in the local average being preserved in the output representation, even on a bi-level output device. This technique adapts to edges, giving less blur than dither matrix techniques. The resulting pattern appears random, with few visual artefacts.

2. Colour Error Diffusion

When error diffusion is applied to colour images, for a binary colour device, the colour space may be represented as a cube, with the axes being red, green and blue for RGB devices, or yellow, magenta, and cyan for YMC devices. On a binary device, only the corners of the cube are able to be represented since a dot of the colour represented by each primary axis is either present or absent. The minimum error representation during the error diffusion process selects the nearest available colour, effectively partitioning the colour cube into 8 smaller cubes as the capture regions for each of the available colours (see figure 1). Since each partition depends only on one colour plane, this implies that each of the RGB or YMC planes may be processed independently.

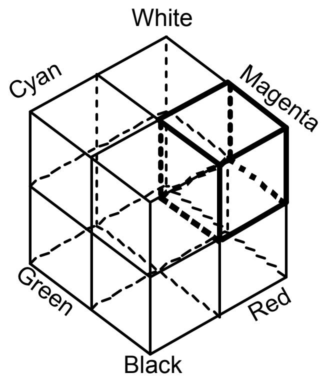

Figure 1: Minimum error division. The capture region for magenta is highlighted.

This process works well for images consisting only of highly saturated colours, that is colours which are close to the outside corners of the colour cube, and result in

0-8186-8183-7/97 $10.00 © 1997 IEEE

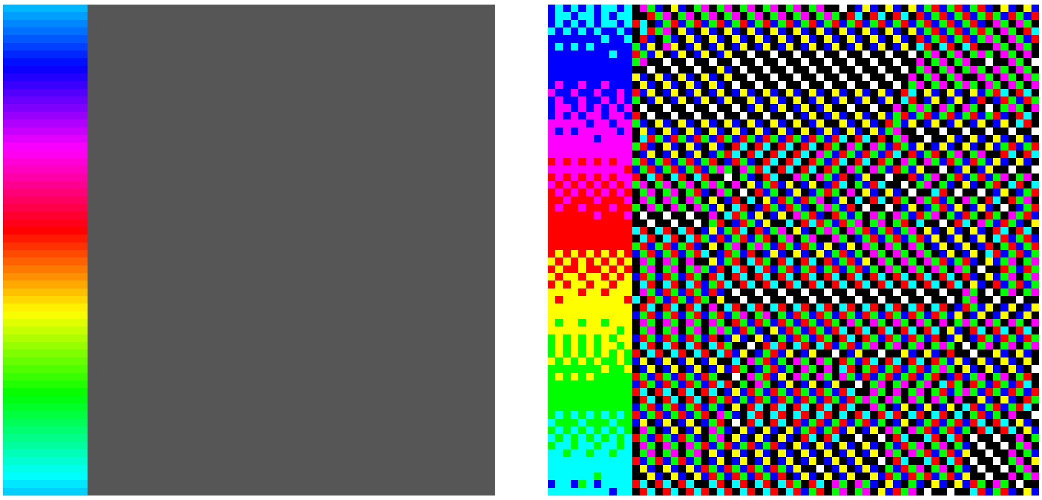

Figure 2: Loss of colour plane synchronisation with error diffusion.

only small errors. However with low saturation images, unless the image consists of pure greys, the errors propagated in each of the colour planes are not synchronised. The result is that instead of a large number of black and white pixels, most of the pixels are coloured. An extreme example of this is shown in figure 2, where a colour stripe is followed by a patch of grey.

To avoid a loss of synchronisation such as that illustrated in figure 2, it is desired to have the dots in the R G and B (or equally in the Y, M and C) planes to coincide as much as possible. This synchronisation between the planes would result in either black or white dots in pure grey regions within the image. In regions of low colour saturation, most of the dot should be either black or white, with a few coloured dots giving the grey a tint.

In general, there should be min(R,G,B) white dots, 1 - max(R,G,B) black dots, and max(R,G,B)-min(R,G,B) coloured dots. The difference max(R,G,B)-min(R,G,B) is closely related to the saturation of the image (saturation is this difference scaled by the intensity [2]). For the purposes of this study, I will define **saturation** as max(R,G,B)-min(R,G,B) and **desynchronisation** as the proportion of coloured pixels in the error diffused image. A test of the effectiveness of any colour resynchronisation algorithm would be how closely desynchronisation follows the saturation in the input image.

3. Resynchronisation method

The low saturation region of the colour cube corresponds to close to the axis between the black and white corners. The section of this line that is most critical is in the centre of the cube. Any small change in position caused by the propagation of errors can cause the pixel value to move from one colour subcube to another, resulting in a change in colour of that point.

Since regions close to the centre are unsaturated, they should be represented primarily by black and white dots in the output. Increasing the size of the black and white capture regions near the centre of the colour cube will reduce the number of pixels in this region being assigned a colour. A modification to the standard Floyd-Steinburg is to have a two-stage threshold process. In the first stage, the threshold is adjusted based on the intensity of the pixel. If the pixel is closer to white than black, the threshold is reduced by an amount ε to make the white region larger. Similarly, if the pixel is closer to black, the threshold is increased by ε to make the black region larger. In the second stage the red, green, and blue components are thresholded as before, with the error being diffused to the unprocessed pixels. This process is represented in pseudocode form in figure 3 (assuming RGB values are in the range 0.0 to 1.0):

```
If (R+G+B) > 1.5
    Threshold = 0.5 - ε
else
    Threshold = 0.5 + ε
If R < Threshold then Rout = 0.0 else Rout = 1.0
If G < Threshold then Gout = 0.0 else Gout = 1.0
If B < Threshold then Bout = 0.0 else Bout = 1.0
Rerror = R - Rout
Gerror = G - Gout
Berror = B - Bout
```

Figure 3: Modified thresholding method for colour error diffusion

It has been reported in the literature [3] that Floyd Steinburg error diffusion produces the correct grey level in monochrome images even if the threshold is different

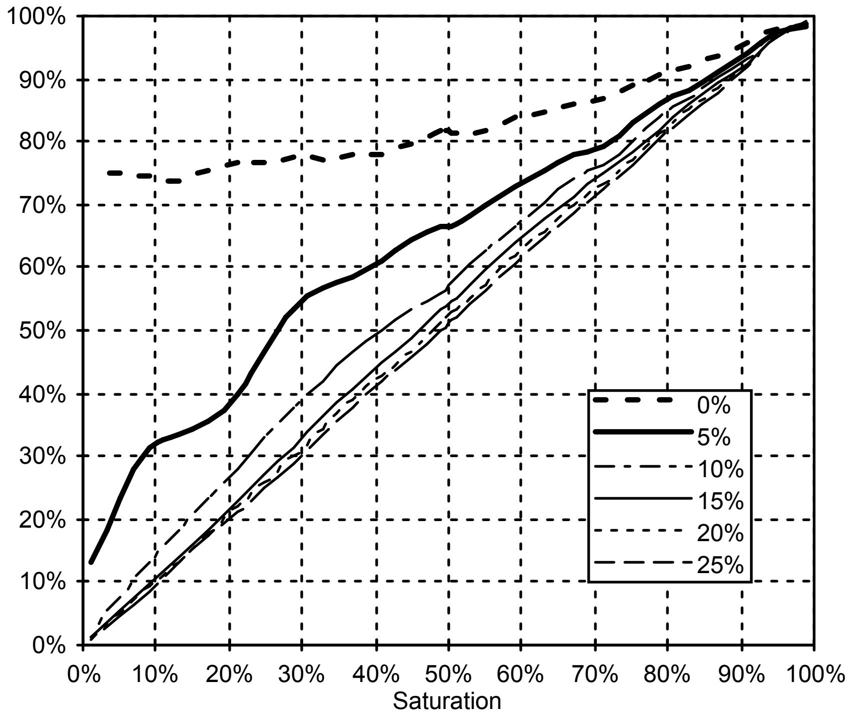

Figure 4: Relationship between desynchronisation and saturation for a range of ε.

from 0.5. This is because any errors resulting from the thresholding are passed on to subsequent pixels. The error soon builds to a level where it is self correcting, regardless of the actual threshold value. This is not to say that the threshold value has no effect on the image produced. Others [4] have modulated the threshold level based on the original pixel value to give edge enhancement.

The algorithm presented here modulates the threshold level based on the image intensity. The increased error introduced by the modulation is diffused to neighbouring pixels, where it is eventually averaged out. In the process, the colour planes become resynchronised.

4. Results

Two aspects of the modified error diffusion algorithm are of particular interest. The first is to determine the level of modulation is required to give good synchronisation. The second is to examine how quickly the colour planes resynchronise when an unsaturated region of an image follows a saturated region.

4.1 Modulation required for synchronisation

To test the relationship between desynchronisation and saturation, a test image was constructed where the hue varied vertically, the saturation varied horizontally from 0 to 100%, and the intensity was 50%. The hue variation ensures that the results are independent of colour. Error diffusion was applied to this test image for a range of threshold modulation levels and by counting the number of black and white pixels in each column, the desynchronisation is able to be determined as a function of saturation. The results are shown in figure 4. The experiment was repeated with the saturation decreasing from 100% to 0% across the image, with almost identical results except for modulation levels below 5%.

Even a small change in the threshold level has a significant effect on the synchronisation characteristic, especially for low saturation levels where it is the most important. As the threshold modulation is increased, the level of synchronisation approaches the ideal case. For threshold deviations greater than 15% ($\varepsilon > 0.15$) there is very little change in the relationship. This is important because it shows that the exact level for ε is not critical.

4.2 Resynchronisation distance

A second test was performed to determine how quickly the individual colour planes become synchronised. For this test, the colour error diffusion algorithm was applied to an image with a transition from a fully saturated colour hue bar to a region of mid grey (0 saturation). The number of pixels that were completely

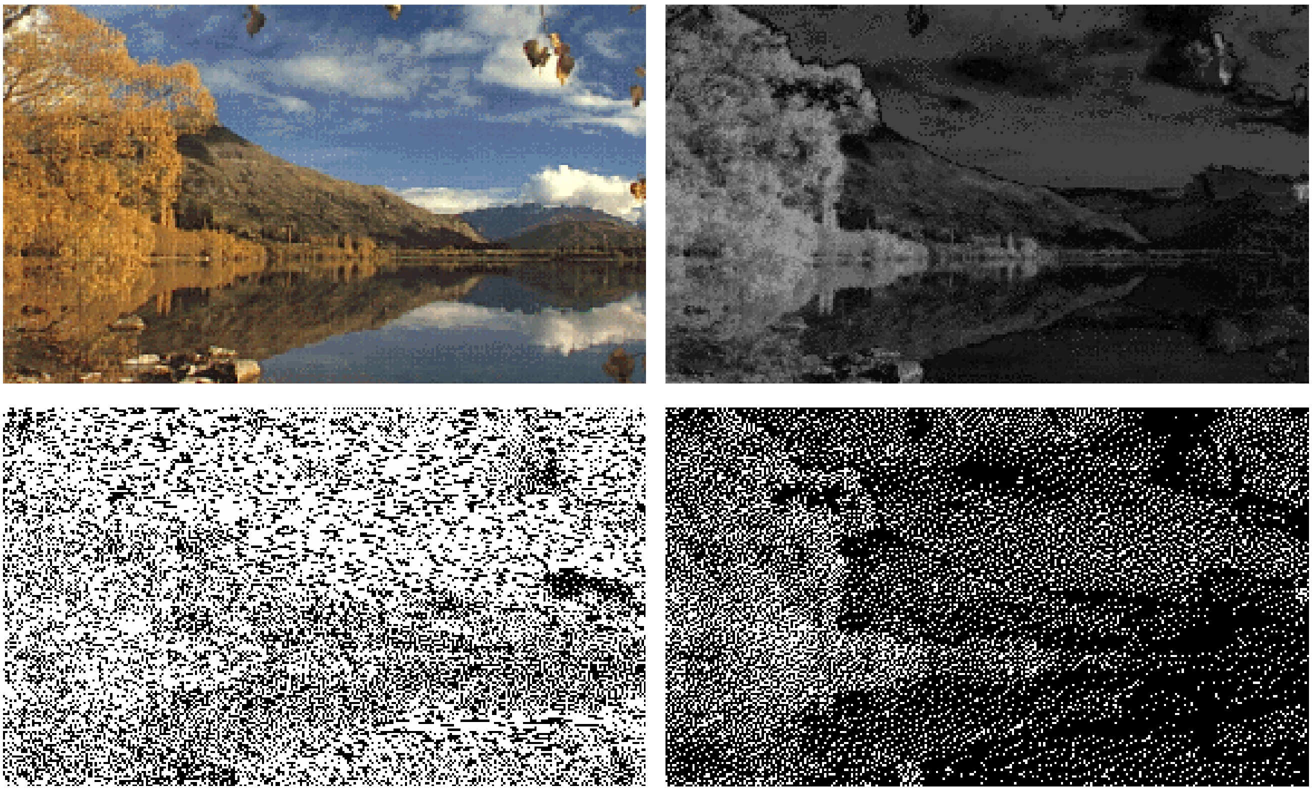

Figure 5: Effect on the saturation of a natural scene. Original image (top left), saturation of the original image (top right), saturation from standard error diffusion (bottom left), saturation from modified error diffusion with ε=15%(bottom right).

black or white were counted as a function of distance from the edge for a range of threshold modulation levels.

The algorithm is extremely effective at resynchronising the colour planes within one or two pixels when the threshold deviation is 15% or greater ($\varepsilon > 0.15$). Increasing ε further has little effect on the synchronisation, since the planes are already synchronised.

The modified error diffusion algorithm was applied to natural image to investigate the effect on saturation. The results are shown in figure 5. The original Floyd-Steinburg algorithm applied independently to each colour plane results in the colour being spread between the output pixels, increasing the overall average saturation at the pixel level.

On a pixel by pixel basis, the saturation in the output image will be either 1 or 0, depending on whether or not the pixel is coloured, or black and white. Since colour desynchronisation should be proportional to saturation, this implies that the average saturation of the image should ideally be preserved by the error diffusion process. This is indeed shown to be the case in figure 5 when the threshold deviation is 15%.

5. Conclusions

Error diffusion applied to each colour plane within colour images results in the pixels becoming desynchronised between the planes. The individual colour planes may be synchronised by adjusting the threshold level based on intensity, to increase the volume of the black and white capture regions within the colour cube.

A threshold modulation level of 15% was effective at resynchronising the individual colour planes. At this level, the planes became resynchronised within one pixel from an edge. The proposed algorithm is therefore extremely effective at preserving local saturation.

References

[1] Floyd R.W. and Steinburg L., "An Adaptive Algorithm for Spatial Grey Scale", *SID Int .Sym. Dig. Tech. papers*,, pp 36-37 1975.

[2] Foley J.D. and van Dam A., *Fundamentals of Interactive Computer Graphics*, Addison-Wesley, Reading, Massachusetts, 1982.

[3] Knuth D.E., "Digital Halftones by Dot Diffusion", *ACM Transactions on Graphics*, **6**, pp 245-273, 1987.

[4] Eschbach R. "Comparison of Error Diffusion Methods for Computer Generated Holograms", *Applied Optics*, **30**, pp 3702-3710, 1991.