Rotated Dispersed Dither: a New Technique for Digital Halftoning

Victor Ostromoukhov, Roger D. Hersch, Isaac Amidror
Swiss Federal Institute of Technology (EPFL)
CH-1015 Lausanne, Switzerland
victor@di.epfl.ch

ABSTRACT

Rotated dispersed-dot dither is proposed as a new dither technique for digital halftoning. It is based on the discrete one-to-one rotation of a Bayer dispersed-dot dither array. Discrete rotation has the effect of rotating and splitting a significant part of the frequency impulses present in Bayer's halftone arrays into many low-amplitude distributed impulses. The halftone patterns produced by the rotated dither method therefore incorporate fewer disturbing artifacts than the horizontal and vertical components present in most of Bayer's halftone patterns. In grayscale wedges produced by rotated dither, texture changes at consecutive gray levels are much smoother than in error diffusion or in Bayer's dispersed-dot dither methods, thereby avoiding contouring effects.

Due to its semi-clustering behavior at mid-tones, rotated dispersed-dot dither exhibits an improved tone reproduction behavior on printers having a significant dot gain, while maintaining the high detail rendition capabilities of dispersed-dot halftoning algorithms. Besides their use in black and white printing, rotated dither halftoning techniques have also been successfully applied to in-phase color reproduction on ink-jet printers.

1 INTRODUCTION

Due to the proliferation of low-cost bi-level desktop printers, converting grayscale images to halftoned black and white images remains an important issue. While significant progress has been made to improve the quality of error-diffusion algorithms, dispersed-dot dither algorithms, which are much faster, have not been substantially improved over the last 20 years. Only recently, a new method for generating well-dispersed dither arrays was reported [14].

This contribution presents a new dispersed-dot dither technique, based on a discrete one-to-one rotation of a conventional dispersed-dot dither threshold array. It will be shown that, compared to Bayer's ordered dither, this new rotated dispersed-dot dither technique produces diagonally oriented small-sized halftone clusters, reduces the power of individual low frequency components, incorporates smooth pattern transitions between neighbouring gray levels and provides an improved tone reproduction curve, halfway between classical clustered-dot and Bayer's dispersed-dot dither techniques.

In section 2, we give a brief survey of the main halftoning techniques and show their respective advantages and drawbacks in terms of computation complexity, tone reproduction behavior and detail resolution. In section 3, the proposed rotated dither method is presented. It is based on a one-to-one discrete rotation of a dither tile made of replicated Bayer dither arrays. In section 4, we compare the proposed rotated dither method with error diffusion and with Bayer's dispersed-dot dither method by showing halftoned images. In section 5, we try to explain why the proposed rotated dither algorithm generates globally less perceptible artifacts than Bayer's by analyzing the frequencies produced by the halftoning patterns at different gray levels. In section 6, we analyze the tone reproduction behavior of the different algorithms and show that the rotated dispersed-dot dither algorithm has a robustness close to that of clustered-dot halftoning, therefore being appropriate for printers with a significant dot gain.

2 BACKGROUND

Both classical clustered-dot halftoning [15] and dispersed-dot dither halftoning algorithms [13] are supported by modern PostScript level-2 printers which provide the means to customize their *dither threshold array* by downloading the threshold values used for the halftoning process. With a given dither threshold array, the halftoning process consists of scanning the output bitmap and, for each output pixel, finding its corresponding locations both in the dither array and in the grayscale input pixmap image, comparing corresponding input image pixel intensity values to dither array threshold values, and accordingly writing pixels of one of two possible output intensity levels to the output image bitmap. Dither array based halftoning is very efficient: only one comparison is needed per output device pixel. Furthermore, output device pixels can be computed independently, which enables the halftoning process to be parallelized and pipelined [10]. Patent literature demonstrates that significant efforts have been made for improving the quality of exact angle clustered-dot dithering techniques [11],[5],[12]. However, the only recent effort known to improve the quality of dispersed-dot dither halftoning is the void-and-cluster method proposed by Ulichney [14].

Error diffusion algorithms, first introduced by Floyd and Steinberg [4], are more computationally intensive methods. They require diffusing the error, i.e. the difference between the output device pixel intensity and the original source image intensity, to a certain number of neighbours. In the traditional Floyd-Steinberg algorithm, parallel computation of output device pixels is not possible since each output pixel depends on the error transferred by its neighbours. The dot diffusion method proposed by Knuth [9] and further improved by Zhang and Webber [16] removes these limitations by

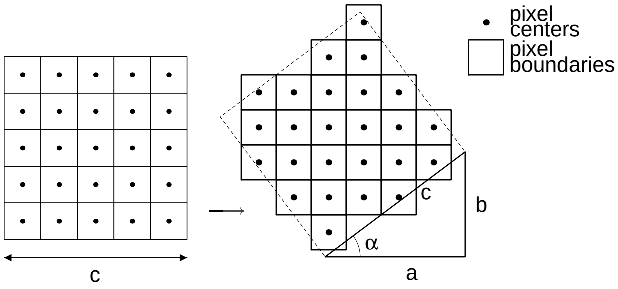

Figure 1: Rotating a square array of c^2 pixels by a Pythagorean angle $\alpha = \arctan(b/a)$, where $c = 5, a = 4, b = 3$.

tiling the output bitmap plane into limited size quadratic neighbourhoods, within which the error is allowed to propagate. Nevertheless, these improved error diffusion halftoning methods still require much more computation than dither array based halftoning, since errors must be propagated to several neighbours and since output pixels cannot always be computed in scanline order.

Generally, clustered-dot dither is preferred to Bayer's dispersed-dot dither because of Bayer's dispersed-dot dither artifacts and poor tone reproduction behavior. We will show that the proposed rotated dispersed-dot dither method generates slightly clustered halftone patterns which offer good tone reproduction behavior and at the same time avoid contouring effects by providing patterns which grow into one another smoothly as the intensity level increases.

3 THE ROTATED DITHER METHOD

Bayer's dispersed-dot ordered dither method has been shown to be optimal in the sense that in each gray level, the lowest frequency is as high as possible [3]. Nevertheless, as Figure 5b shows, low frequency components given by the dither array period size are quite strong for a large number of gray levels. These low frequency halftoning artifacts are well perceived, since at increasing intensity levels they switch back and forth from horizontal and vertical to diagonal directions. The transition between one intensity level and the next one creates an abrupt pattern change which appears in the halftoned image as a false contour.

Since human eye sensitivity to gratings decreases considerably at oblique orientations [2], we can make the perceived halftoning artifacts less visible by rotating the dispersed dither array. Moreover, we will show in section 5 that the discrete rotation of the dispersed dither array decreases the power of the individual low-frequency components by distributing their low-frequency energy over a larger set of frequencies.

The discrete rotation we intend to apply to dither arrays has some similarities with the rotation of bitmap images. For the sake of simplicity, we will therefore also use in this context the term "pixel" to denote a simple dither array element. The rotating task consists of finding a discrete rotation, which generates a "rotated Bayer threshold array" whose threshold values are exactly the values of the original Bayer dither threshold array. Let us consider rotations of binary pixel grids composed of unit size pixel squares. Exact rotation of a square pixel grid by an angle α around the position (x_0, y_0) can be described by the following transformation applied to the pixel centers (x, y):

$$x' = (x - x_0)\cos\alpha - (y - y_0)\sin\alpha + x_0 \qquad (1)$$

$$y' = (x - x_0)\sin\alpha + (y - y_0)\cos\alpha + y_0$$

In the general case, more than one rotated original pixel center may fall within a single pixel's square boundary in the destination grid, and some destination pixels may remain empty [6]. We therefore need a one-to-one discrete rotation which unambiguously maps the set of original dither array elements into the new set formed by the rotated dither array.

Let us consider the continuous boundary of a square array of c^2 discrete pixels (Figure 1). It can be shown that if this boundary is rotated by a Pythagorean angle $\alpha = \arctan(b/a)$, where a and b are Pythagorean numbers satisfying the Diophantine equation $a^2 + b^2 = c^2$, the resulting rotated square boundary contains the same number of pixel centers as the original pixel array (Figure 1). Therefore, a discrete one-to-one rotation can be obtained by rotating with a Pythagorean angle and by an appropriate one-to-one mapping between the set of dither elements belonging to the original square and the set of dither elements belonging to the rotated square.

Such a discrete one-to-one rotation is obtained by rotating with a Pythagorean angle $\alpha = \arctan(b/a)$, where $c = 5, a = 4$, $b = 3$ and $\alpha = \arctan(3/4) = 36.87°$ and by applying rounding operations. Let us assume that (i, j) is the coordinate system of the original dither array and (x, y) the coordinate system of the rotated dither array, and that i_0, j_0, and respectively x_0, y_0 are integer values defining the location of the given original square dither array, respectively the location of the rotated dither array:

$$x = \text{round}\left(\frac{a}{c} * (i - i_0) - \frac{b}{c} * (j - j_0)\right) + x_0 \qquad (2)$$

$$y = \text{round}\left(\frac{b}{c} * (i - i_0) + \frac{a}{c} * (j - j_0)\right) + y_0$$

For this discrete one-to-one rotation, the distance between corresponding rotated and rounded pixel centers is either zero or equal to $1/\sqrt{5} = 0.447214$ (see Figure 3). In order to apply this discrete one-to-one rotation to a square Bayer dither threshold array, we have to consider the Bayer threshold array D^n of size $n x n$, replicated c times vertically and horizontally, since its side length must be an integer multiple of Pythagorean hypotenuse c (Figure 2a). We will denote the dither tile obtained this way by D^{c*n}. In Figures 2 and 4 we use as example the Bayer dither threshold array of size $n = 4$.

The discrete one-to-one rotation described by equations (2) applied to a dither tile D^{c*n} made of a replicated Bayer dither threshold array yields a rotated dither tile R^{c*n} (Figure 2b), which also paves the plane like the original tile D^{c*n}. The elements of the rotated dither tile R^{c*n} have the same dither values as the corresponding original dither elements from the dither tile D^{c*n}. The subpixel displacements expressing the difference between continuous and discrete rotations are shown in Figure 3. The circle at the top left of Figure 3 shows the repetitive subpixel displacement patterns. Pixel centers of the same color within one oblique row in tile R^{c*n} of Figure 3 correspond to pixel centers in one horizontal row of the original tile D^{c*n}.

The rotated dither tile R^{c*n} includes irregularly arranged dither sub-arrays paving the dither plane (one of these sub-arrays is highlighted in blue in Figure 4). Such a diagonally oriented dither sub-array is easily transformed, according to Holladay's algorithm [7], into an equivalent rectangular dither tile paving the plane and containing the same number of dither threshold values (Figure 4).

Let us recall the computational effort required for dither array based halftoning. Once the dither tile containing the dither threshold values has been generated, it can be stored in memory and used for halftoning the grayscale images. Halftoning time is determined by the time required to scan the output image bitmap space, pixel by pixel, maintaining a pointer both at the corresponding places in the grayscale source image and in the Holladay rectangular dither tile and performing comparisons between corresponding dither threshold and input pixel intensity values. Computation time is proportional to the number of output bitmap pixels and is much smaller than the time needed by error-diffusion algorithms requira)

0	8	2	10	0	8	2	10	0	8	2	10	0	8	2	10	0	8	2	10
12	4	14	6	12	4	14	6	12	4	14	6	12	4	14	6	12	4	14	6
3	11	1	9	3	11	1	9	3	11	1	9	3	11	1	9	3	11	1	9
15	7	13	5	15	7	13	5	15	7	13	5	15	7	13	5	15	7	13	5
0	8	2	10	0	8	2	10	0	8	2	10	0	8	2	10	0	8	2	10
12	4	14	6	12	4	14	6	12	4	14	6	12	4	14	6	12	4	14	6
3	11	1	9	3	11	1	9	3	11	1	9	3	11	1	9	3	11	1	9
15	7	13	5	15	7	13	5	15	7	13	5	15	7	13	5	15	7	13	5
0	8	2	10	0	8	2	10	0	8	2	10	0	8	2	10	0	8	2	10
12	4	14	6	12	4	14	6	12	4	14	6	12	4	14	6	12	4	14	6
3	11	1	9	3	11	1	9	3	11	1	9	3	11	1	9	3	11	1	9
15	7	13	5	15	7	13	5	15	7	13	5	15	7	13	5	15	7	13	5
0	8	2	10	0	8	2	10	0	8	2	10	0	8	2	10	0	8	2	10
12	4	14	6	12	4	14	6	12	4	14	6	12	4	14	6	12	4	14	6
3	11	1	9	3	11	1	9	3	11	1	9	3	11	1	9	3	11	1	9
15	7	13	5	15	7	13	5	15	7	13	5	15	7	13	5	15	7	13	5
0	8	2	10	0	8	2	10	0	8	2	10	0	8	2	10	0	8	2	10
12	4	14	6	12	4	14	6	12	4	14	6	12	4	14	6	12	4	14	6
3	11	1	9	3	11	1	9	3	11	1	9	3	11	1	9	3	11	1	9
15	7	13	5	15	7	13	5	15	7	13	5	15	7	13	5	15	7	13	5

b)

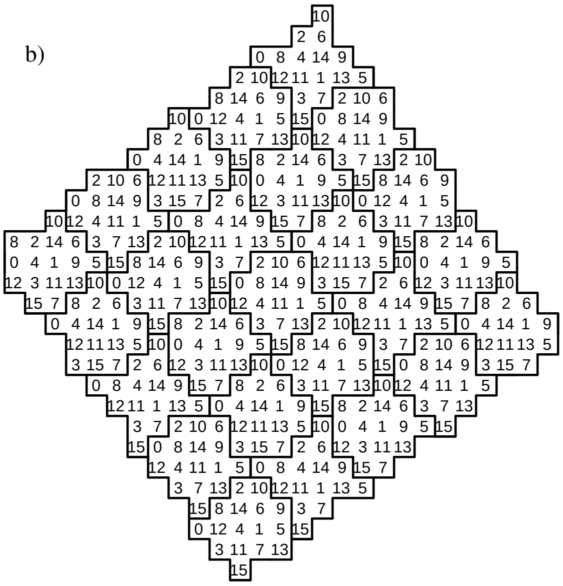

Figure 2: (a) Original replicated dither threshold array D^{c*n} and (b) rotated dither tile R^{c*n} obtained by the discrete rotation of tile D^{c*n}.

ing at each output bitmap pixel the error to be distributed to several neighbouring pixel elements [16].

4 EXPERIMENTAL RESULTS

Let us compare the proposed rotated dither method with standard Bayer dispersed-dot dither, Floyd-Steinberg error diffusion and clustered-dot dither. As can be seen from Figure 5, rotated dispersed-dot halftoning generates less texture artifacts than the standard Bayer method. Moreover, as it can be seen in the gray wedges at the top of the images, the tone reproduction of rotated dither is much better than that of the Bayer dither and slightly better than that of error diffusion.

As can be seen from Figure 5, the detail rendition capabilities of rotated dither are slightly better than those of Bayer dither, and much better than those of clustered-dot dither. It is well known that error diffusion has an inherent sharpening behavior and therefore provides better detail rendition capabilities [8]; nevertheless, it creates visible patterns in constant intensity areas.

In the grayscale wedges produced by rotated dither at consecutive gray levels, especially at intensity levels close to 50%, texture changes produced by rotated dither are smoother than in Floyd-Steinberg's error-diffusion or in Bayer's dispersed-dot dither methods (Figure 6).

Rotated dither halftoning techniques can also be applied both to single orientation multi-color in-phase printing and to increase the number of perceived colors on displays with a limited number of intensity levels.

5 FREQUENCY DOMAIN ANALYSIS

In order to compare the proposed rotated dither algorithm with Bayer's dispersed-dot dither algorithm, we analyze their respective halftone patterns at various intensity levels by comparing their frequency amplitude spectra. It is well known that the human eye is most sensitive to the lowest frequency components of screen dot patterns, especially if they have a horizontal or vertical orientation. Bayer demonstrated that his ordered dither algorithm minimizes the occurrence of low-frequency components. However, he didn't take into account the amplitudes of the spectral frequencies, nor did he consider the anisotropic behavior [2] of the eye's contrast sensitivity function (CSF). In this section, we will show that the rotated dispersed-dot dither algorithm produces halftone patterns with a larger number of frequency components than Bayer's.

To compare the two dither algorithms, we consider halftone cells produced according to the example shown in section 3, representing gray levels $0, \frac{1}{16}, \frac{2}{16}, .., \frac{15}{16}$. In that example, the Holladay rectangle paving the corresponding rotated dither tile has a size of 20 x 4 pixels. By choosing a sample array of a size which is an integer multiple of the horizontal and vertical replication period of the dither rectangle paving the plane, we ensure that the frequencies present in the Discrete Fourier Transform (DFT) of the sample array are located exactly on the spatial-frequency sampling grid, thereby avoiding leakage effects and ensuring that the spectral impulses fall exactly on the center of DFT impulses [1].

In the examples shown in Figure 9, we consider 80x80 pixel sample halftone arrays created according to Bayer's 4x4 dispersed-dot dither array and according to the rotated dither method, using 4x4 dispersed-dot dither arrays, replicated 5 times and rotated. We compare the amplitude spectrum of the respective halftone arrays for several gray levels. Figure 9 shows the halftone patterns as well as their corresponding DFT impulse amplitudes at grayscale levels $\frac{1}{16}, \frac{3}{16}, \frac{5}{16}$ and $\frac{7}{16}$. The dot surfaces in the spectra are proportional to the amplitude of the corresponding frequency impulses.

The following observations can be made:

– Compared to Bayer's dither, rotated dither incorporates additional lower frequency components having, however, a lower amplitude than the original Bayer frequency component.

– Bayer's original frequency components are present in the rotated dither patterns; they are simply rotated by $\alpha = \arctan(3/4)$, the angle used for the discrete one-to-one rotation. However, their amplitude is considerably lower than that of the corresponding Bayer frequency components. Figure 10 shows the amplitude spectrum for halftone at intensity level $\frac{3}{16}$ generated by the Bayer and the rotated dither methods. The Bayer frequency impulses present in the frequency spectrum of the rotated halftone pattern are marked by rings.

The rotated dither method rotates the frequency impulses present in the Bayer halftone array and splits one part of their amplitude into many lower amplitude distributed impulses. The power of the original Bayer main frequency components is therefore reduced, and additional low energy components are created. Moreover, since the main frequency components are rotated, they are less perceptible to the human eye than the horizontal and vertical components present in most of Bayer's halftone patterns.

6 TONE REPRODUCTION BEHAVIOR

The tone reproduction behavior of a given halftoning algorithm is heavily dependent on the dot gain behavior of the considered printer. At levels darker than 50%, dispersed-dot dither as well as error-diffusion halftoning algorithms tend to generate one pixel wide elongated white surface areas which may shrink considerably due to dot gain. We will show that the clustering behavior of rotated dither at mid-tones has a positive impact on its tone reproduction capabilities.

Let us compare the tone reproduction behavior of the rotated dispersed-dot dither algorithm with Bayer's dispersed-dot dither algorithm, error diffusion and clustered-dot halftoning. For this purpose we measure and plot the tone reproduction behavior of a variable intensity grayscale wedge printed on a black and white laser printer. Figure 7 shows the tone reproduction behavior for the considered halftoning techniques and Figure 8 the grayscale wedges printed at 300 dpi on a commercial laser printer.

If we compare the rotated dither method with Bayer's dither and error-diffusion, Figures 7 and 8 clearly show that for printers with a certain dot gain, rotated dispersed-dot dither has a behavior closer to clustered-dot dither, especially at mid-tones. It is therefore a good candidate for dispersed-dot printing on laser or ink-jet printers (300–800 dpi) having a significant dot gain.

While the clustered-dot dither method is the most robust in terms of reproduction behavior at large dot gains, the new rotated dispersed-dot dither method offers both a favorable reproduction behavior and good detail rendition capabilities.

7 CONCLUSION

This contribution proposes a new high-speed dispersed-dot dither algorithm whose dither array is obtained by discrete rotation of a conventional Bayer dither array. Discrete rotation attenuates the visible frequency components produced by halftoning, and spreads the power of the visible low frequency artifacts onto additional frequencies, thereby reducing the perception of the produced artifacts.

Furthermore, when rendering images at smoothly increasing intensity levels, rotated dispersed-dot dither generates less contouring effects than Bayer's dither method and less artifacts than Floyd-Steinberg's error-diffusion method.

Due to its additional lower-frequency components, the proposed rotated dispersed-dot algorithm has a more pronounced clustering behavior than both Bayer's dispersed-dot dither and error-diffusion algorithms. It therefore exhibits an improved tone reproduction behavior on printers having a significant dot gain. Rotated dispersed-dot dither is nearly as robust as clustered-dot dither while offering much higher detail rendering capabilities.

Rotated dither halftoning techniques can be applied both to single orientation multi-color in-phase printing and to increase the number of perceived colors on display devices with a limited number of intensity levels.

ACKNOWLEDGMENTS

We would like to thank the HP Research Labs (Palo Alto, CA) and the Swiss National Fund (grant No. 21-31136.91) for supporting the project.

REFERENCES

[1] E.O. Brigham, *The Fast Fourier Transform and its Applications.* Prentice-Hall, UK, 1988.

[2] F.W. Campbell, J.J. Kulikowski, J. Levinson, The effect of orientation on the visual resolution of gratings, *J. Physiology*, London, 1966, Vol 187, 427-436.

[3] B.E. Bayer, An Optimum Method for Two-Level Rendition of Continuous-Tone Pictures, *IEEE 1973 International Conference on Communications,* Vol. 1, June 1973, 26-11–26-15.

[4] R.W. Floyd, L. Steinberg, An Adaptive Algorithm for Spatial Grey Scale, *Proc. SID,* 1976, Vol 17(2), 75-77.

[5] Gall, Winrich, "Method and Apparatus for Producing Half-Tone Printing Forms with Rotated Screens on the Basis of Randomly Selected Screen Threshold Values", U.S. Patent No. 4700235 (1987), Assignee: Dr. Ing. Rudolf Hell GmbH. (Fed. Rep. of Germany).

[6] R.D. Hersch, Raster Rotation of Bilevel Bitmap Images, *Eurographics'85 Proceedings,* (Ed. C. Vandoni), North-Holland, 1985, 295-308.

[7] Holladay T. M., "An Optimum Algorithm for Halftone Generation for Displays and Hard Copies," *Proceedings of the Society for Information Display,* 21(2), 1980, 185-192.

[8] K.T. Knox, Edge Enhancement in Error Diffusion, *SPSE's 42nd Annual Conf.,* May 1989, 56-79.

[9] D.E. Knuth, Digital Halftones by Dot Diffusion, *ACM Trans. on Graphics,* 6(4), 1987, 245-273.

[10] M. Morgan, R.D. Hersch, V. Ostromoukhov, Hardware Acceleration of Halftoning, Proceedings SID International Symposium, Anaheim, in *SID 93 Digest,* May 1993, Vol XXIV, 151-154.

[11] Rosenfeld, Gideon, "Screened Image Reproduction", U.S. Patent No. 4456924 (1984). Assignee: Scitex Corporation Ltd. (Israel).

[12] Troxel, D.E., "Method and Apparatus for Generating Digital, Angled Halftone Screens Using Pixel Candidate Lists and Screen Angle Correction to Prevent Moire Patterns", U.S. Patent No. 5124803 (1992). Assignee: ECRM.

[13] R. Ulichney, *Digital Halftoning,* The MIT Press, Cambridge, Mass., 1987.

[14] R. Ulichney, The void-and-cluster method for dither array generation, *IS&T/SPIE Symposium on Electronic Imaging Science & Technology, Proceedings Conf. Human Vision, Visual Processing and Digital Display IV,* (Eds. Allebach, Rogowitz), SPIE Vol. 1913, 1993, 332-343.

[15] J.A.C. Yule, *Principles of Colour Reproduction,* John Wiley & Sons, NY (1967).

[16] Y. Zhang, R.E.Webber, Space Diffusion: An Improved Parallel Halftoning Technique Using Space-Filling Curves, Proceedings of SIGGRAPH'93, In *ACM Computer Graphics,* Annual Conference Series, 1993, 305-312.

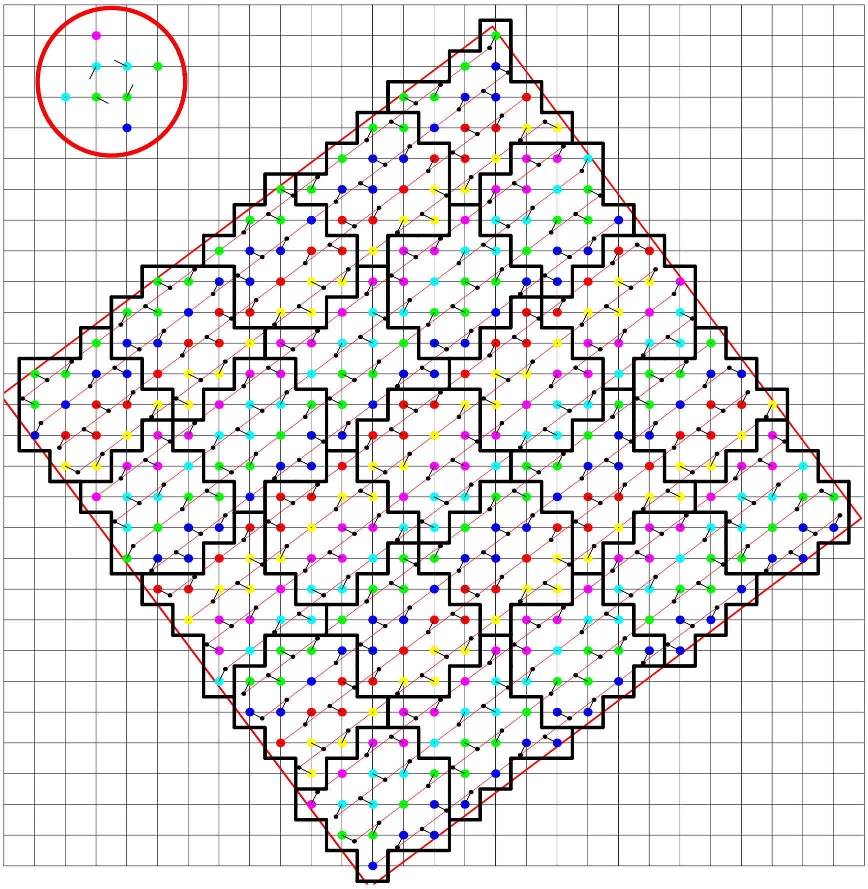

Figure 3: Subpixel displacements due to the difference between continuous and discrete rotation.

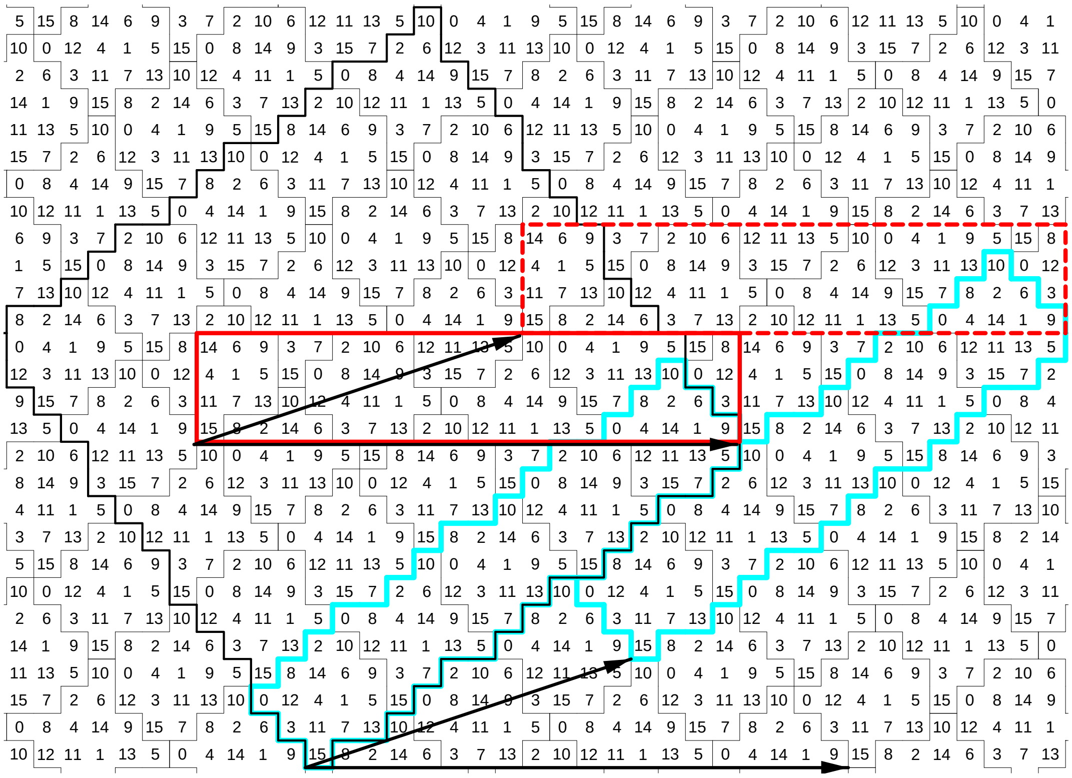

Figure 4: Rotated dither sub-tile (blue) and corresponding Holladay rectangle (red).

Figure 5: Grayscale image, halftoned at 150 dpi with (a) the 16x16 rotated dither algorithm, (b) 16x16 Bayer's dither algorithm, (c) a conventional diagonally oriented clustered-dot dither array and (d) error diffusion.

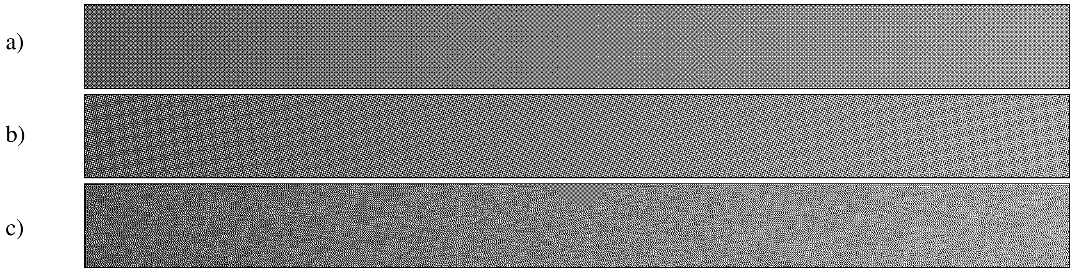

Figure 6: Intensity ramp at intensity levels between 37.5% and 62.5% for (a) 16x16 Bayer dither, (b) 16x16 rotated dither and (c) error-diffusion (300 dpi).

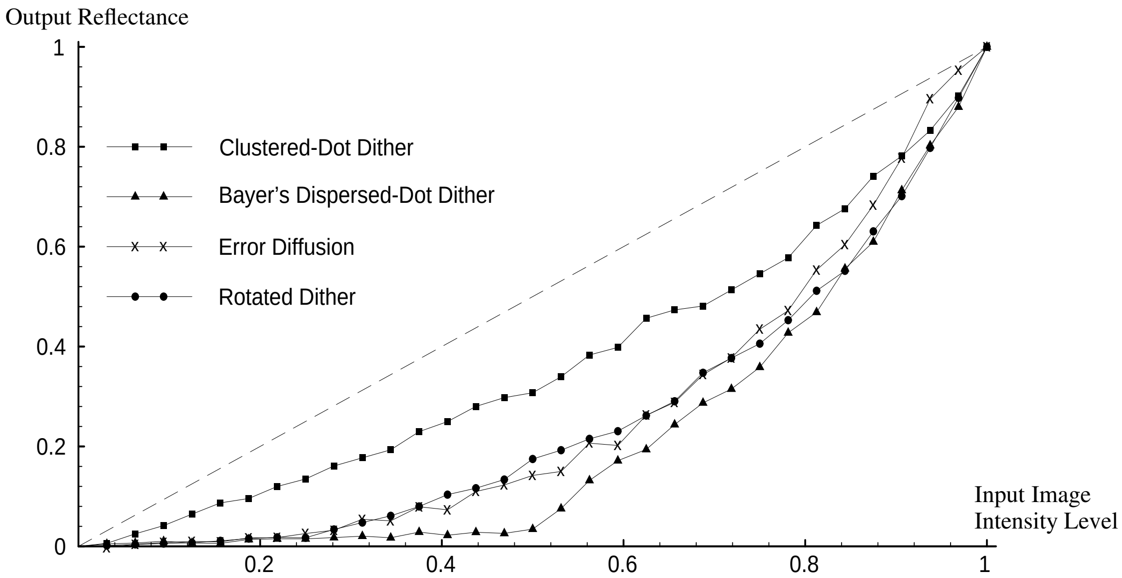

Figure 7: Tone reproduction curves obtained from density measurements of wedges printed at 300 dpi on a laser printer.

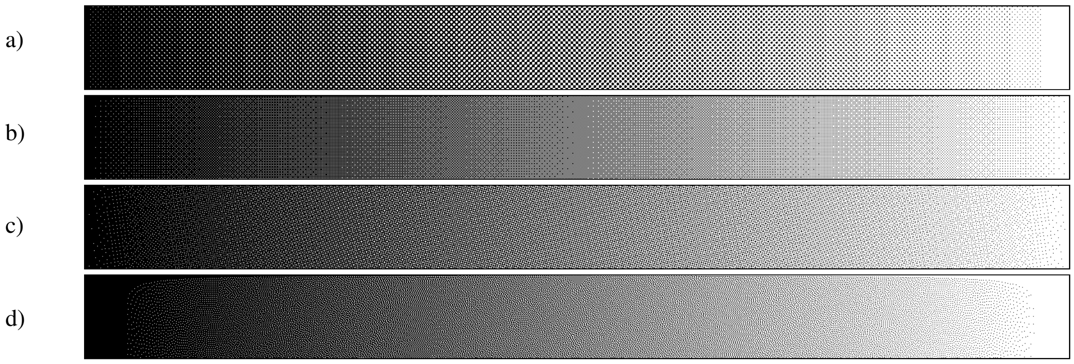

Figure 8: Tone reproduction behavior of a a grayscale wedge halftoned with (a) a conventional diagonally oriented clustered-dot dither array (dot area = 32 pixels), (b) Bayer's 16x16 dispersed-dot dither, (c) 16x16 rotated dispersed-dot dither and (d) error diffusion, printed at 300 dpi.

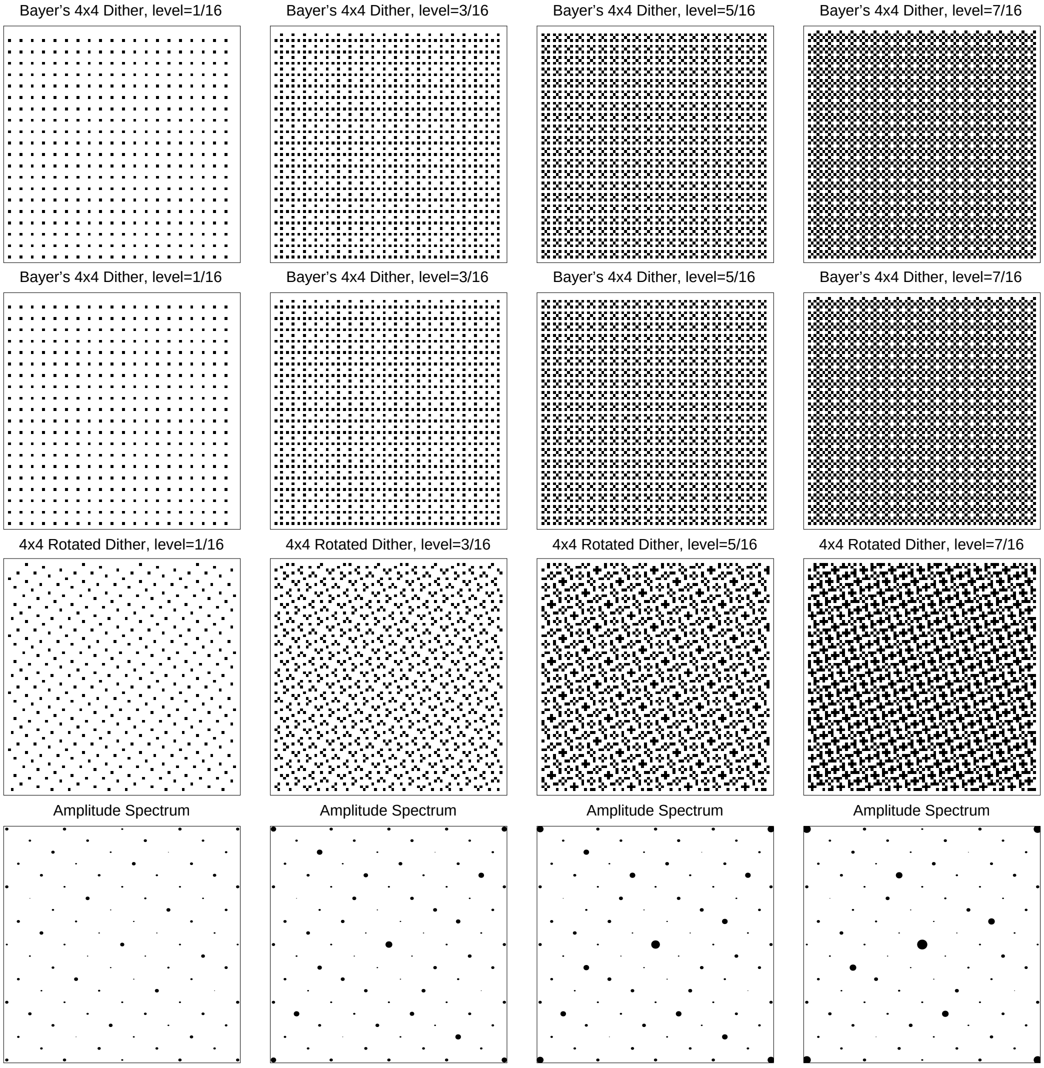

Figure 9: Halftone patterns and corresponding DFT impulse amplitudes at gray levels $\frac{1}{16}$, $\frac{3}{16}$, $\frac{5}{16}$ and $\frac{7}{16}$.

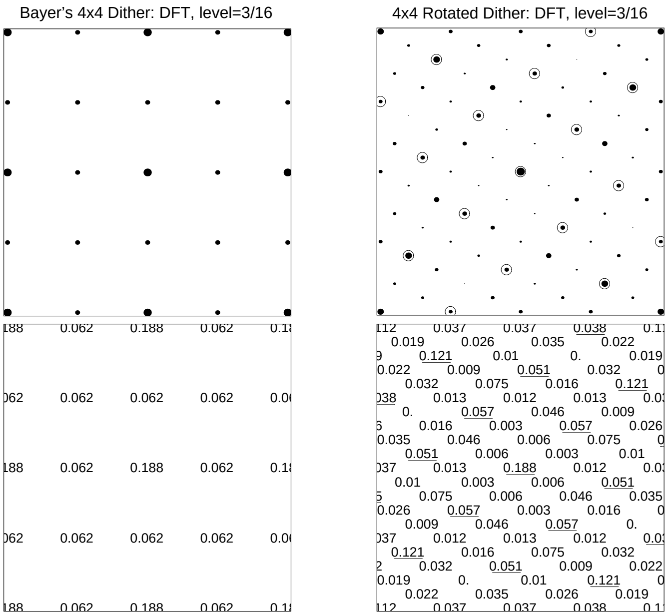

Figure 10: Amplitude spectrum at gray level $\frac{3}{16}$.