Ranked Dither for Robust Color Printing

Maya R. Gupta and Jayson Bowen

Dept. of Electrical Engineering, University of Washington, Seattle, USA;

ABSTRACT

A spatially-adaptive method for color printing is proposed that is robust to printer instabilities, reproduces smooth regions with the quality of ordered dither, reproduces sharp edges significantly better than ordered dither, and may be less susceptible to moire. The new method acts in parallel on square, non-overlapping blocks of each color plane of the image. For blocks with low spatial activity, standard ordered dither is used, which ensures that smooth regions are printed with acceptable quality. Blocks with high spatial activity are halftoned with a proposed variant of dither, called ranked dither. Ranked dither uses the the same ordered dither matrix as standard dither, but the ranks of the thresholds are used rather than the thresholds themselves. Ranked dither is more sensitive than ordered dither to edges and more accurately reproduces sharp edges. Experiments were done with standard ordered dither masks of size 130, 130, 128, 144 for the cyan, magenta, yellow, and black planes respectively. Both on-screen and in-print, the results were sharper halftones. The entire process can be implemented in parallel and is not computationally expensive.

Keywords: dither, halftoning, color

1. INTRODUCTION

Different halftoning methods can make a significant difference in the printed image quality of one-bit and two-bit printers, particularly if there are problems printing isolated dots, or misregistration of color planes. Ordered dither is a standard halftoning choice when robustness to such printing noise and variations is a goal. Ordered dither can create smoother flat regions than error diffusion, but the resulting halftones generally appear less sharp than error diffusion halftones. In this paper, we propose a spatially-adaptive halftoning method for robust color printing. The goal of spatially-adaptive halftoning is to use halftoning well-suited for smooth areas in smooth areas, and to use a halftoning method better suited for edges in areas where there are edges. We compare the proposed *ranked dither* to standard ordered dither, with respect to the goals of being robust to printer instabilities, reproducing smooth regions without graininess or artifacts, and reproducing sharp edges.

To ensure robustness to printer color plane misregistration and to minimize memory needed, most commercial halftoning algorithms process each color plane independently. Standard dither compares each color plane of an image to a repeated array of threshold values. Ordered dither uses ordered threshold values that create clusters of dots for more robustness to instabilities in the printing process. The ordered dither clusters of dots are regularly spaced, but at different angles for each color plane. Besides being relatively robust to the variations in the printing process, ordered dither can create smoother flat regions than error diffusion. Ordered dither smoothes edge information however. A thorough review of dither and other common halftoning techniques can be found in Kang's book on digital halftoning.[1]

2. RELATED WORK

We discuss halftoning algorithms related to the ideas proposed.

Send correspondence to M. R. G.: gupta@ee.washington.edu

2.1. Hybrid halftoning

Hybrid halftoning algorithms use multiple halftoning methods for different regions of the image. For example, in a recent patent,[2] different halftoning methods are applied to different regions based on an activity index which measures spatial activity. The patent teaches how to smoothly transition between error diffusion and standard dither as the image content changes. The spatially-adaptive method proposed in this paper differs in that the different halftoning methods are ordered dither and ranked dither, and that the transitions between the two methods are pre-defined to occur at rectangular boundaries. Also, the transitions between the two methods occur rapidly over small regions of the image.

Another recent patent details a method called replacement halftoning,[3] which combines dither and error diffusion in a more complicated manner. Replacement halftoning dithers the image but also error diffuses a virtual color channel. In certain regions of the image, the dithered output is replaced by the error diffused output. The virtual color channel is a combination of the color planes (for example, the maximum color value at each pixel). Again, the different halftoning methods are different, and in replacement halftoning the decision to switch between methods is made pixel-by-pixel, and the processing is done jointly over the color planes.

A number of other hybrid halftoning algorithms exist, which are similarly differentiated by how the transitions between methods are implemented, and what halftoning methods are used.

2.2. Halftoning with multiple dither matrices

A number of researchers have investigated halftoning with multiple dither matrices, where large dither matrices are used for smooth image areas, and smaller dither matrices are used for edge regions. In this way, a dynamic trade-off is attempted between the spatial and color resolution of the halftone. Work by Ostromoukhov et al.[4] uses multiple dither matrices to reduce banding in slowly varying color regions. In that work, a decision about which dither matrix to use is made for each pixel in turn based on local gradient information. Similarly, in work by Wandell et al.,[5] a "busyness" value is computed for each pixel based on local gradient information. The busyness value for each pixel is used to select one out of a set of dither matrices, with a goal of better reproduction of edges and texture. However, they show that this pixel-by-pixel decision leads to objectionable artifacts.

A recent patent for multi-level printing uses multiple dither processes[6] by dithering each pixel with two (or more) dither matrices, and then weighting the outputs. There has also been recent research into using multiple dither matrices based on luminance values for display on monitors.[7]

2.3. Block Color Quantization

Ranked dither groups pixels with the same value for each color plane in each window, calculates the corresponding ideal percentage of coverage for each group, and decides which pixels to turn on based on a dither mask. A similar approach is taken in previous work by Gupta, BCQ halftoning.[8–10] BCQ and ranked dither both group/cluster pixels within a window, calculate the number of dots to turn-on for each group/cluster, and use a dither matrix to determine the placement of the halftone dots. However, significant image quality differences arise due to which pixel values are used, how they are grouped, when this halftoning is applied, and how dither matrices are used to determine placement.

First, BCQ works on the color planes jointly, as opposed to treating each color plane separately. However, to be robust to printing instabilities it is better to halftone each color plane independently. Second, BCQ clusters RGB pixel values for a block (window) of the image, putting pixels with similar values into each group. The idea in that work is that only a few clusters should be needed to adequately represent all visually differentiable image regions within a block. In contrast, ranked dither groups only pixels with the same value, which can result in many groups for a window. For BCQ, each block contains pixels with different (but similar) values, so the mean RGB value is calculated for each cluster and translated into CMYK ink coverage, which determines how many CMYK pixels to turn-on for the cluster. Then, for BCQ, one dither mask is used to determine which C,M,Y, and K pixels to turn-on. In ranked dither, the color planes are separated, and different ordered dither matrices are used for each color plane to determine which pixels to turn-on for the final halftone. These ordered dither matrices do not need to be the same size as the rectangular windows, and can lie at the conventional angles for the different color planes (C and M at 15 degrees, Y at 0 degrees, and K at 45 degrees). Because BCQ uses a single dither matrix that is the same size as the window and at 0 degrees, BCQ halftones can appear blocky in slowly-varying regions unless the dpi is relatively high. Lastly, BCQ is applied to every block, while we propose only applying ranked dither to regions of high spatial activity.

3. A NEW SPATIALLY-ADAPTIVE METHOD USING ORDERED DITHER AND RANKED DITHER

In this section we describe the proposed spatially-adaptive halftoning method, which switches automatically between standard ordered dither and ranked dither, which is explained in Section 3.2. All processing is done independently for each color plane. Consider a particular color plane of the image, and a threshold image the same size as the image, but made up of ordered dither mask thresholds. As is standard, different ordered dither masks are used for each color plane to create clusters of dots at different angles. For ordinary dither, one would compare the color plane of the image to the ordered dither thresholds to determine which dots to print.

Instead, each color plane of the image is partitioned into non-overlapping windows; in this paper we use 12×12 pixel windows. Using non-overlapping windows preserves the ability to implement the halftoning as a parallel process. For each window, a decision is made as to whether there is intense spatial activity (such as an edge or texture). If there is not substantial spatial activity within the window, then standard ordered dither is applied: the image pixel is compared to its corresponding ordered dither threshold to determine if a dot should be printed there. However, if there is substantial spatial activity within the window, then ranked dither is used. (Note that, in general, the dither matrix size and shape will be different than the window size or shape.)

3.1. Determining Level of Spatial Activity

A decision function is needed to determine whether there is significant spatial activity in a window to classify it as an edge/texture window. A variety of different measures of spatial activity could be used, we found a maximum block difference worked well. Each color plane of the image is processed independently for speed and robustness. For a given window, the maximum block difference is calculated by dividing the window into equal blocks of pixels. For example, using 12×12 pixel windows, we divided each window into nine 4×4 subblocks. Calculate the average pixel value for each block, then determine the maximum difference between any two block averages. If this maximum difference is above a threshold, the window is classified as an edge/texture window.

The threshold is key; if the threshold is set too low, then small rough patches may appear in an area that should have been smooth. If the threshold is set too high, edges will be halftoned by ordered dither and hence be smoothed. Each color plane should have a different threshold set for it. For example, a black color plane would require a lower threshold than a lighter cyan color plane. Different hardware, inks, or toners would require different threshold constants.

3.2. Ranked Dither

Ranked dither is a new method for halftoning regions with high spatial activity. Ranked dither can use the same dither matrix used to halftone smooth regions, but the ranks of the thresholds are used rather than the thresholds themselves. Figure 1 and 2 show the steps of the ranked dither process. In the top left of Figure 1 is an example 4×4 window from an image, which is smaller than our recommended window size for ranked dither, and is used only to illustrate the ranked dither process. In the top right of Figure 1, the 4×4 window is shown divided into groups of pixels that have the same pixel value. The bottom left of Figure 1 shows the corresponding ordered dither thresholds for that window of that color plane of the image. The ordered dither thresholds are grouped corresponding to the groups of the pixels, shown in the bottom right of Figure 1.

Each group's common pixel value determines the percentage of ink for that group. A number of ink dots for each group is determined by multiplying the percentage of ink for the group by the number of group pixels. For example, the group of value 127 has six pixels, resulting in $round((127/255) * 6) = 3$ halftone dots for the group. The group of value 80 has three pixels, resulting in $round((80/255) * 3) = 1$ halftone dots for that group. The group of value 20 results in $round((20/255) * 6) = 0$ dots, and the group with value 128 yields one dot. To determine which dots are turned on, the dither thresholds corresponding to each group are ordered, and if d dots

20	20	128	127
127	20	20	127
80	20	20	127
80	80	127	127

Example 4×4 window of one colorplane.

20	20	128	127
127	20	20	127
80	20	20	127
80	80	127	127

The window grouped by pixel value.

105	120	135	150
90	15	30	165
75	60	45	180
240	225	210	195

Corresponding ordered dither thresholds for the window.

105	120	135	150
90	15	30	165
75	60	45	180
240	225	210	195

The ordered dither thresholds grouped according to the window's pixel values.

Figure 1. First ranked dither process steps are shown.

are to be turned on, then the smallest d dither thresholds in the group are assigned the d dots. An example is shown in the top left of Figure 2 for the group with value 127; three of the group's spots are marked on with a 1 for dot, and the remaining three group spots are marked with a 0 for no-dot. The same process is done for each group, the resulting ranked dither halftone is shown in the top right of Figure 2. For comparison, the ordered dither halftone is shown in the bottom of Figure 2.

This example shows how ranked dither can do a better job of reproducing edges. This is because the ranked dither considers each group separately, ensuring that each group receives a representative number of dots. For

			1
1			1
			0
		0	0

Ranked dither halftone evaluated only for the group with value 127.

0	0	0	1
1	0	0	1
1	0	0	0
0	0	0	0

Ranked dither halftone.

			0
1			0
			0
		0	0

Ordered dither halftone evaluated only for the group with value 127.

0	0	0	0
1	1	0	0
1	0	0	0
0	0	0	0

Ordered dither halftone.

Figure 2. Next ranked dither process steps are shown (continued from previous figure).

example, the ranked dither halftone has three dots for the six pixels with value 127, while the ordered dither halftone has only one dot for the six pixels with value 127.

Like ordered dither, the ranked dither halftones tend to result in clusters of dots. Because the dither thresholds of the ordered dither matrix are used, the ink dots which are placed for ranked dither tend to be adjacent, because similar image values tend to be adjacent. The resulting clusters of dots keep the number of *isolated dots* small in ranked dither halftones, but the ranked dither halftones will have slightly more isolated dots than ordered dither halftones. (An isolated dot is any halftone pixel which is turned on, but which is surrounded by eight neighboring pixels which are not turned on. Isolated dots are difficult to print with some print engines.)

The proposed spatially-adaptive dither only applies ranked dither to high spatial-activity regions. For regions with less spatial activity, ranked dither will yield correct average colors values, but its responsiveness to edges can translate into a responsiveness to noise. This responsiveness to noise can result in an unsmooth appearance for smooth natural regions.

4. EXPERIMENTS

Experiments compared standard ordered dither with the proposed spatially-adaptive ordered and ranked dither combination, as described in Section 3. For both methods, the ordered dither was implemented with a quad microcluster threshold array dither matrix from Kang's book on halftoning [1, p. 346]. (Note that the cyan screen in the book contains an error: threshold 107 should be 71, 123 should be 107, and a threshold of 123 should be placed above 113.) The threshold arrays contain 130, 130, 144, and 128 levels for cyan, magenta, yellow, and black respectively. The same ordered dither matrix was used for rank dither.

The thresholds to decide if the maximum block difference was large enough to classify a window as an edge window or a smooth window were 30 for the C, M, and Y color planes, and a threshold of 8 for the K color plane, making the spatially-adaptive halftoning more sensitive to edges in black.

The images used were converted from their native RGB colorspace to CMYK using Adobe Photoshop's sRGB to SWOP CMYK converter. The CMYK images were halftoned. Then, to view the images, the CMYK halftones were converted to sRGB using the Adobe Photoshop conversion from SWOP CMYK to sRGB.

5. RESULTS

Example results are shown in Figures 3, 4, and 5. Figures should be viewed from some distance to get a feel for what the halftones would look like printed at a reasonable dpi. Edges are clearly sharper in the spatially-adaptive ordered dither/ranked dither method (the bottom halftone in each figure), for example the red arm of the horse-rider in Figure 3. Smooth areas of the spatially-adaptive method look like the ordered dither halftone, for example in the background of Figure 4, because if a window was judged to be smooth, ordered dither was used.

In the spatially-adaptive halftones the normal rosette pattern of the ordered dither can be seen, but in none of the test images could we see an artifact pattern from the square windows. Also, the rosette pattern from the dither tends to be slightly less visible because it is broken up by the ranked dither windows, for example the yellow background in Figure 4.

The spatially-adaptive halftoning uses two repeating patterns: the regular pattern of the ordered dither mask, and the regular square windows used for the spatial adaptation. Often when different repeating patterns are used moire problems arise. In fact, we saw less moire with the spatially-adaptive halftoning, we conjecture this is because the ranked dither simply does a better job at capturing edges. As an example, Figure 5 shows the halftones for a cyan zoneplate image. The halftones are rendered in black and white for ease of viewing.

6. CONCLUSIONS

On-screen and printed results show that the ranked dither achieves sharper edges than ordered dither. By spatially-adapting the halftoning to use either ordered or ranked dither, both smooth and sharp regions can be reproduced with high quality. Using non-overlapping windows for the spatial adaptation enables the adaptation to be implemented in parallel. Ranked dither has no computationally intensive steps. In conclusion, a spatially-adaptive mix of ordered and ranked dither is a practical halftoning solution for fast printing and can lead to high image quality while creating halftones that are robust to printer instabilities.

REFERENCES

1. H. R. Kang, *Digital Color Halftoning*, SPIE Press, Bellingham, 1999.
2. Q. Lin, "Adaptive halftoning based on image content." US Patent # 5970178, 1999.
3. M. Gupta and H. Takahashi, "Replacement halftoning." US Patent # 6851783, 2005.
4. V. Ostromoukhov and S. Nehab, "Halftoning with gradient-based selection of dither matrices." US Patent # 5701366, 1997.
5. H. Z. Hel-Or, X. M. Zhang, and B. A. Wandell, "Adaptive cluster dot dithering," *Journal of Electronic Imaging* **2**(8), pp. 133–144, 1999.
6. K. Kritayakirana, D. Tretter, and Q. Lin, "Adaptive halftoning method and apparatus." US Patent # 6760126, 2004.
7. S. Daly and X. Feng, "Methods and systems for adaptive dither structures." US Patent Pending # 50174360, 2004 (filed).
8. M. Gupta, M. Gormish, and D. Stork, "Block color quantization: A new method for color halftoning," *Proc. of the IEEE Intl. Conf. on Image Proc.* , 2000.
9. M. Gupta, "Improved block quantization halftoning algorithm." US Patent Pending, 2001 (filed).
10. M. Gupta, M. Gormish, K. Berkner, and D. Stork, "A block color quantization method for color halftoning." US Patent Pending, 2000 (filed).

Ordered Dither

Spatially-adaptive Ordered and Ranked Dither

Figure 3. Example comparing ordered dither (top) with using ordered and ranked dither (bottom). Image is a 150×150 pixel crop from an Indian painting.

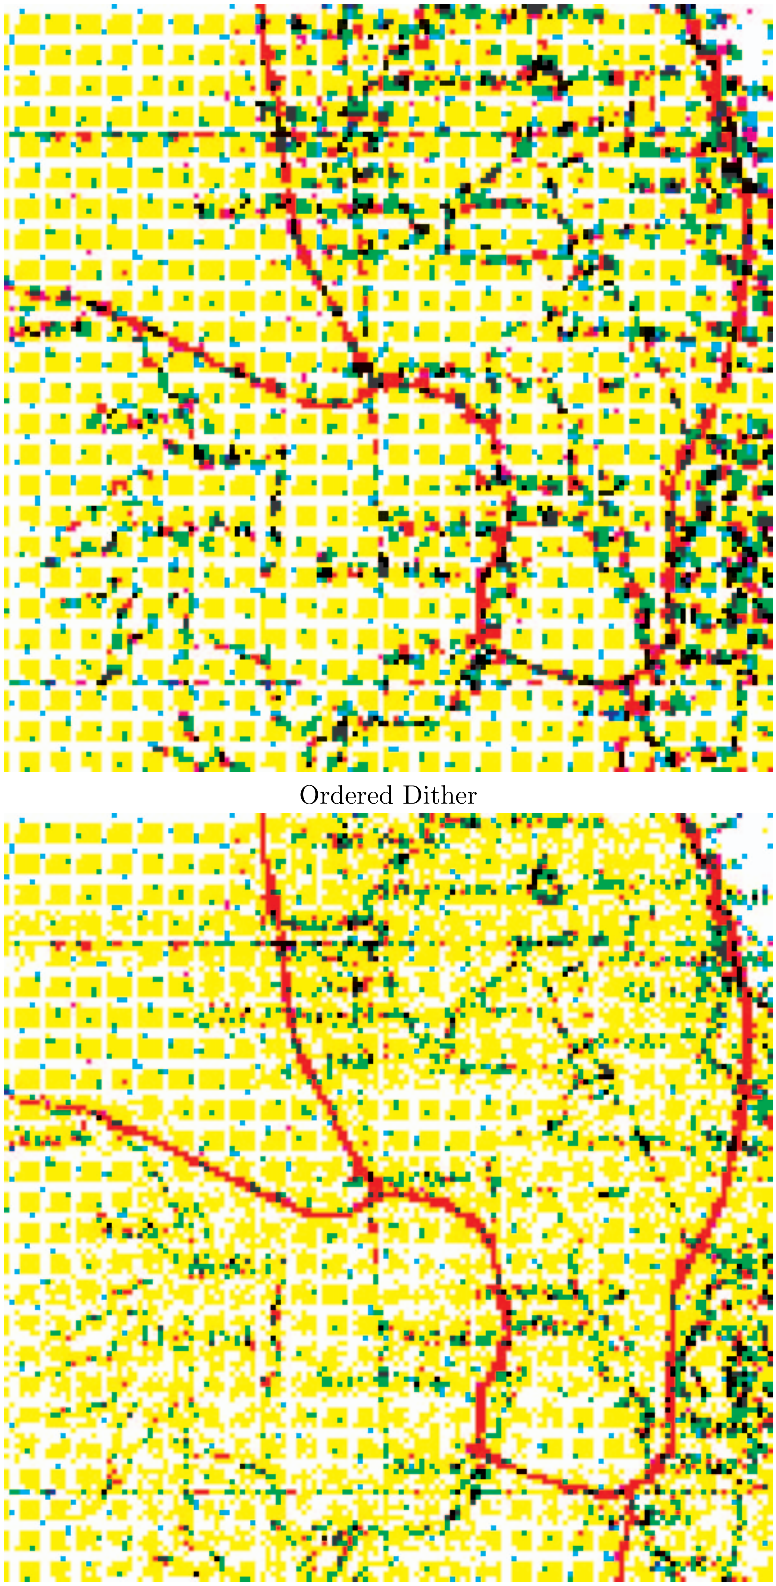

Ordered Dither

Spatially-adaptive Ordered and Ranked Dither

Figure 4. Example comparing ordered dither (top) with using ordered and ranked dither (bottom). Image is a 150×150 pixel crop from a map of Africa.

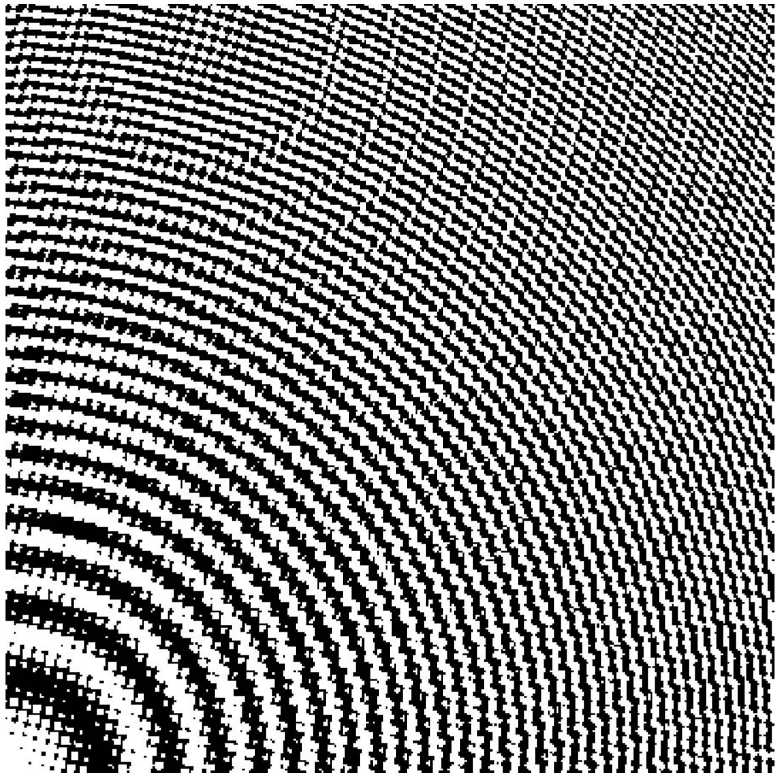

Ordered Dither

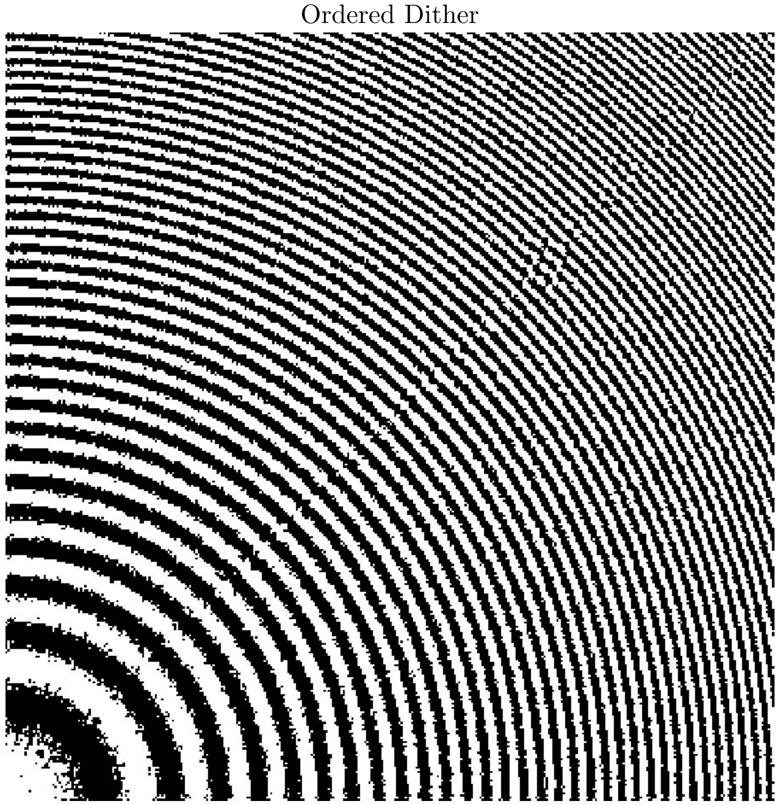

Spatially-adaptive Ordered and Ranked Dither

Figure 5. Example comparing ordered dither (top) with using ordered and ranked dither (bottom). Image is a quarter of a cyan zoneplate.