Supercell dither masks with constrained blue noise interpolation

Chai Wah Wu, Charles P. Tresser*, Gerhard R. Thompson, Mikel J. Stanich**
IBM T. J. Watson Research Center, Yorktown Heights, NY, USA
****IBM Global Financial Markets Research Center, Hawthorne, NY, USA***
*****IBM Printing Systems Division, Boulder, CO, USA***

Abstract

When halftoning images using dither halftoning algorithms, traditional clustered dot dither algorithms trade off cluster cell size versus the number of gray levels. Supercell clustered dot dither algorithms are used to obtain a small cluster cell size while allowing a large number of gray levels to be rendered. However, commonly used supercell clustered dot dither algorithms have unpleasant patterns at certain gray levels. We present an algorithm to use constrained blue noise to correct such unpleasant patterns. We also illustrate how this algorithm can generate nonlinear dither masks which can compensate for dot gain without the use of a separate tone reproduction curve.

1. Introduction

In traditional clustered dot dither halftoning methods, there is a tradeoff between cluster cell size and the number of gray levels in the dither mask; the smaller the cell size, the smaller the number of gray levels [1]. For example, a 2 by 4 cluster cell can accommodate 9 gray levels (Fig. 1).

16	48	239	207
80	112	144	175

2 x 4 halftoning cell

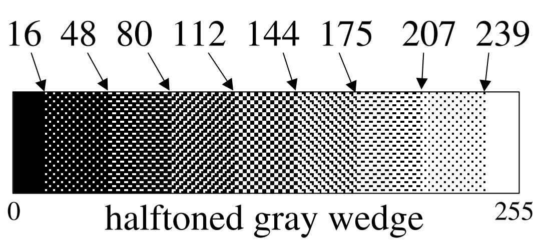

Figure 1: Dither pattern generated by a 2x4 cell. A gray wedge is rendered into 9 distinct halftone patterns.

A small cluster cell generates halftoned images with sharper edges and finer details, whereas a large number of gray levels prevents contouring and posterization. To allow a large number of gray levels while maintaining a small cluster size, supercell techniques are used, in which several halftone cells are grown out of synch to obtain many gray levels. However, commonly used supercell clustered dot dither generally have unpleasant patterns at certain gray levels. For instance, in Fig. 2 we show a pattern generated by a commonly used supercell clustered dot dither mask with a 4x8 cell size. One can see a “grid” pattern which can be unpleasant and objectionable in the halftoned image.

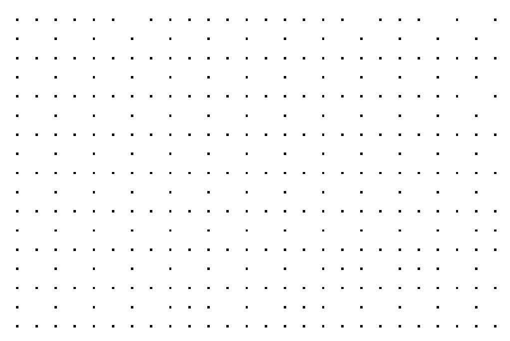

Figure 2: Dither pattern generated by a supercell dither with 4x8 cells.

In this paper, we present a method to generate a supercell clustered dither without such unpleasant patterns [2, 3]. The proposed method uses constrained blue noise to replace the unpleasant patterns with more pleasant patterns. The method can incorporate HVS (human visual system) based weighting functions. We also show how nonlinear dither masks can be created which can be used to create calibrated halftones without the use of a tone reproduction curve.

2. Halftoning and dither

Most digital printers today can print in only a limited number of colors. Digital halftoning is a technique for printing a picture using small dots with a limited number of colors such that it appears to consist of many colors when viewed from a proper distance. For example, a picture of black and white dots can appear to display various shades of gray when viewed from some distance. For simplicity, we will only discuss binary or monochrome printers in this paper where at any pixel location, the printer can either print a black dot of ink or toner or nothing at all.

The fastest and most commonly used methods for digital halftoning in laser printers are *dithering* algorithms which use threshold arrays, also called dither matrices or dither masks. A dither mask M is considered as a matrix of numbers of size $m \times n$. Each input image pixel $I(i,j)$ is compared with $M(i \bmod m, j \bmod n)$. If $I(i,j)$ is less than $M(i \bmod m, j \bmod n)$, then the output pixel $O(i,j) = 1$ and a black dot is printed, otherwise $O(i,j) = 0$ and no dot is printed (Fig. 3).

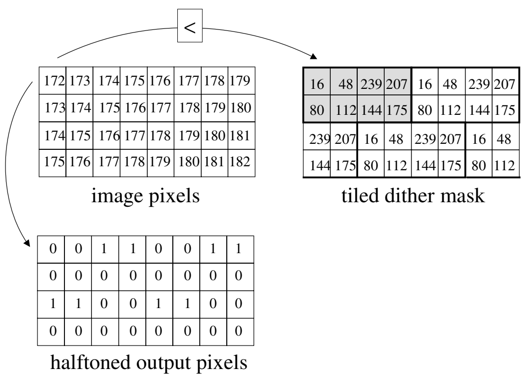

Figure 3: Dither algorithm.

The goal is to create a dither mask such that natural images result in nice halftone images. This is a difficult criteria to achieve. A more tractable design goal is to create a dither mask such that for each gray level $0 \le g \le 1$, the corresponding image of uniform gray ($I(i,j) = g$ for all i, j) is rendered into a nice halftone pattern O_g by the dither mask[1]. In particular, for a dither mask M the halftone pattern O_g can be defined as follows: O_g has a black dot in the (i,j)-th position if and only if $M(i,j) > g$. We denote this set of halftone patterns by $H = \{O_g\}$. Thus one way to approach the dither mask design problem is to define this set of halftone patterns H. The number of black pixels in the halftone pattern O_g is proportional to the corresponding gray level g if a linear tone reproduction curve is assumed. In Section 9.2 we discuss the case when this relation is nonlinear. Because of the way the dither algorithm is defined, the patterns in H satisfy the *stacking condition*: the pattern of black pixels of O_{g_1} for a lighter gray level g_1 is a subset of the pattern of black pixels of O_{g_2} for a darker gray level $g_2 \le g_1$. It is clear that the stacking condition implies that a dither matrix of size $m \times n$ can render at most $mn + 1$ distinct gray levels.

In clustered dot dither, the patterns form ever growing clusters of dots. These types of dither patterns are generally used on laser printers as they are more easily printed by the xerographic process. Furthermore, clustered dots also minimize the effects of dot gain.

3. Supercell dither masks

To increase the number of gray levels that can be rendered while preserving the cluster size, a supercell clustered mask can be constructed as follows. A collection of dither masks (each containing one or more halftone cells) are tiled to form a larger dither mask. The dots in these cells are grown asynchronously generating many gray levels while preserving the resolution of a single cell (Fig. 4). In many supercell clustered masks used, the asynchronous behavior among these cells generates unpleasant patterns (Fig. 2). The purpose of this paper is to present an algorithm for removing such unpleasant patterns.

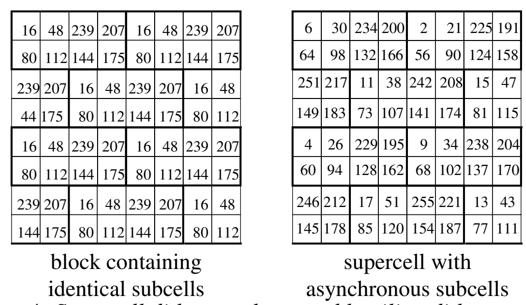

Figure 4: Supercell dither mask created by tiling dither masks.

4. Correcting unpleasant patterns between two good patterns

We first discuss how to fix unpleasant patterns between two good patterns which we want to preserve. Suppose that we are given gray levels $g_1 > g_2 > g_3$, where the patterns for g_1 and g_3 are pleasant, while the pattern for g_2 is not. The algorithm will replace g_2 with a new pattern while preserving the patterns for g_1 and g_3. Let us denote the pattern of black pixels for g_i by P_i. Because of the stacking conditions, the pixels in P_1 are a subset of the pixels in P_2 which in turn are a subset of the pixels in P_3. Therefore the pattern for P_2 is obtained by setting the pixels in P_1 (which we will call the *fixed* pixels) to black and setting a subset of the free pixels [2] to black and rearranging the free pixels.

[1] We assume that $0 \le g \le 1$ where $g = 0$ denotes black and $g = 1$ denotes white.

[2] The *free* pixels are defined as the pixels which are in P_3 but not in P_1.

This is illustrated in Figure 5. The rearranging will occur by a constrained blue noise algorithm, which rearranges only the free pixels and ensures that the fixed pixels do not move (Section 5). Furthermore, depending on the gray level and application, an additional clustering constraint is imposed to ensure the black free pixels are adjacent to the fixed pixels. This is to prevent the occurrence of isolated pixels in the halftone patterns.

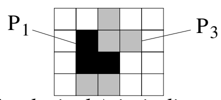

Figure 5: *P_1 (the fixed pixels) is indicated by the black pixels whereas P_3 is indicated by the black and gray pixels. The free pixels are indicated by the gray pixels. P_2 is obtained by setting the fixed pixels to black, and setting some of the free pixels to black. The free pixels are rearranged using a constrained blue noise algorithm to obtain a pleasant pattern.*

5. Constrained blue noise

Blue noise has been found to be pleasant to the eye and can be defined as having a spectral distribution where the low frequencies are attenuated [4]. For illustrative purposes, we use in this paper the void and cluster method proposed by Ulichney [5] to generate blue noise with a specific number of black pixels although other methods of blue noise generation can be used. The void and cluster method generates blue noise by moving black pixels in dense clusters of black pixels to dense clusters of white pixels. The density measure is calculated by circular convolution of the pattern with the impulse response of a lowpass filter. In practice this convolution is done via a FFT. The reason for a circular convolution is because in general the dither mask is much smaller than the image and the dither mask is tiled to halftone an image. This is expressed by the modulo operation in Section 2. The pixels with the highest density are moved to the pixels with the lowest density. The difference between the highest density and the lowest density is a measure of how good the blue noise is. We extend this method to a *constrained blue noise* algorithm by specifying the restricted subset of pixels which can be moved and by restricting the locations where these pixels can be moved to. The highest density pixel and lowest density pixel are chosen from these restricted sets respectively.

6. Supercell dither mask with reduced artifacts

Given the components above, we are now ready to describe a procedure for generating supercell dither masks which reduces the artifacts present in traditional supercell dither masks.

Step 1: Construct an m by n initial dither mask M_0 either by taking an existing supercell dither mask or by tiling several single cell dither masks. This starting dither mask could be chosen to be a dispersed dither, stochastic dither or clustered dot dither. A clustered dot dither mask can be generated using a spot function which can be chosen to be circular, elliptical, line, etc., depending on the application.
Step 2: Inspect the patterns O_g produced by M_0 and select the "good" gray levels for which the corresponding pattern O_g is pleasant. Let $h_1 > h_2 > \ldots > h_k$ be this set of "good" gray levels. The determination of pleasant patterns can either be done automatically by looking at large low-frequency content at angles which are more noticeable or it can be done by visual inspection.
Step 3: Let Δg denote the smallest increment in gray level, i.e., $\Delta g = \frac{1}{255}$ for 8-bit grayscale. For each pair of adjacent "good" gray levels $g_a = h_i > h_{i+1} = g_b$, we construct the pattern for $g' = g_a - \Delta g$ by using the algorithm in Section 4. After the pattern $O_{g'}$ for g' is created, g_a is set to g', the set of patterns $H = \{O_g\}$ is updated with $O_{g'}$, and the procedure is repeated.
Step 4: After the set of patterns $H = \{O_g\}$ is created, the final dither matrix M is created by setting $M(i, j) = g$ if O_g has a white pixel in the (i, j)-th position and $O_{\hat{g}}$ has a black pixel in the (i, j)-th position for all $\hat{g} < g$.

7. Low pass filter and human visual system

The choice of the impulse response of the low pass filter depends on the model of the human visual system (HVS). Some choices include Gaussian filters and Butterworth filters [5, 6] or exponential functions [7, 8, 9]. Many of these models also take into account the fact that our visual systems is not isotropic with respect to angle; our visual system is less sensitive to luminance changes in the diagonal directions.

8. Examples

In Figs. 6a and 6b we show two segments of a gray wedge halftoned by a traditional supercell clustered dot dither mask and by a supercell clustered dot dither mask created using the proposed algorithm respectively. In both cases, the cluster cell size is 3x6. Note that many of the unpleasant patterns occur in the lightest part of the wedgearea.

By choosing the "good" patterns as the pattern with all white pixels ($g = 1$), the pattern with all black pixels ($g = 0$) and a checkerboard pattern at 50% gray $g = 0.5$, we can create a stochastic (or FM) dither mask using the proposed algorithm which have a checkerboard pattern at 50%.

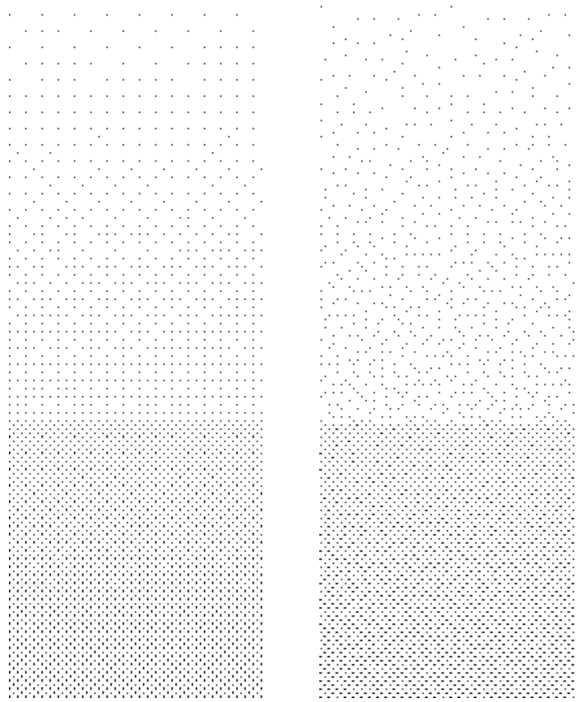

(a) (b)

Figure 6: *Segments of a gray wedge halftoned with (a) traditional supercell dither. (b) proposed algorithm.*

9. Applications

Besides the application discussed above to remove unpleasant patterns, the proposed algorithm can also be used in other applications.

9.1. Increasing the number of gray levels

The above algorithms can be used to increase the number of gray levels which can be rendered. For instance an initial dither mask which can only render 33 gray levels can be increased to render 1024 gray levels. The size of the supercell dither mask gives an upper bound for the number of gray levels which can be rendered.

9.2. Calibrated dither masks

In the previous discussion we assume that all the gray levels are uniformly dispersed between white and black. This assumes a linear relationship between the input gray level and the printed output gray level, i.e., the number of white pixels in a pattern O_g is proportional to the gray level g. This is not true when nonlinear effects such as dot gain or dot overlap are introduced. A common solution to this problem is to first apply a tone reproduction curve on the input data to compensate for this nonlinear effect and then use the linear halftone dither mask. This generally reduces the number of gray levels and does not compensate completely for the nonlinear effect, especially when this effect is large. Next, we show how we can compensate for such nonlinear effects directly by modifying the threshold array. One advantage of this method is that the number of gray levels is preserved and even strong nonlinearity can be compensated.

First, the good patterns are printed and their density measured. This provides a set of data points for the compensation curve. The data points are then interpolated by, for example, a spline to generate a compensation curve. This curve is then used to determine the gray levels at which patterns are needed. In other words, Δg in Section 6 is not constant, but is determined by this curve. Due to the stacking constraint of the dither algorithm, the curve generated by the interpolation should be monotonic. If the measured density curve is not monotonic, this has to be approximated by a monotonic curve.

The proposed algorithm can be used in color printers where multiple masks are used. Several of these techniques has been employed in the halftone design of IBM's line of laser printers, ranging from the desktop models to production print machines.

References

[1] H. R. Kang, *Optical Color technology for electronic imaging devices*. Bellingham, WA: SPIE Press, 1997.

[2] G. R. Thompson, C. P. Tresser, and C. W. Wu, "Clustered aperiodic mask." US Patent 5,917,951, 1999.

[3] G. R, Thompson, C. P. Tresser, and C. W. Wu, "Multicell clustered mask with blue noise adjustments." US Patent 6,025,930, 2000.

[4] R. A. Ulichney, "Dithering with blue noise," *Proc. of the IEEE*, vol. 76, pp. 56–79, 1988.

[5] R. Ulichney, "The void-and-cluster method for dither array generation," in *Proc. of SPIE*, vol. 1913, pp. 332–343, SPIE, 1993.

[6] M. Yao and K. J. Parker, "Modified approach to the construction of the blue noise mask," *J. of Elec. Imag.*, vol. 3, no. 1, pp. 92–97, 1994.

[7] R. Näsänen, "Visibility of halftone dot textures," *IEEE Trans. Syst. Man, Cyb.*, vol. 14, pp. 920–924, 1984.

[8] J. Sullivan, L. Ray, and R. Miller, "Design of minimal visual modulation halftone patterns," *IEEE Trans. Syst. Man, and Cyb.*, vol. 21, pp. 33–38, 1991.

[9] J. Sullivan, R. Miller, and G. Pios, "Image halftoning using a visual model in error diffusion," *J. of Opt. Soc. of America A*, vol. 10, no. 8, pp. 1714–1724, 1993.