Halftoning with Image-Based Dither Screens

Oleg Verevka John W. Buchanan †

Department of Computing Science
University of Alberta, Edmonton, Alberta, Canada
{oleg,juancho}@cs.ualberta.ca

† Currently at Electronic Arts Inc.
Burnaby, British Columbia, Canada
juancho@ea.com

Abstract

Continuous tone images must be halftoned to be displayed on binary output devices such as printers. The ordered dither algorithm is a popular approach to halftoning. This algorithm uses a threshold matrix to approximate gray scale values. The arrangement of thresholds in the matrix determines texture artifacts introduced into the halftoned image. Thus, the challenge of research in ordered dithering is to find a matrix that results in the least visible texture artifacts.

In this work we control the halftoning texture by generating a threshold matrix from an arbitrary image. We demonstrate that processing images using adaptive histogram equalization results in pixel distributions similar to traditional dither screens. Ordered dithering with the resulting threshold matrix enables us to define texture in the halftoned image. We control the appearance of this texture by a combination of the ordered dither algorithm with an error diffusion process.

We present applications of the image-based dither screens to both photorealistic and artistic rendering. In the case of photorealistic tone reproduction this technique preserves textures and edges of the original image. The ability to define an arbitrary texture enables us to introduce a variety of artistic effects. A halftoned image can be embossed with another image, texture, or text. Also, halftoning with textures clipped from the existing art works approximates the look of traditional illustration media.

Key words: Halftoning, texture processing, artistic screening, non-photorealistic rendering.

1 Introduction

Texture is an important channel of visual information. In realistic images (e.g. photographs) texture helps us to understand surface properties, directions, depth, etc. In addition to this geometric information, artists often use textures to convey other types of information such as importance, focus, completeness, etc. [13]. Thus, it is necessary to control texture in both photorealistic and non-photorealistic computer rendering. In this paper we investigate texture control in halftoned images. Our algorithm enables us to define a halftoning texture, and is thus applicable to photorealistic and artistic image generation.

Texture is an inherent artifact of the halftoning process. It is introduced into the image as a result of approximating gray tones by black and white dots. Reduction of halftoning texture artifacts is the primary goal of research in conventional tone reproduction with halftoning. Two problems persist:

1. The appearance of halftoning texture in image regions with smooth gradients;
2. Distortion of image features (edges and textures) by the halftoning texture.

In this work we do not attempt to hide the halftoning texture. Instead, we show how the halftoning texture can be controlled.

The ability to introduce and control texture in a halftoned image is closely related to the non-photorealistic rendering (NPR) area of computer graphics. A number of NPR techniques produce binary images Most of them are tailored to imitate traditional illustration media: line drawings [20], engravings [14], pen-and-ink illustrations [25, 18]. These systems approximate image tones by procedurally defined textures.

In this research we adapt traditional halftoning techniques to artistic rendering. A similar approach was taken by Ostromoukhov and Hersch [17] where the ordered dither algorithm was extended to allow a designer to define dither screens of a desired geometric shape. Buchanan [4] introduced various halftoning textures by altering parameters of the clustered error diffusion method. Unfortunately, these techniques generate a limited style of textures and texture effects are hard to control.

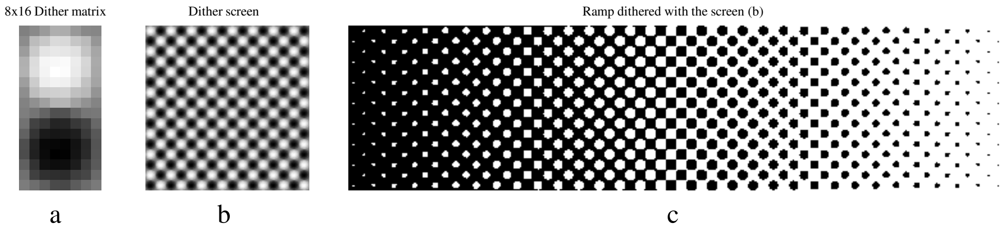

Figure 1: Dither matrix (a) is tiled to produce a dither screen (b). The arrangement of thresholds in the matrix controls the appearance of texture introduced into the halftoned image (c).

The primary goal of our research is to derive a halftoning texture from an arbitrary image. Our contribution is not limited to artistic rendering. Control of the introduced halftoning texture addresses the problem of texture distortions in photorealistic tone reproduction. A number of approaches were studied in the past. Jarvis *et al.* [10] suggested the use of a pre-processing filter for enhancing image edges. Adaptive error diffusion techniques that account for sharp image variations were investigated [21, 5]. Buchanan *et al.* [3] presented a post-processing algorithm that draws edges in the already halftoned image. While the use of edge enhancement techniques (e.g. [10, 5]) often results in improved texture rendering, the problem of texture reproduction has been rarely addressed. Thus, our work is a contribution in this area.

Our approach is based on a combination of ordered dithering with the error diffusion method. In the following section we present an overview of the relevant halftoning material. We discuss the main properties of a good dither screen: *uniform homogeneous threshold distribution*. We then show that adaptive histogram equalization produces dither screen with these properties. The appearance of background textures is controlled by an error diffusion technique. Further, we demonstrate applications of our algorithm to both photorealistic and artistic rendering.

2 Halftoning process

In this section we review the previous halftoning results that became a foundation of our method. These are: *error diffusion* and *ordered dithering* [22].

Error diffusion is a region based halftoning algorithm. The output value is determined by the input gray values, the threshold and the weighted errors diffused from the neighboring regions. The error associated with this approximation is propagated to un-processed neighbor regions, thus compensating for the approximating error over a region. Error diffusion algorithms differ in the size and shape of approximation regions and in the order of their traversal. Most techniques use single pixel regions with a variety of processing paths [7, 26, 22, 11, 16]. Others use multiple pixel regions to generate clustered dot halftoning [24, 5].

Unlike error diffusion algorithms, ordered dither is a point process. The value of the output pixel is determined by the threshold of the dither screen. Typically, a dither screen is constructed from a dither matrix tiled over the surface of an image (Figure 1a,b). The arrangement of threshold values within a dither matrix defines the halftoning texture. Since small size Bayer's [1] matrices are often used, the resulting images suffer from strong regular texture artifacts (Figure 1c). The goal of the recent research in ordered dither is to reduce periodic artifacts by generating large stochastic dither matrices: blue noise masks [15, 19], void-and-cluster screens [23].

We explore the ability of the ordered dither algorithm to define halftoning texture by the arrangement of thresholds in a dither screen. In the following section we discuss the properties of good dither screens, and we describe a method for converting an arbitrary gray scale image into a dither screen with these properties.

3 Conversion of a texture into a dither matrix

Arbitrary gray scale textures can be used as threshold matrices in the ordered dither algorithm. However, resulting halftoning is often a poor approximation of the continuous tone image. Consider the bark texture presented in Figure 2. When this gray scale image is used as a dither matrix the ordered dither fails to reproduce a simple ramp (black on the left, white on the right). Thus, the texture must be transformed to satisfy essential[1] threshold distribution properties of threshold matrices developed by the previous research [1, 15, 23]:

1. *Uniform distribution of threshold values.*
   A dither screen should contain the same number of

[1] Only threshold distribution properties that effect the accuracy of tone reproduction are mentioned here. Other properties including the ability to tile textures may also influence the appearence of halftoned images.

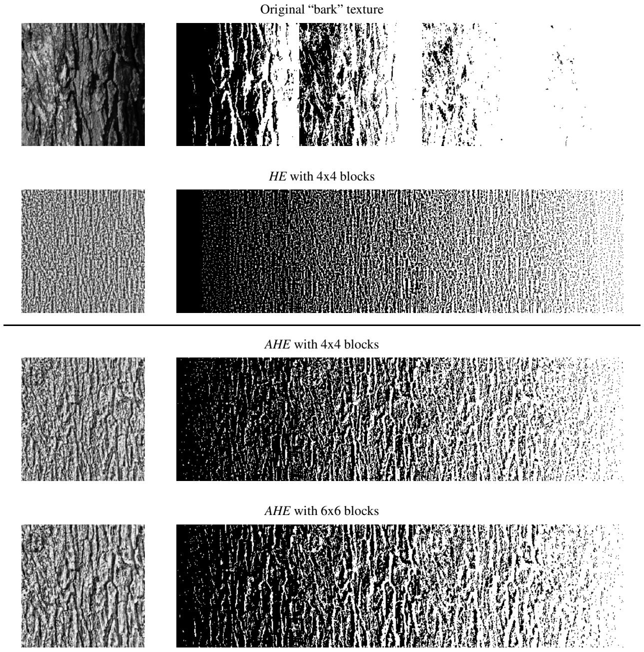

Figure 2: Halftoning ramp with bark texture dither screens (left column). Halftoning with unmodified texture fails to reproduce the gray ramp (top). Processing texture blocks with histogram equalization algorithm approximates uniform and homogeneous spatial distribution properties of threshold values and results in acceptable reproduction of the ramp. Unfortunately, some texture features are destroyed. The *AHE* algorithm preserves local characteristics of the input texture, while approximating the desired threshold distribution properties. The images are printed at 300 dpi.

pixels with the same values. In the 8x16 dither matrix (Figure 1a) all the values between 0 and 127 appear exactly once. This property enables the uniform[2] reproduction of the maximum range of gray tones.

2. *Homogeneous spatial distribution of threshold values.*
   Pixels with the same threshold values should be uniformly spread throughout the dither screen. Thus, gray tones are approximated in the same fashion in different regions of the image. The homogeneity property is satisfied automatically when a small dither matrix is tiled to generate a dither screen (Figure 1b).

[2]Non-uniform distribution of threshold values that accounts for the specifics of output hardware was suggested in previous research[19].

The task of uniform distribution of pixel values (property 1) has been addressed in image processing by the *histogram equalization* (*HE*) algorithm [9]. The *HE* transform is based on global cumulative distribution of image values $H(g)$. The gray-scale value g is transformed as follows: $\bar{g} = H(g)$.

In order to achieve the spatial homogeneity of threshold values (property 2), one can split an image into small blocks and apply the *HE* algorithm independently to each block. That is, the pixels of the block are transformed using cumulative distribution of values $h_b(g)$ within this block. Thus, each block becomes equivalent to a small dither matrix within a large dither screen. Since this dither screen has both desired properties outlined above the resulting halftoned images are also good approximations of the input gray scale values. Unfortunately, the use of histogram equalization leads to strong blocking artifacts and destroys some features of the input texture.

These limitations can be resolved by *adaptive histogram equalization* (*AHE*) approaches developed for local contrast enhancement. While many versions of the *AHE* exist [8], we have chosen the following algorithm:

1. The input texture is split into small blocks.
2. Cumulative distributions $h_b(g)$ are computed for each block.
3. Local distribution $h(g)$ for a pixel within a block is computed as a linear approximation between $h_b(g)$ of neighboring blocks.
4. The texture is transformed using local distributions: $\bar{g} = h(g)$

The resulting dither screen is free from blocking artifacts. Also, the desired properties of threshold value distributions are closely approximated. The size of blocks in the splitting process controls the size of texture features in the halftoned images (Figure 2). Large blocks correspond to strong texturing effects. We found that the block sizes between 3x3 and 8x8 produce the best results.

4 Texture control

We have presented an algorithm that enables us to introduce a desired texture into the halftoned image. In this section we discuss a method to control contrast and dot clustering of this texture.

The main advantage of ordered dither halftoning methods is that the distribution of the pixels is defined by a dither matrix. Unfortunately when a bad approximation for a pixel is generated there is no way to compensate for the error. This means that the structure of the dither screen is clearly visible. Error diffusion techniques, on

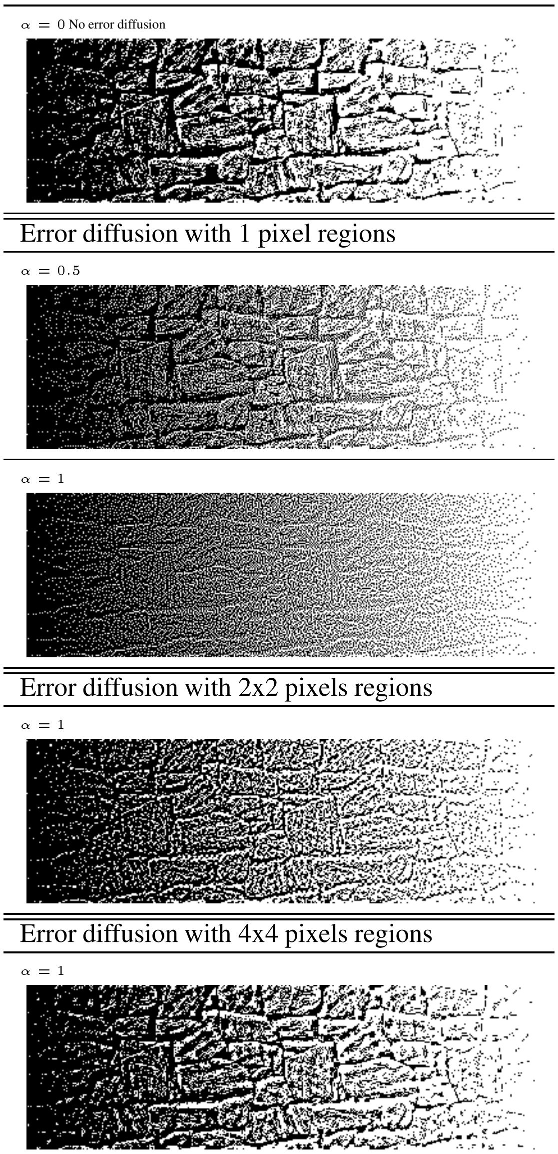

Figure 3: A combination of ordered dither with parameterized error diffusion allows as to control the appearance of the halftoning texture. By increasing the amount of errors diffused into a neighboring region (parameter α) we improve the accuracy of tone reproduction and reduce texturing effects. Dot clustering and improved preservation of texture features for large α is achieved by processing image by pixel blocks. Test images are printed at 150 dpi.

the other hand, attempt to compensate for quantization errors by propagating the error to unprocessed neighbor regions. The advantages of using error diffusion in combination with ordered dither were studied before [2, 22]. This combination not only improves the accuracy of tone reproduction but also reduces texture artifacts of the ordered dither.

We use a parameterized error diffusion [6] method to control the contrast of the halftoning texture. The parameter α where $\alpha \in [0, 1]$ determines the influence of approximation errors diffused from the neighbor regions on the output value of the current pixel. The input pixel g is set to 0 or 1 using the following formula:

$$b = \begin{cases} 0 & \text{if } g + \alpha E \leq t \\ 1 & \text{otherwise} \end{cases} \qquad (1)$$

where E is the sum of errors diffused into the current pixel and t is the corresponding threshold value from the dither screen. While any error diffusion technique can be used in this application, we have implemented Floyd-Steinberg algorithm [7] with serpentine processing of rows. Halftoning results for various values of the α parameter are presented in Figure 3. Maximal error diffusion corresponds to $\alpha = 1$ and produces the softest textures. We found that in most cases $\alpha \in (0.2, 0.8)$ is a good choice for this parameter.

This error diffusion algorithm can be generalized to process images in pixel blocks. In this case a sum of quantization errors is computed over a block of pixels and is distributed to the neighbor blocks. The output pixel is determined using the accumulated error E weighted by the number of pixels in the processing block. It turns out that processing images in a block fashion leads to clustering of dots in the resulting halftoning. Dot clustering is often desired to compensate for dot overlap in laser printers. Also, we found that error diffusion over pixel blocks preserves small texture features for large values of α (Figure 3). However, the use of large (e.g. 4x4 blocks) in the error diffusion process may introduce block artifacts into the halftoning texture.

To summarize, we demonstrated how an arbitrary image-based texture can be used as a dither matrix in halftoning. A combination of ordered dither with an error diffusion process enables us to control the appearance of this texture in the resulted image. In the rest of the paper we present applications of our technique to photorealistic tone reproduction and artistic rendering.

5 Photorealistic tone reproduction

In this section we apply texture control techniques to photorealistic halftoning of images. The focus of this investigation is to improve rendering of the original image

Non-clustered halftoning

a: Floyd Steinberg

b: Floyd Steinberg with edge enhancement

c: Partial error diffusion $\alpha = 0.7$

Halftoning with image-based dither screen combined with error diffusion $\alpha = 1$.

d: HE with 4x4 regions

e: CLAHE with 4x4 regions

Floyd-Steinberg [7] error diffusion “washes out” edges and textures of the original image (image **a**). Image enhancement (Jarvis *et al.* [10]) reveals image textures at the expense of shifts in contrast and over emphasized sharp edges (image **b**). Partial error diffusion increases local contrast but does not enhance the display of image textures (imge **c**). Halftoning with image-based dither screen results in a balanced rendering of textures and edges (image **d**). However, halftoning with dither screens produced by the AHE algorithm suffers from noise amplification in uniform areas of the image. We control noise artifacts by generating dither screens with the CLAHE algorithm (image **e**). The images are printed at 150 dpi.

Clustered halftoning

f: Clustered 3x3 dither

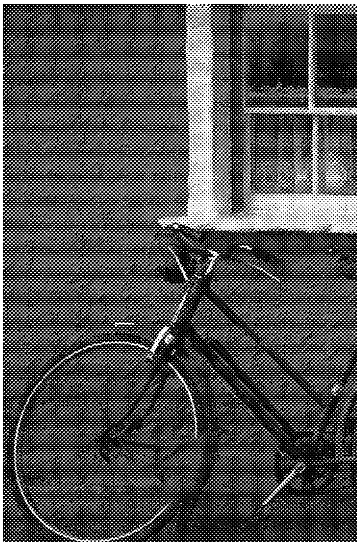

g: Image-based dither screening

Original image fragment

Image-based dither screening combined with error diffusion over pixel blocks results in clustered halftoning. Unlike conventional clustered dither (image **f**) our technique (image **g**) aligns dot clusters with image features thus highlighting edges and textures. These images are printed at 300 dpi.

Figure 4: Application of image-based screening to photorealistic tone reproduction. Dither screen is generated from the reproduced image.

textures. Since the ability to set isolated pixels differs between the output hardware we address halftoning with both dispersed and clustered dot allocation.

Floyd-Steinberg error diffusion [7] is often used when isolated pixels can be set independently. This algorithm creates good tone approximation and does not introduce strong regular artifacts. However, this error diffusion scheme "washes out" edges and textures of the original image (Figure 4a). Pre-processing step presented by Jarvis *et al.* [10] reveals edges and textures in the error diffused image (Figure 4b). This image enhancement approach is a compromise between over exaggeration of strong edges and rendering of textures. Moreover, strong image enhancement leads to noise amplification in the smooth shaded areas.

Our approach is to use a dither screen generated from the image to be reproduced. Thus, by a combination of the error diffusion with the ordered dither algorithm we guarantee that original image textures and edges are preserved in the halftoned image (Figure 4d). When compared to image enhancement results, our technique also preserves the contrast of the original image. Unfortunately, the AHE algorithm suffers from strong noise artifacts. We address this problem by the use of the contrast limited AHE (CLAHE) algorithm suggested by Zuiderveld [27]. The CLAHE algorithm limits histogram equalization in homogeneous image regions. Halftoning with CLAHE generated dither screens preserves edges and textures without noise amplification in the smooth areas of the image (Figure 4e). Therefore, our technique is a feasible alternative to image enhancement [10] in the halftoning applications that do not require dot clustering.

Unfortunately, many hard copy devices (e.g. laser printers) suffer from dot gain — bleeding of inks/toner into neighboring regions. Dot clustering partially compensates for this flaw. Halftoning with clustered dither screens is a compromise between good tone approximation and the introduction of regular texture artifacts that destroy small image details (Figure 4f). Halftoning with image-based dither matrix and error diffusion over pixel blocks produces clustered dot dither and preserves image textures (Figure 4g). Also, this algorithm aligns dot clusters along image edges and reveals small image details (e.g. spikes on the wheels, curtain ornaments).

Overall, our texture control techniques allow us to enhance reproduction of image features. We generate a dither screen from the reproduced image and is thus guarantee the appearance of its edges and textures in the resulting halftoning. This approach is a departure from traditional ordered dithering. Unlike previous research, we do not attempt to find a dither matrix that suits halftoning of all images. Instead we generate a dither screen that improves reproduction of a particular image.

6 Artistic screening

In the previous section we used image-based dither screens to enhance reproduction of textures and edges. The goal of our research here is to enhance halftoned image in a non-photorealistic style. In particular, we use dither screens to introduce a desired texture into the halftoned image. The halftoning texture in this case may carry additional information in the form of another image or text. Therefore, for the purposes of artistic rendering dither screens can be generated from a variety of sources including a photograph, a texture sample, or scanned text (Figure 5 left).

The need for multiple textures in the context of non-photorealistic rendering was discussed by Winkenbach and Salesin [25]. Image-based dither matrix approach allows us to introduce similar effects with *texture fusion*. The goal of texture fusion is to vary textures across the dynamic range of the image. We merge textures T_1 and T_2 into one matrix using the following formula:

$$T = \begin{cases} T_1 & \text{if } T_1 \leq t \\ T_2 & \text{otherwise} \end{cases} \quad (2)$$

where $0 < t < 1$ is the threshold intensity at which only one texture is visible. We found that merging multiple textures is essential for approximation of many traditional illustration techniques. For example painterly effects are achieved by combining oil brush strokes and canvas texture. Cross hatching can be simulated by merging two fragments of a scanned pencil art work (Figure 5 right).

Overall, image-based dither matrix approach has a number of advantages over the previous NPR research:

1. Our approach is not limited to any particular style or art medium.

2. In other NPR techniques an artist must design the elements of the halftoning layer (procedural texture in pen-and-ink [25]; screen elements in artistic screening [17]). In our approach the texture does not need to be designed by an artist, but can be extracted from an arbitrary image.

3. By varying textures we can reproduce a number of effects generated by the previous techniques: cross hatching and multiple textures of pen-and-ink [25]; texturing with text in artistic screening by Ostromoukhov and Hersch [17].

4. Control of texture strength and accuracy of tone reproduction with the error diffusion method was not investigated before in the NPR literature.

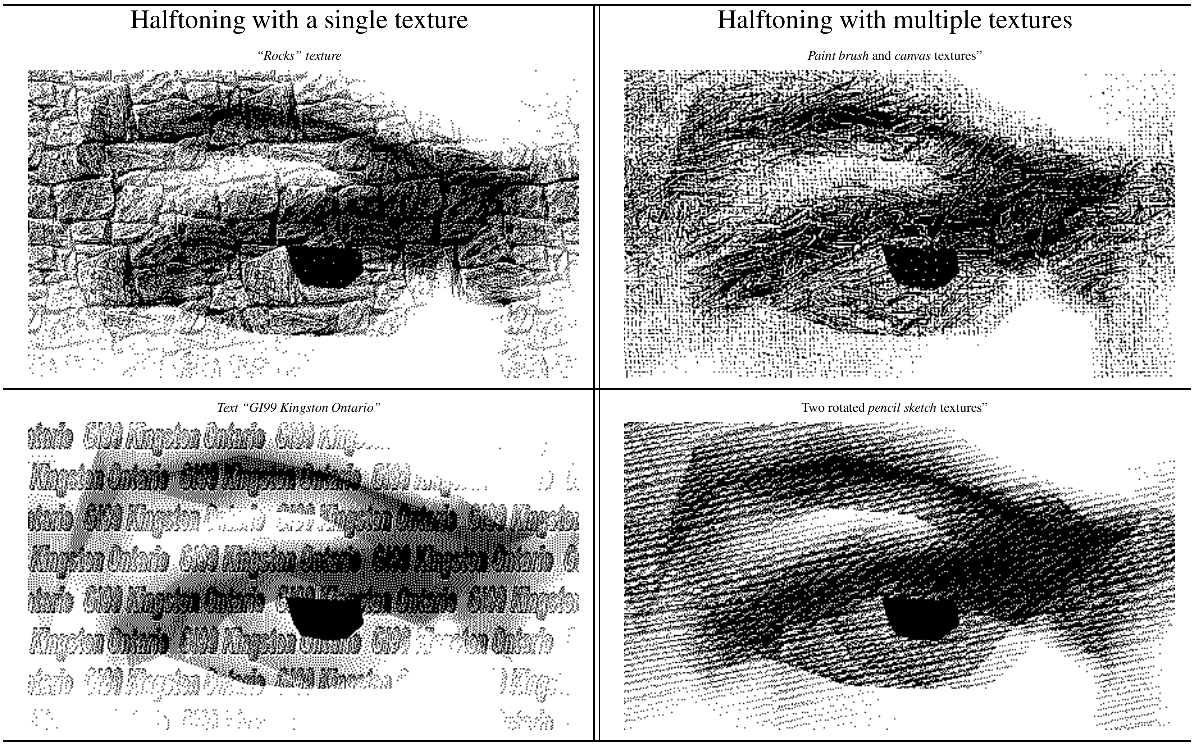

Figure 5: Non-photorealistic effects are introduced by a variation of dither screens. The possible sources for a dither screen are: a complete image, a texture sample, or scanned text. Also a dither screen may contain more than one texture. In the images on the right textures vary for the dark and light parts of the gray scale range. Thus, cross hatching and artistic media effects are introduced. These images are printed at 150 dpi.

The previously developed non-photorealistic rendering techniques have a number of advantages over our algorithm. In particular, both artistic screening [17] and pen-and-ink illustrations [25] allow fine control over the shape of textures. Whereas textures in the image-based dither screens are dependent on the noise in the source image. Also, the procedural nature of pen-and-ink textures enables designers to fine tune individual pen strokes. Such control is impossible in the image-based textures we use.

7 Conclusion

The display of gray scale images on binary devices requires a halftoning step such as ordered dithering. The challenge of halftoning research is to find a dither matrix that results in the least visible artifacts. We took a different approach and developed techniques to control texture defined by an arbitrary image.

Our approach is a combination of the ordered dither with the parameterized error diffusion method. We have shown that an image processed by the adaptive histogram equalization algorithm has approximately the same statistical distribution of values as traditional dither screens. Therefore, we use AHE processed images in the ordered dithering to introduce a desired texture. The strength of this texture and its clustering properties are controlled by parameters of the error diffusion method.

Our approach to control of the halftoning texture is a contribution to both photorealistic and non-photorealistic rendering. Halftoning with image derived dither screens preserves edges and textures of the original image. The ability to produce an arbitrary halftoning texture lands itself to artistic rendering. We have illustrated how dither screens generated from a variety of sources can enhance image presentation and approximate the look of traditional illustration styles.

8 References

[1] B. E. Bayer. An optimum method for two-level rendition of continious tone pictures. *IEEE International Conference on Communications*, 1:26–11–26–15, June 1973.

[2] C. Billotet-Hoffman and O. Bryngdahl. On the error diffusion technique for electronic halftoning. *Proceedings of the Society for Information Display*,

24(3):253–258, 1983.

[3] J. W. Buchanan, Lisa M. Streit, and O. Versyovka. Edge enhancement issues in half-toning. In *Graphics Interface '98*, pages 209–216, June 1998.

[4] J. W. Buchanan. Special effects with half-toning. In *Proceedings of Eurographics 96*, pages 97–108, 1996.

[5] J. W. Buchanan and O. Verevka. Edge preservation with space-filling curve half-toning. In *Graphics Interface '95*, pages 75–82, May 1995.

[6] Reiner Eschbach and Keith T. Knox. Error-diffusion algorithm with edge enhancement. *Journal of the Optical Society of America A*, 8(12):1844–1850, December 1991.

[7] Robert W. Floyd and Louis Steinberg. An adaptive algorithm for spatial greyscale. *Proceedings of the Society for Information Display*, 17(2):75–77, 1976.

[8] J. M. Gauch. Investigations of image contrast space defined by variations on histogram equalization. *CVGIP: Graphical Models and Image Processing*, 4(54):269–280, July 1992.

[9] A. K. Jain. *Fundamentals of Digital Image Processing*. Prentice Hall, 1989.

[10] J.F. Jarvis, C.N. Judice, and W.H. Ninke. A survey of techniques for the display of continuous tone pictures on bilevel displays. *Computer Graphics and Image Processing*, 5:13–40, March 1976.

[11] Donald E. Knuth. Digital halftones by dot diffusion. *ACM Transactions on Graphics*, 6(4):245–273, 1987.

[12] W. Leister. Computer generated copper plates. *Computer Graphics Forum*, 13(1), 1994.

[13] J. Landsdown and S. Schofield. Expressive rendering: A review of nonphotorealistic techniques. *Computer Graphics and Applications*, 15(3):29–37, 1995.

[14] W. Leister. Computer generated copper plates. *Computer Graphics Forum*, 13(1), 1994.

[15] Theophano Mitsa and Kevin J. Parker. Digital halftoning using a blue noise mask. In *ICASSP 91: 1991 International Conference on Acoustics, Speech, and Signal Processing*, volume 2, pages 2809–2812, Toronto, Canada, May 1991. IEEE.

[16] Avi C. Naiman and David T. W. Lam. Error diffusion: Wavefront traversal and contrast considerations. In *Graphics Interface '96*, pages 78–86. May 1996.

[17] Victor Ostromoukhov and Roger D. Hersch. Artistic screening. In *SIGGRAPH 95 Conference Proceedings*, pages 219–228. Addison Wesley, August 1995.

[18] Michael P. Salisbury, Sean E. Anderson, Ronen Barzel, and David H. Salesin. Interactive pen–and–ink illustration. In *Proceedings of SIGGRAPH '94*, pages 101–108. July 1994.

[19] Mark A. Schulze and Thrasyvoulo N. Pappas. Blue noise and model-based halftoning. In *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display V*, volume 2179, pages 182–194, February 1994.

[20] T. Strothotte, B. Preim, A. Raab, J. Schumann, and D. R. Forsey. How to render frames and influence people. *Computer Graphics Forum*, 13(3):455–466, 1994. Eurographics '94 Conference issue.

[21] Stefan Thurnhofer and Sanjit K. Mitra. Nonlinear detail enhancement of error-diffused images. In *Proceedings, SPIE: Human Vision, Visual Processing, and Digital Display V*, volume 2179, February 1994.

[22] Robert Ulichney. *Digital Halftoning*. MIT Press, 1987.

[23] Robert Ulichney. The void-and-cluster method for dither array generation. In *Proceedings, SPIE: Human Vision, Visual Processing, and Digital Display IV*, volume 1913, pages 332–343, February 1993.

[24] Luiz Velho and Jonas de Miranda Gomes. Digital halftoning with space filling curves. In Thomas W. Sederberg, editor, *Computer Graphics (SIGGRAPH '91 Proceedings)*, volume 25, pages 81–90, July 1991.

[25] Georges Winkenbach and David H. Salesin. Computer–generated pen–and–ink illustration. In *Proceedings of SIGGRAPH '94*, pages 91–100, July 1994.

[26] I.H. Witten and R.M. Neal. Using peano curves for bilevel display of continuous-tone images. *IEEE Computer Graphics and Applications*, 2(3):47–52, May 1982.

[27] Karel Zuiderveld. Contrast limited adaptive histogram equalization. In Paul Heckbert, editor, *Graphics Gems IV*, pages 474–485. Academic Press, Boston, 1994.