Frequency Analysis of Ordered Dither

Robert A. Ulichney

Digital Equipment Corporation
550 King St., LKG1-2/C13, Littleton, MA 01460-1289

Abstract

Ordered dither, the class of digital halftoning techniques which uses a periodic array of thresholds, is generalized for both rectangular and hexagonal grids, by means of the spatial method of recursive tessellation, a sub-tiling algorithm. The nature of the texture patterns so produced is illustrated and examined in the frequency domain, revealing several insights within and between the classes of rectangular and hexagonal grids. A simple explicit expression is derived which allows the use of the rectangular DFT to compute a hexagonal Fourier transform, maintaining continuous-space dimensions. Digitally produced examples are included.

1 Introduction

Digital halftoning addresses the problem of rendering continuous tone images on displays that are capable of only producing binary pixels. Several methods have been reported and survey papers are available [1,2,3,4]. This review essentially summarized the results found in [4].

Halftoning algorithms fall into two major categories: those that operate on only the pixel itself (point operations) and those that additionally work on the neighborhood around a pixel (neighborhood operations). Neighborhood techniques generally yield better looking halftoned pictures at the expense of more computation and storage. Examples of such techniques are 2-D pulse-density modulation [5], and various approaches to error diffusion [6,7,8,9]

In this paper, the point operation known as ordered dither is considered. Once a suitable threshold array is determined, the extreme simplicity of implementing halftoning by ordered dither makes it an important practical design candidate. An ordered dither algorithm generates a binary halftone image by comparing pixels from an original continuous-tone image to a threshold value from a deterministic, periodic array. The thresholds are historically referred to as "ordered" to contrast with methods whose thresholds are random.

Ordered dithering techniques can further be divided into two classes by the nature of the dots produced, clustered and dispersed. The most well known halftoning method, the classical or printer's screen, is generated by a clustered-dot ordered dither algorithm mimicking the optical process used for over 100 years in the printing industry. The minimum size of a clustered dot has traditionally been limited by the minimum printing plate area that will hold a dot of ink.

When the image is to be produced on a device that can faithfully display every binary pixel, the preferred choice of ordered dither is one that generates dispersed rather than clustered dots. Dispersed-dot threshold arrays yield high frequency fidelity and illusions of constant gray regions better than do clustered-dot arrays of the same resolution and threshold array period.

The design of dispersed-dot ordered dither threshold arrays has been studied on square grids by Limb [10], Lippel [11,12], and most notably by Bayer. Halftoning with a particular homogeneous threshold matrix has become known as "Bayer's dither" after his famous 1973 proof [13] of optimality with respect to minimizing low frequency texture. In fact, in a wide sense, "ordered dither" has come to mean Bayer's dither. More recently, ordered dither on hexagonal grids has been independently studied in two efforts [14,15].

Although the design of optimally homogeneous dispersed-dot ordered dither threshold arrays can be approached in many ways, in this paper a spatial domain method of recursive tessellation is presented which generalizes the design for both rectangular and hexagonal grids and for all possible periods; two classes of period shape are possible and will be referred to as "even" and "odd".

The most characteristic aspect distinguishing dithering methods is the way areas of constant gray are presented; it is here that certain patterning becomes disturbing. To better understand and examine this feature, dither patterns will be explicitly analyzed and illustrated in the frequency domain. Special consideration will be given to the fact that visual frequency response is in fact anisotropic with orientation. It will be argued that for rectangular grids, only odd period arrays should be used; but if the choice between grid type is available, a hexagonal grid is superior.

Since this effort includes non-rectangular grids, the next section focuses on defining some characteristics of sampling grids. This will be especially used in section 4 in deriving relations for the continuous-space Fourier transform.

1.1 Periodic Sampling Grids

A periodic sampling grid can be expressed as a two dimensional impulse train,

$$\sum_{\mathbf{n}}^{\infty} \delta(\mathbf{x} - \mathbf{V}\mathbf{n}), \tag{1}$$

where the Sampling Matrix, $\mathbf{V} = [\mathbf{v}_1 \vdots \mathbf{v}_2]$, is composed of two linearly independent Sampling Vectors,

$$\mathbf{v}_1 = \begin{bmatrix} v_{11} \\ v_{21} \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} v_{12} \\ v_{22} \end{bmatrix} \tag{2}$$

with reference coordinate system, $\mathbf{x}$, and index vector, $\mathbf{n}$.

Figure 1 shows a periodic sampling grid in its most general form. The sampling vectors, $\mathbf{v}_1$ and $\mathbf{v}_2$, can be thought of as grid generating vectors.

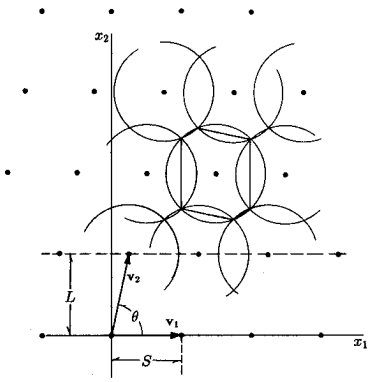

Figure 1: General periodic sampling grid with associated pixel shape.

Image data is almost always organized in lines for digitizing, storing, and displaying. It is convenient to refer to the coordinates, x_1 and x_2, as the *sample* and *line* directions, respectively. Most often (but not always) the sample axis refers to the horizontal direction, and the line axis refers to the vertical dimension. The sample period, S, is the distance between grid points in the sample direction, and the line period, L, is the distance between lines. In terms of the sampling vectors, $S = |\mathbf{v}_1|$ and $L = |\mathbf{v}_2| \cos\theta$, where θ is the angle between $\mathbf{v}_1$ and $\mathbf{v}_2$.

Figure 1 also illustrates a natural way of defining pixel shape. Pixel shape is defined as the smallest circumscribing polygon about a given grid point constructed from the perpendicular bisectors of lines between that point and all other grid points. It should be noted that the area of such pixels is always $S \times L$ regardless of its shape.

A general periodic sampling grid, that is, one with $\mathbf{v}_2$ uncontrained, has two shortcomings. Firstly, there is a lack of symmetry in the pixel shape and in the neighborhood surrounding it. Secondly, the offset for every line on an orthogonal coordinate system is different. This fact adds considerable complexity to simple image operations like cropping and scaling, or to the design of displays.

Semiregular grids are defined as Periodic sampling grids whose corresponding pixel shapes are symmetric about at least two axes. There are two classes of semiregular grids:

1. Rectangular grids, where $\mathbf{v}_1 \cdot \mathbf{v}_2 = 0$, (that is, $\mathbf{v}_1 \perp \mathbf{v}_2$).
2. Semiregular hexagonal grids, where $\mathbf{v}_1 \cdot \mathbf{v}_2 = |\mathbf{v}_1|^2/2$, (or $|\mathbf{v}_2| \cos\theta = |\mathbf{v}_1|/2$).

Note that semiregular hexagonal grids require only one offset of exactly $S/2$ for every other line on an orthogonal coordinate system. Such grids have been referred to as "offset sampled" or "quincuncial". The familiar case of rectangular grids requires no offset.

A useful metric which completely describes the shape of pixels on semiregular grids is aspect ratio. Aspect ratio is defined as the ratio of the sample period, S, to the line period, L. That is, $\alpha = S/L$.

The shape of the pixel on a rectangular grid is a regular polygon, a square, at only one aspect ratio, $\alpha = 1$. The shape of pixels on semiregular hexagonal grids is much more interesting. There are two cases where the pixel is a regular hexagon, for $\alpha = \frac{2}{\sqrt{3}}$ (hexagonal grid of the first kind) and $\alpha = 2\sqrt{3}$ (hexagonal grid of the second kind), and one special case where it is square, $\alpha = 2$.

2 Method of Recursive Tessellation

In this section, a method for deriving optimally homogeneous threshold arrays for regularly shaped periods of grid points (tiles) on regular grids is presented.

2.1 Tessellating Regular Grids

Two cases will be worked out in detail, rectangular with $\alpha = 1$ and hexagonal with $\alpha = \frac{2}{\sqrt{3}}$, hereby referred to as the "rectangular" and "hexagonal" cases. The results can later be applied to the remaining two cases by rotation. Rotate the hexagonal result by 90° (or 30°) for hexagonal grids of the second kind ($\alpha = 2\sqrt{3}$); for hexagonal grids with $\alpha = 2$, rotate the rectangular result by 45°.

The fundamental period or tile of the threshold array can have an integral power of 2 elements for the rectangular case, and an integral power of 3 elements for the hexagonal case. The power of 2 or 3 is defined as the *order*, η, of the array. Figure 2 shows the first 8 orders of rectangular tiles and Figure 3 shows the first 5 hexagonal orders. Families of periods are referred to as "even" or "odd" depending on η.

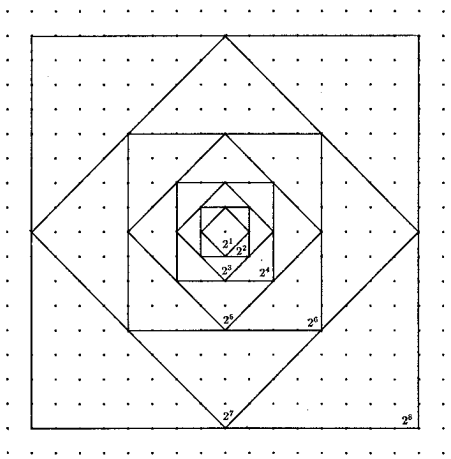

Figure 2: The first 8 orders of rectangular tiles. The number of elements unique to each tile is 2^η.

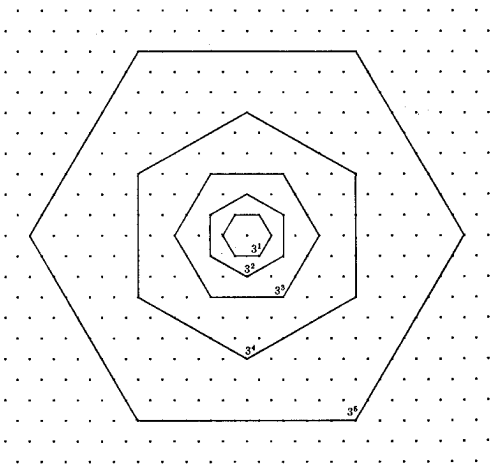

Figure 3: The first 5 orders of hexagonal tiles. The number of elements unique to each tile is 3^{η}.

Once a grid and period order is selected, all of two-space is tiled. Since rectangular tiles share each vertex with 4 other tiles, and each edge with 2 other tiles, only 1 vertex and 2 edges are unique to each tile. By a similar argument, only 2 vertices and 3 edges are unique to each hexagonal tile. The outlined edge and vertices in figure 4 show those points which are not part of the unique period of samples for examples of both odd and even orders.

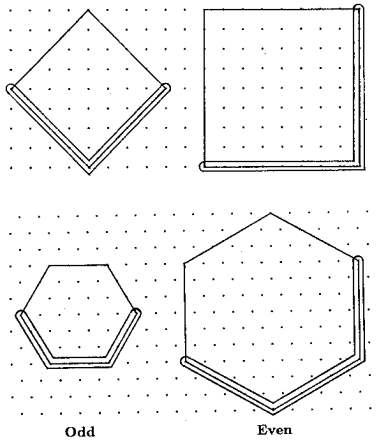

Figure 4: Odd and even spatial periods for rectangular grids (top) and hexagonal grids of the first kind (bottom). Boundary locations which are not part of the unique period are outlined.

2.2 Generation of Threshold Arrays

The goal is to order the samples from 1 to 2^{η} for the rectangular case or 1 to 3^{η} for the hexagonal case in such a way that as each successive position is numbered (turned on), the total two-dimensional ensemble of "on" positions is as homogeneously arranged as possible. When used as threshold arrays, the corresponding arrangement of output binary dots will be dispersed as homogeneously as possible for each gray level to be simulated.

The algorithm for generating the threshold array pivots on the fact that starting with a tile of a given order, η, the center (prime) point and all the vertices can act as center points for a retiling with tiles of order $\eta - 1$. The vertices of these tiles can further act as center points for another retiling of order $\eta - 2$, and so on. All of the vertices at each stage of this recursive tessellation are numbered before the next tessellation takes place. Breaking down the plane in this way provides a mechanism for locating the family of points that are exactly in the center of the voids between points of the higher order families.

The ordering of the family of points within a given stage or "subtessellation", is governed by an offset vector between the central (prime) point and any one of the vertices of its circumscribing tile. Figure 5 illustrates the process for a

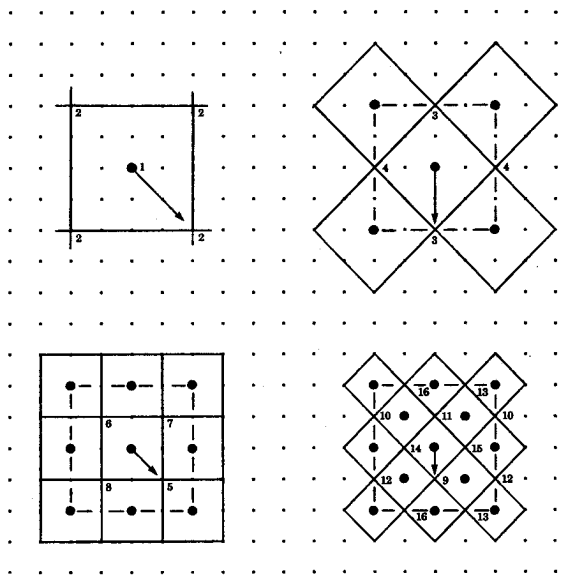

Figure 5: The 4 stages of recursive tessellation of a fourth order rectangular tile.

fourth order rectangular array. For an array of order η, there are η stages or subtessellations. At each stage, i, 2^i points are numbered by placing the tail of an offset vector on the center of each tile to locate one vertex in that tile. The offset vector can be initially oriented to point at any vertex, but must not change orientation within a given stage. Each vertex located with the offset vector is assigned a number equal to that of the center point plus the number of points assigned in the preceding stage, or 2^{i-1}.

The same procedure applies to the hexagonal case, except that each stage requires 2 passes, and the total number of points assigned in each stage is 3^i. The example illustrated in Figure 6 shows the method of recursive tessellation for a hexagonal tile of order 3. Two passes are required for each of the 3 stages. Note that on each second pass, the tail of the offset vector is placed on points assigned in the first pass instead of on tile centers.

Figure 7 shows the resulting rectangular threshold arrays for $\eta_r = 1$ through 6. These arrays are the same as described by Bayer [13]. The first 4 orders of hexagonal threshold arrays are shown in Figure 8. Recall that because the array edges are shared when tiling the plane, two edges of each rectangular tile and three edges of each hexagonal tile are not unique (Figure 4).

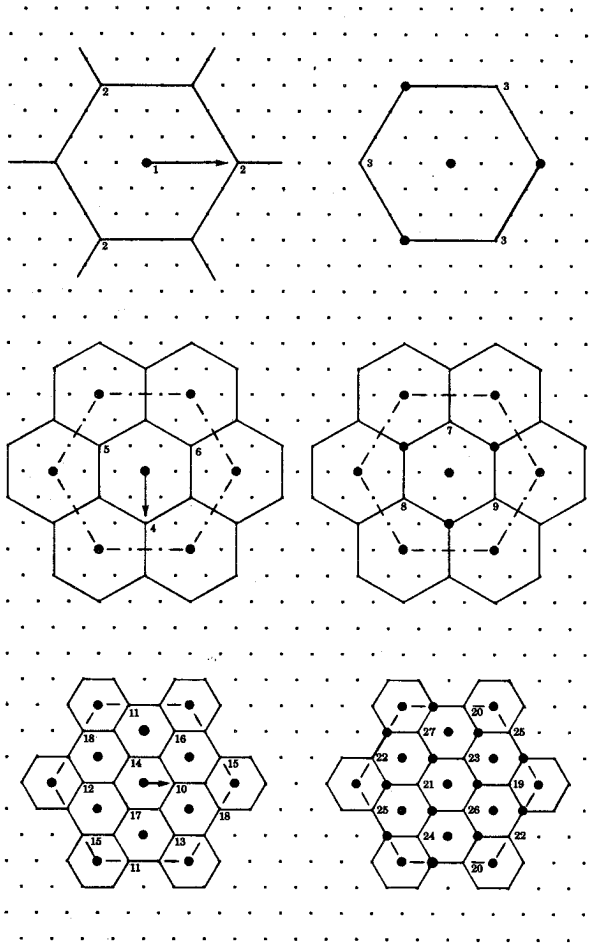

Figure 6: Hexagonal recursive tessellation of a third order hexagonal tile. Two passes are required for each of the three stages.

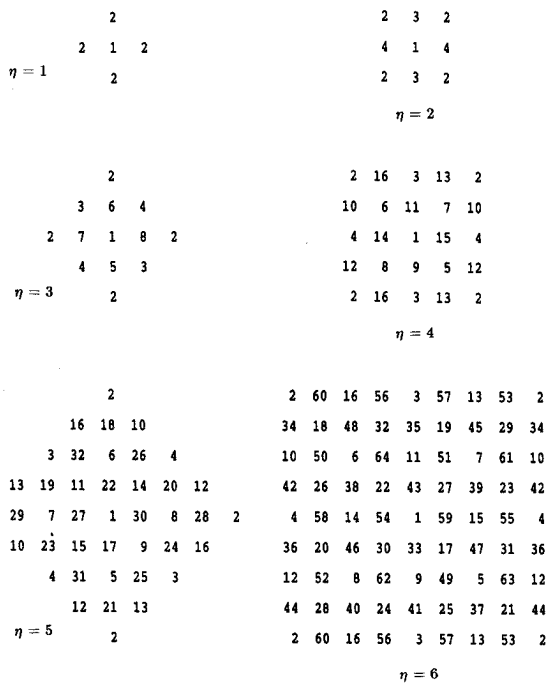

Figure 7: Rectangular threshold arrays. The edges of each array are shared in the tiling.

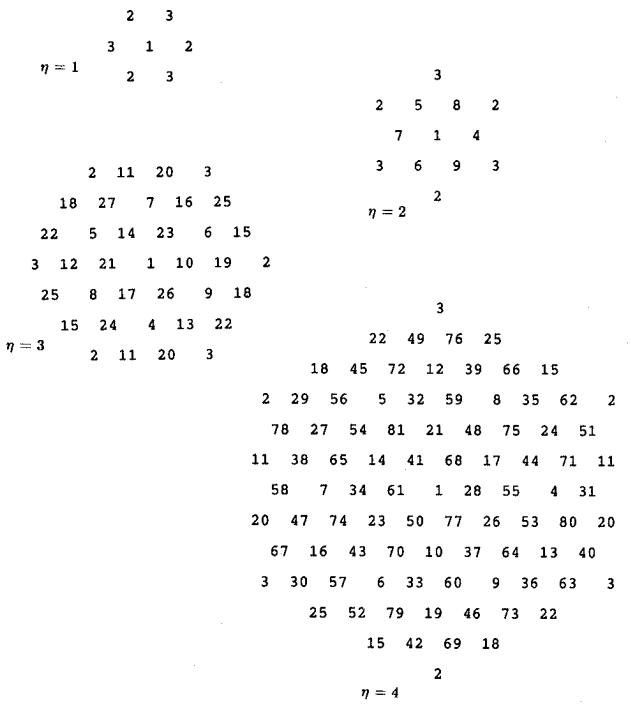

Figure 8: Hexagonal threshold arrays. The edges of each array are shared in the tiling.

As a practical matter, for operating on rectangularly shaped images, 2 periods of these arrays have to be packed together to form rectangularly periodic matrices. The exception, of course, is rectangular threshold arrays with η even, which are already rectangularly periodic. Figure 9 shows how this packing is done for odd ordered rectangular and hexagonal arrays, and Figure 10 illustrates this for even ordered hexagonal arrays. Note that rectangular tessellation of these packed matrices preserves the periodicity of the originally shaped tiles.

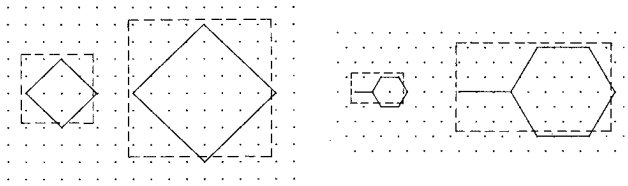

Figure 9: Packing two odd ordered periods for rectangular storage.

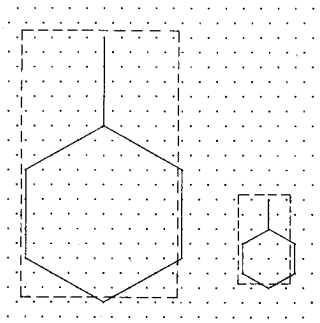

Figure 10: Packing two even ordered periods for rectangular storage.

3 Examples

To assure a fair comparison, each of the images shown in this section has an equal number of samples per unit area. Figure 11 shows the result of halftoning a wrapped one-dimensional gray scale ramp on a rectangular grid with each of the first 8 orders. For economy of space, only one fifth of the ramp (from 30% to 50% gray) is shown. The full ramps from 0 to 100% gray can be seen in [4]. In each case, $2^{\eta_r}+1$ gray levels (including white) are simulated. Similarly, the first 5 hexagonal orders are shown in Figure 12. Figure 13 and Figure 14 illustrate the result of halftoning a scanned image.

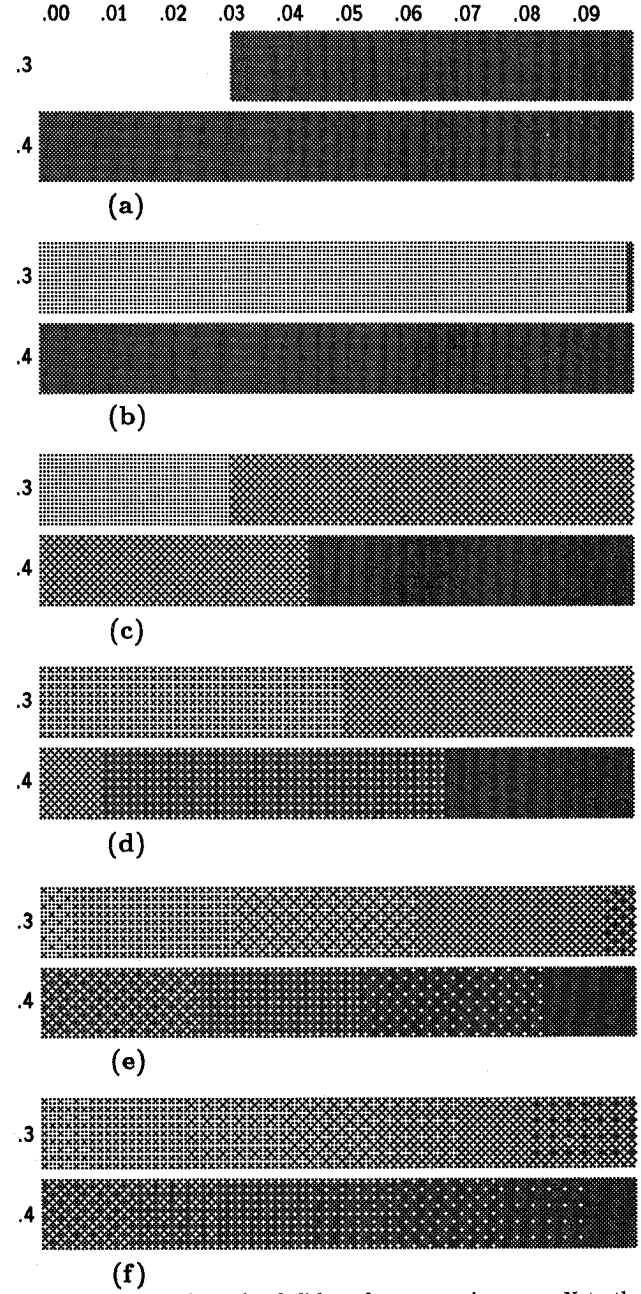

Figure 11: Rectangular ordered dither of a gray scale ramp. Note that only one fifth of the ramp (from 30% to 50% gray) is shown. (a) $\eta_r = 1$. (b) $\eta_r = 2$. (c) $\eta_r = 3$. (d) $\eta_r = 4$. (e) $\eta_r = 5$. (f) $\eta_r = 6$. (g) $\eta_r = 7$. (h) $\eta_r = 8$.

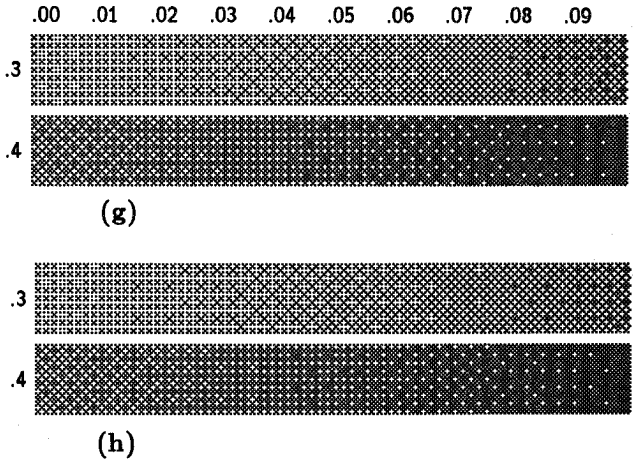

Figure 11 (continued)

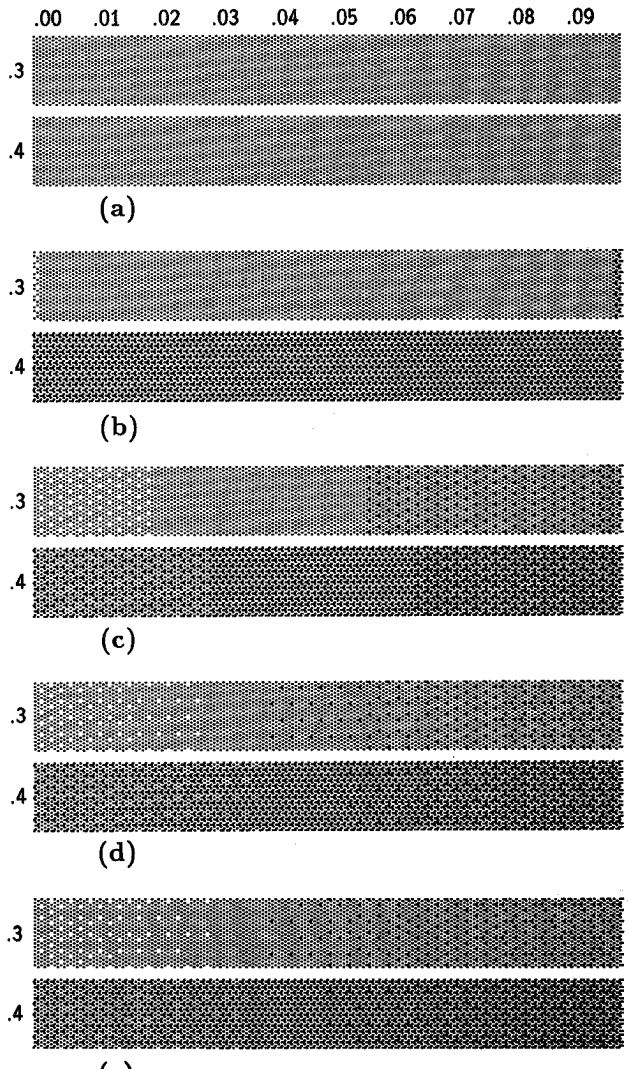

Figure 12: Hexagonal ordered dither of a gray scale ramp. Note that only one fifth of the ramp (from 30% to 50% gray) is shown. (a) $\eta_h = 1$. (b) $\eta_h = 2$. (c) $\eta_h = 3$. (d) $\eta_h = 4$. (e) $\eta_h = 5$.

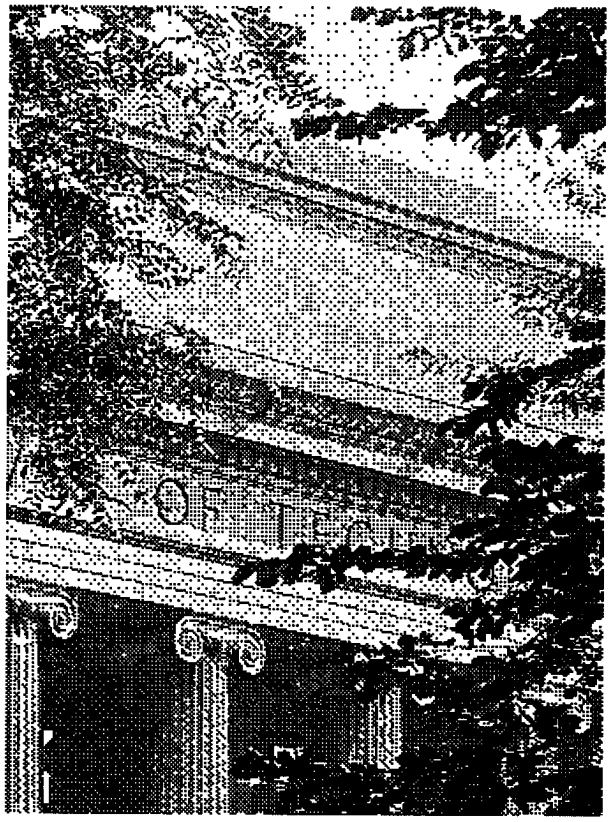

Figure 13: Rectangular ordered dither of a scanned picture. 257 (or 2^8+1) levels of gray, $\eta_r = 8$.

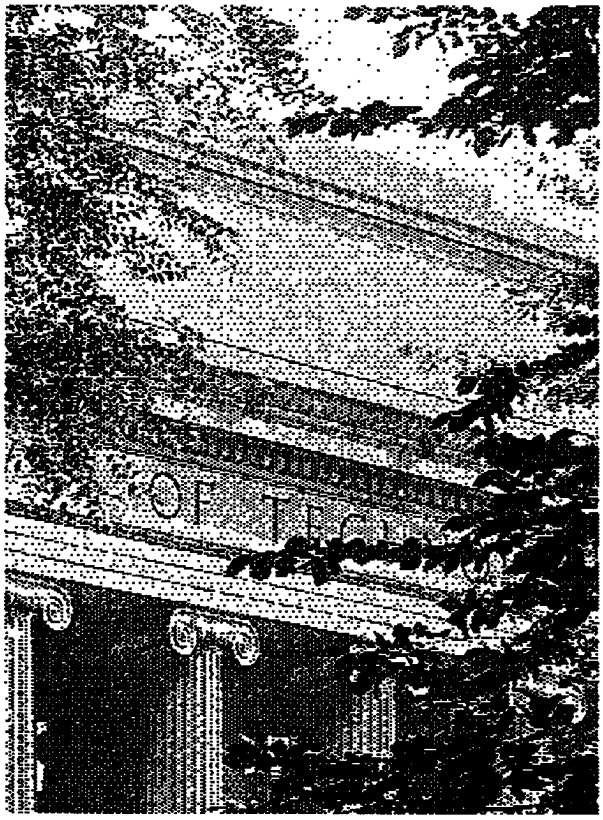

Figure 14: Hexagonal ordered dither of a scanned picture. 244 (or 3^5+1) levels of gray, $\eta_h = 5$.

These examples demonstrate the tradeoff between different threshold array orders. Note that "zeroth order" halftoning would correspond to using a single fixed threshold. For small orders, spurious contours due to texture changes in the patterns result from the insufficient number of gray levels. Images produced with large order threshold arrays suffer from the appearance of low frequency patterns in areas of uniform gray. The choice of optimum order, η, depends on the resolution in cycles/degree with which the image is to be displayed.

Upon comparing rectangular and hexagonal images produced with comparable number of gray levels, the images on the hexagonal grid possess patterns that are less disturbing than those on the rectangular grid. Since the dawn of halftoning over a century ago, printers have recognized that images looked better if the classical screen was oriented at a 45 degree angle. This phenomenon of anisotropic frequency response of the visual system has since been quantified. Studies have shown [16,17,18] that a sharp asymmetry exists with regard to our ability to detect fine gratings as a function of orientation; human vision is more acute for horizontal and vertical orientations than for screens oriented at a 45° angle.

The preponderance of horizontal and vertical patterns seen in rectangular ordered dither is not present in the hexagonal case, which probably explains its more pleasing appearance.

As was mentioned at the onset of this section, these results can be applied to the two other regular grids, hexagonal grids with $\alpha = 2\sqrt{3}$ and $\alpha = 2$. While rotating the rectangular results by 45° for application on a hexagonal grid with $\alpha = 2$ is straightforward, it may not be clear why the hexagonal results in this section produced on a hexagonal grid of the first kind need to be rotated 90° for application on a hexagonal grid of the second kind ($\alpha = 2\sqrt{3}$). Figure 15 illustrates why this rotation is necessary.

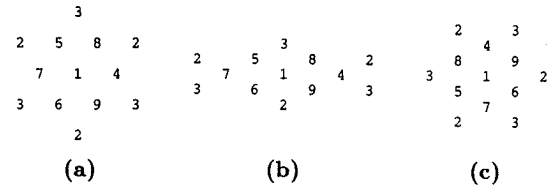

Figure 15: Importance of hexagonal kind. (a) Second order threshold array for a regular hexagonal grid of the first kind ($\alpha = \frac{2}{\sqrt{3}}$). (b) Incorrect use of the above data on a regular hexagonal grid of the second kind ($\alpha = 2\sqrt{3}$). (c) Correct translation of threshold values for a hexagonal grid of the second kind.

4 Fourier Analysis

The Fourier transform has been used in the past to evaluated halftone patterns. For example, it has been used to study classical halftone screens [19], and image induced artifacts due to interaction with clustered and dispersed-dot ordered dither [20] on square grids. In this section, a process for computing and displaying the Fourier transform of ordered dither in terms of the continuous-space dimensions of the grid are presented for the purpose of comparing rectangular and hexagonal methods of various orders.

The most characteristic feature of a halftone technique is the texture generated in areas of uniform gray. The rendition of high frequency detail depends primarily on how sharp the image was (or to what extent high pass filtering was performed) prior to halftoning. In general, some degree of presharpening will usually produce higher quality halftoned pictures.

The best measure of the virtues of a halftone algorithm, then, is its ability to render areas of uniform gray. To present insight into the nature of the periodic output of ordered dither "exposure plots" of composite Fourier transforms are introduced in this section.

The binary output from halftoning fixed gray level inputs by means of ordered dither will be periodic with the same spatial period as the threshold array. The spatial period will be specified by two vectors, $\mathbf{p}_1$ and $\mathbf{p}_2$, in terms of the spatial sampling vectors, $\mathbf{v}_1$ and $\mathbf{v}_2$ as described in section 1.

The spatial periods can be thought of as tiles which cover all of two-space. There are two basic types of period shape, even and odd; as mentioned in section 2, this corresponds to the order, η. Even periods are replicated by period vectors that are collinear with the sampling vectors,

$$\begin{aligned} \mathbf{p}_1 &= N\mathbf{v}_1 \\ \mathbf{p}_2 &= N\mathbf{v}_2, \end{aligned} \tag{3}$$

for some integer, N. Odd periods are defined by

$$\begin{aligned} \mathbf{p}_1 &= N(\mathbf{v}_1 + \mathbf{v}_2) \\ \mathbf{p}_2 &= N(\mathbf{v}_2 - \mathbf{v}_2). \end{aligned} \tag{4}$$

It is important to note that two odd periods on a rectangular grid and three odd periods on a hexagonal grid can always be packed into one even period.

The derivation of a general expression for the Fourier transform of periodic patterns with an even period is addressed in the following section.

4.1 Continuous-space Fourier Transform Computation

For nonrectangular grids, it is not immediately clear how a discrete Fourier transform can be computed, or what its dimensions in continuous frequency space are. So, a method of computation along with an explicit expression in continuous space is sought.

Probably the first formalization of the Fourier representation of nonrectangularly sampled spaces was by Petersen and Middleton in 1962 [21]. Mersereau later specifically addressed hexagonally sampled signals with his derivation of the Hexagonal Discrete Fourier Transform (HDFT) [22]. But this expression is complicated, primarily because hexagonal periods can not in general be rearranged to repeat in a rectangular (rhomboidal) fashion.

For the case where the number of elements in a hexagonal period, or any shaped period, on a general periodic sampling grid is a perfect square, the following is a proof inspired by the multidimensional sampling theorem [23] that the canonical rectangular DFT can be used to compute its Fourier transform.

The two-dimensional continuous space Fourier transform, $C(f_1, f_2)$, of an image, $c(x_1, x_2)$, can be expressed as

$$\begin{aligned} C(f_1, f_2) &= \mathcal{F}\{c(x_1, x_2)\} \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} c(x_1, x_2) e^{-j2\pi(f_1 x_1 + f_2 x_2)} dx_1 dx_2 \\ c(x_1, x_2) &= \mathcal{F}^{-1}\{C(f_1, f_2)\} \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} C(f_1, f_2) e^{j2\pi(f_1 x_1 + f_2 x_2)} df_1 df_2 \end{aligned}$$

or, more conveniently in vector form:

$$C(\mathbf{f}) = \mathcal{F}\{c(\mathbf{x})\} = \int_{-\infty}^{\infty} c(\mathbf{x}) e^{-j2\pi \mathbf{f}^\top \mathbf{x}} d\mathbf{x} \tag{5}$$

$$c(\mathbf{x}) = \mathcal{F}^{-1}\{C(\mathbf{f})\} = \int_{-\infty}^{\infty} C(\mathbf{f}) e^{j2\pi \mathbf{f}^\top \mathbf{x}} d\mathbf{f} \tag{6}$$

The units of the frequency components in these expressions are cycles/unit-length (as opposed to radians/unit-length).

One important identity that will be used is the transform of a two-dimensional impulse train:

$$\mathcal{F}\left\{\sum_{\mathbf{n}}^{\infty} \delta(\mathbf{x} - \mathbf{An})\right\} = \frac{1}{|\det \mathbf{A}|} \sum_{\mathbf{m}}^{\infty} \delta(\mathbf{f} - \mathbf{Bm}) \tag{7}$$

where $\mathbf{B}^\top \mathbf{A} = \mathbf{I}$,

$\det \mathbf{A} = a_{11}a_{22} - a_{12}a_{21}$,

and $\sum_{\mathbf{n}}^{\infty}$ represents $\sum_{n_1=-\infty}^{\infty}\sum_{n_2=-\infty}^{\infty}$.

Note that the spatial pixel area, SL, is equal to $|\det \mathbf{A}|$, the frequency "pixel area" is $|\det \mathbf{B}|$, and $|\det \mathbf{A}| = |\det \mathbf{B}|^{-1}$. The matrix $\mathbf{B}$ can be expressed explicitly as

$$\mathbf{B} = (\mathbf{A}^{-1})^\top = \frac{1}{\det \mathbf{A}} \begin{bmatrix} a_{22} & -a_{21} \\ -a_{12} & a_{11} \end{bmatrix}. \tag{8}$$

Recalling that $\mathbf{V}$ has been defined as the *spatial sampling matrix* (equations (1) and (2),

$$\mathcal{F}\left\{\sum_{\mathbf{n}}^{\infty} \delta(\mathbf{x} - \mathbf{Vn})\right\} = \frac{1}{|\det \mathbf{V}|} \sum_{\mathbf{m}}^{\infty} \delta(\mathbf{f} - \mathbf{Um}) \tag{9}$$

where $\mathbf{U} = [\mathbf{u}_1 \vdots \mathbf{u}_2] = (\mathbf{V}^{-1})^\top$ is the *frequency baseband replication matrix.*

Suppose $c(\mathbf{x})$ is a periodic array of weighted impulses on the sampling grid of equation (1) with an $N \times N$ rhomboidal shaped period. It could be expressed as

$$c(\mathbf{x}) = \left(\sum_{\mathbf{n}}^{N} I[\mathbf{n}]\delta(\mathbf{x} - \mathbf{Vn})\right) * \sum_{\mathbf{l}}^{\infty} \delta(\mathbf{x} - \mathbf{Pl}) \tag{10}$$

where "$*$" denotes two-dimensional convolution,

$$\sum_{\mathbf{n}}^{N} \text{ represents } \sum_{n_1=0}^{N}\sum_{n_2=0}^{N},$$

$I[\mathbf{n}]$ is a discrete-space image, and $\mathbf{P} = [\mathbf{p}_1 \vdots \mathbf{p}_2] = N\mathbf{V}$ is the *spatial period replication matrix.*

Note that

$$\mathcal{F}\left\{\sum_{\mathbf{l}}^{\infty} \delta(\mathbf{x} - \mathbf{Pl})\right\} = \frac{1}{|\det \mathbf{P}|} \sum_{\mathbf{k}}^{\infty} \delta(\mathbf{f} - \mathbf{Qk}) \tag{11}$$

where $\mathbf{Q} = [\mathbf{q}_1 \vdots \mathbf{q}_2] = (\mathbf{P}^{-1})^\top$ is the *frequency sampling matrix.* $\mathbf{Q}$ also equals $\frac{1}{N}\mathbf{U}$, since $\mathbf{P} = N\mathbf{V}$ and $\mathbf{U} = (\mathbf{V}^{-1})^\top$.

The Fourier transform of this image is

$$\begin{aligned} C(\mathbf{f}) &= \mathcal{F}\{c(\mathbf{x})\} \\ &= \left(\int_{-\infty}^{\infty} \sum_{\mathbf{n}}^{N} I[\mathbf{n}]\delta(\mathbf{x}-\mathbf{V}\mathbf{n})\, e^{-j2\pi \mathbf{f}^\top \mathbf{x}} d\mathbf{x}\right) \qquad (12) \\ &\quad \times \frac{1}{|\det \mathbf{P}|} \sum_{\mathbf{k}}^{\infty} \delta(\mathbf{f}-\mathbf{Q}\mathbf{k}) \\ &= \left(\sum_{\mathbf{n}}^{N} I[\mathbf{n}] e^{-j2\pi \mathbf{f}^\top \mathbf{V}\mathbf{n}}\right) \frac{1}{|\det \mathbf{P}|} \sum_{\mathbf{k}}^{\infty} \delta(\mathbf{f}-\mathbf{Q}\mathbf{k}) \qquad (13) \end{aligned}$$

This expression is zero everywhere except at $\mathbf{f} = \mathbf{Q}\mathbf{k}$, so the frequency term in the exponent becomes

$$\mathbf{f}^\top = \mathbf{k}^\top \mathbf{Q}^\top$$

where $\mathbf{Q}^\top = \frac{1}{N}\mathbf{V}^{-1}$ since $\mathbf{Q} = (\mathbf{P}^{-1})^\top$ and $\mathbf{P} = N\mathbf{V}$.

Thus, equation (13) becomes simply

$$C(\mathbf{f}) = \frac{1}{|\det \mathbf{P}|} \sum_{\mathbf{k}}^{\infty} \left(\sum_{\mathbf{n}}^{N} I[\mathbf{n}] e^{-j\frac{2\pi}{N}\mathbf{k}^\top \mathbf{n}}\right) \delta(\mathbf{f}-\mathbf{Q}\mathbf{k}) \qquad (14)$$

or,

$$C(\mathbf{f}) = \frac{1}{|\det \mathbf{P}|} \sum_{\mathbf{k}}^{\infty} I[\mathbf{k}]\delta(\mathbf{f}-\mathbf{Q}\mathbf{k})$$

where, in scalar notation,

$$I[\mathbf{k}] = I[k_1, k_2] = \sum_{n_1=0}^{N} \sum_{n_2=0}^{N} I[n_1, n_2] e^{-j\frac{2\pi}{N}(k_1 n_1 + k_2 n_2)} \qquad (15)$$

is recognized as the familiar two-dimensional discrete Fourier transform. Equation (14) satisfies the need for an explicit expression in continuous frequency space and a simple mechanism for computation.

For completeness, the inverse Fourier transform can be shown to be [4]

$$c(\mathbf{x}) = \sum_{\mathbf{n}}^{\infty} I[\mathbf{n}]\delta(\mathbf{x}-\mathbf{V}\mathbf{n}) = \left(\sum_{\mathbf{n}}^{N} I[\mathbf{n}]\delta(\mathbf{x}-\mathbf{V}\mathbf{n})\right) * \sum_{\mathbf{l}}^{\infty} \delta(\mathbf{x}-\mathbf{P}\mathbf{l}) \qquad (16)$$

where in scalar notation

$$I[\mathbf{n}] = I[n_1, n_2] = \frac{1}{N^2} \sum_{k_1=0}^{N} \sum_{k_2=0}^{N} I[k_1, k_2] e^{j\frac{2\pi}{N}(k_1 n_1 + k_2 n_2)} \qquad (17)$$

is the two-dimensional inverse discrete Fourier transform.

This applies to periodic grids in general. For the case of semiregular grids considered in this paper, the matrices involved can be expressed in terms of the sample and line periods.

For the rectangular case:

$$\mathbf{V} = \begin{bmatrix} S_r & 0 \\ 0 & L_r \end{bmatrix} \quad \mathbf{U} = \begin{bmatrix} (S_r)^{-1} & 0 \\ 0 & (L_r)^{-1} \end{bmatrix} \qquad (18)$$

$$\mathbf{P} = \begin{bmatrix} NS_r & 0 \\ 0 & NL_r \end{bmatrix} \quad \mathbf{Q} = \begin{bmatrix} (NS_r)^{-1} & 0 \\ 0 & (NL_r)^{-1} \end{bmatrix} \qquad (19)$$

For the semiregular hexagonal case:

$$\mathbf{V} = \begin{bmatrix} S_h & S_h/2 \\ 0 & L_h \end{bmatrix} \quad \mathbf{U} = \begin{bmatrix} (S_h)^{-1} & 0 \\ (-2L_h)^{-1} & (L_h)^{-1} \end{bmatrix} \qquad (20)$$

$$\mathbf{P} = \begin{bmatrix} NS_h & NS_h/2 \\ 0 & NL_h \end{bmatrix} \quad \mathbf{Q} = \begin{bmatrix} (NS_h)^{-1} & 0 \\ (-2NL_h)^{-1} & (NL_h)^{-1} \end{bmatrix} \qquad (21)$$

The associated vectors are displayed in figures 16 and 17. Recalling that equation (14) is only valid for periods

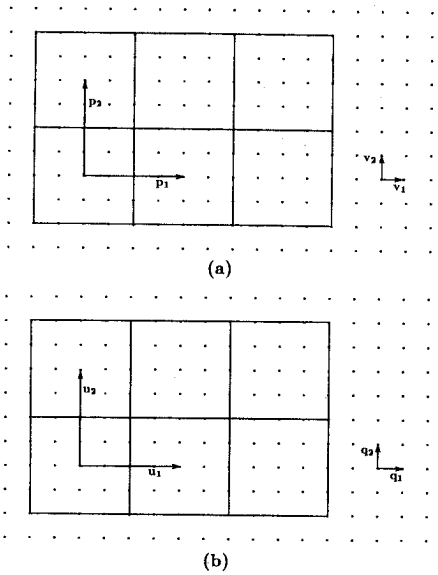

Figure 16: Rhomboidal tiling of a rectangular array coincides with the rectangularly shaped period when the period is even. (a) Spatial domain. (b) Frequency domain.

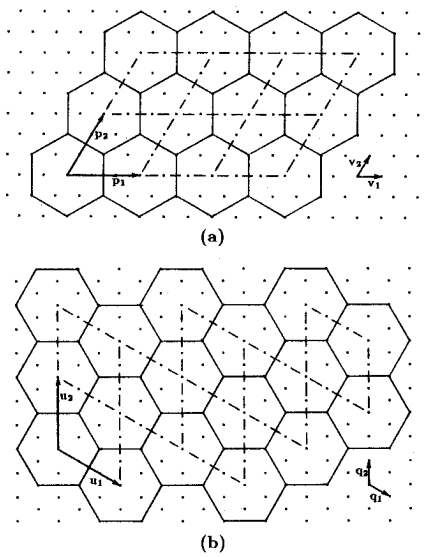

Figure 17: Rhomboidal tiling of a hexagonal array with an even period. (a) Rhomboidal and hexagonal tiling of equivalent data in the spatial domain. (b) Resulting rhomboidal tiling in the frequency domain; the baseband has hexagonal shape.

which have a perfect square number of elements, only even tiles meet that condition. This is most certainly true for rectangular tiles as shown in Figure 16, but Figure 17 illustrates why the condition can be met for hexagonal grids. The rhomboidal tiles shown cover the plane with precisely the same data as the hexagonal tiles.

In frequency space, the rhombus described by the vectors, $\mathbf{u}_1$ and $\mathbf{u}_2$, define how the Fourier transform is tiled but does not necessarily describe the shape of the baseband. In the hexagonal case (Figure 17), the baseband has the hexagonal shape as shown. It is interesting to note that when the period is even, a spatial hexagonal grid of the first kind has a transform on a frequency hexagonal grid of the second kind (and vice versa).

For the case where the period is odd, equation (14) can still be used by invoking the Similarity Theorem [24, p. 370]. Figure 18 shows how 2 odd rectangular tiles and 3 odd hexagonal tiles can be packed into a single even period.

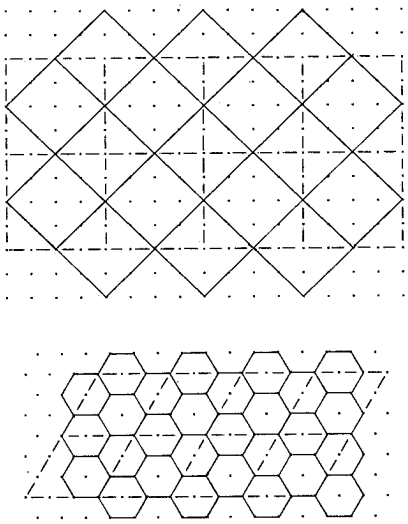

Figure 18: Packing odd periods into an even period.

The Similarity Theorem states that when exactly K periods are transformed as one period, each nonzero frequency component in the resulting DFT will be accompanied by K zero coefficients and have a magnitude K times as large as the DFT of a single period. So, packing odd tiles in this way only requires the additional step of dividing the resulting DFT by K.

This step is automatically handled by the normalization term in equation (14), $|\det \mathbf{P}|^{-1}$, where

$$|\det \mathbf{P}| = Z \times (\text{pixel area}) = (\text{spatial period area}).$$

where Z is the number of elements in the period used.

4.2 Composite Fourier Transform

The DFT is a strictly valid description of the Fourier transform *only* for periodic arrays of weighted impulses of which equation (10) is one. It is impossible for a signal to be both of finite spatial extent *and* truly band limited. When the DFT is applied to real signals, some degree of spectral overlap must be accepted. In well designed systems, this error is made inappreciable, but never zero.

In this section, the Fourier transform of the images resulting from halftoning a two-dimensional plane of one gray level are examined. $I[\mathbf{n}]$ in this case happens to consist of only 1's and 0's. Such images are precisely as described by equation (10), and thus the results are theoretically exact. (This could be the only practical use of the DFT that can make that claim.)

The output binary image is capable of rendering one of $Z+1$ gray levels for ordered dither with a Z element threshold array. The binary image, $I[\mathbf{n};g]$, resulting from halftoning a continuous-tone image consisting of one constant gray level, g, with a given threshold array has a DFT equation (15) denoted by $I[\mathbf{k};g]$.

The phase of $I[\mathbf{n}]$ relative to an origin is not important. Of interest is the magnitude of $I[\mathbf{k};g]$, which will provide insights into the relative distribution of energy in frequency space. To summarize this information over all gray levels for a given threshold matrix, an average is used. The *composite DFT* is defined as having frequency components,

$$I_\Sigma[\mathbf{k}] = \frac{1}{Z+1}\sum_{\text{all } g}|I[\mathbf{k};g]| \tag{22}$$

Plugging $I_\Sigma[\mathbf{k}]$ into equation (14) yields a specification of the location and magnitude of Fourier transform impulses. This *composite Fourier transform* is defined as

$$C_\Sigma(\mathbf{f}) = \frac{1}{|\det \mathbf{P}|}\sum_{\mathbf{k}}^{\infty} I_\Sigma[\mathbf{k}]\delta(\mathbf{f}-\mathbf{Q}\mathbf{k}) \tag{23}$$

Note that in all cases, the zero frequency term of $C_\Sigma(\mathbf{f})$ will be an impulse with area

$$\frac{I_\Sigma[\mathbf{0}]}{|\det \mathbf{P}|} = \frac{1}{2SL}, \tag{24}$$

since $I_\Sigma[\mathbf{0}]$ will always equal $Z/2$, the average number of black output pixels in a period.

A reasonable means of displaying $C_\Sigma(\mathbf{f})$, is with dots of an area proportional to the magnitude of the impulses. Such a display is similar to a photograph of an ideal optical Fourier transform consisting of points of light of different intensity; the resulting size of the exposed points would be proportional to their magnitude. Thus, the name "exposure plots" will be used to describe them.

4.3 Exposure Plots

Figure 19, displays exposure plots of $C_\Sigma(\mathbf{f})$ for $\eta_r = 1$ through 8 on a rectangular spatial grid ($\alpha = 1$). At the bottom of each exposure plot, a scale defining the actual dimensions (in cycles/unit-length) of the plot is terms of the original sample period, S, is provided.

Similarly, Figure 20, displays exposure plots for $\eta_h = 1$ to 5 where the sample grid was a regular hexagonal grid of the first kind ($\alpha = \frac{2}{\sqrt{3}}$). The composite transforms exhibit 4 fold symmetry for the rectangular case and 6 fold symmetry for the hexagonal case. The numerical values of $I_\Sigma[\mathbf{k}]$ for each of these cases is given in [4].

Consistent with the fair comparison approach of the last section, the number of spatial samples per unit-area are held constant, and the exposure plots are all generated with the same dimensional scale. Note that rectangular and hexagonal images with the same spatial pixel area have basebands of equal area.

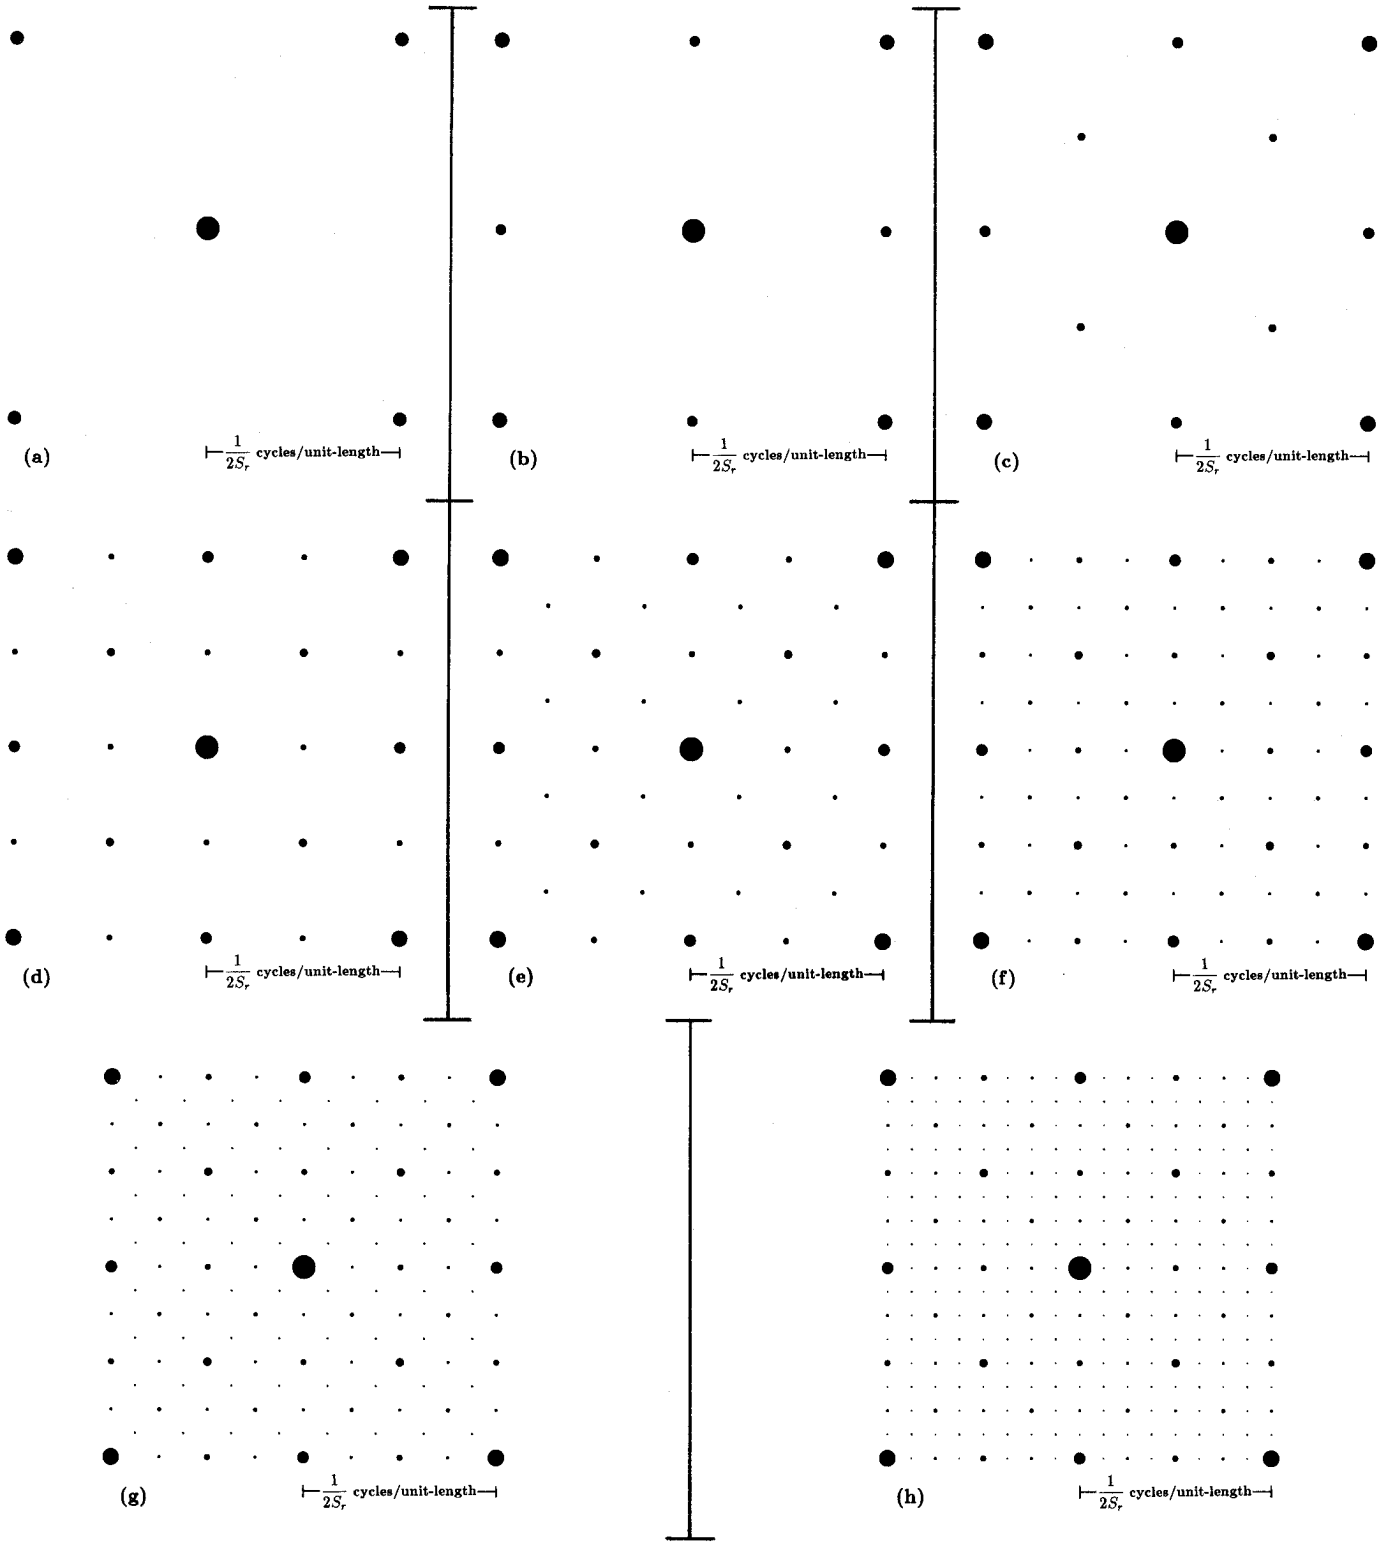

Figure 19: Composite Fourier transforms of rectangular ordered dither.
(a) Average of 3 patterns, $\eta_r = 1$.
(b) Average of 5 patterns, $\eta_r = 2$.
(c) Average of 9 patterns, $\eta_r = 3$.
(d) Average of 17 patterns, $\eta_r = 4$.
(e) Average of 33 patterns, $\eta_r = 5$.
(f) Average of 65 patterns, $\eta_r = 6$.
(g) Average of 129 patterns, $\eta_r = 7$.
(h) Average of 257 patterns, $\eta_r = 8$.

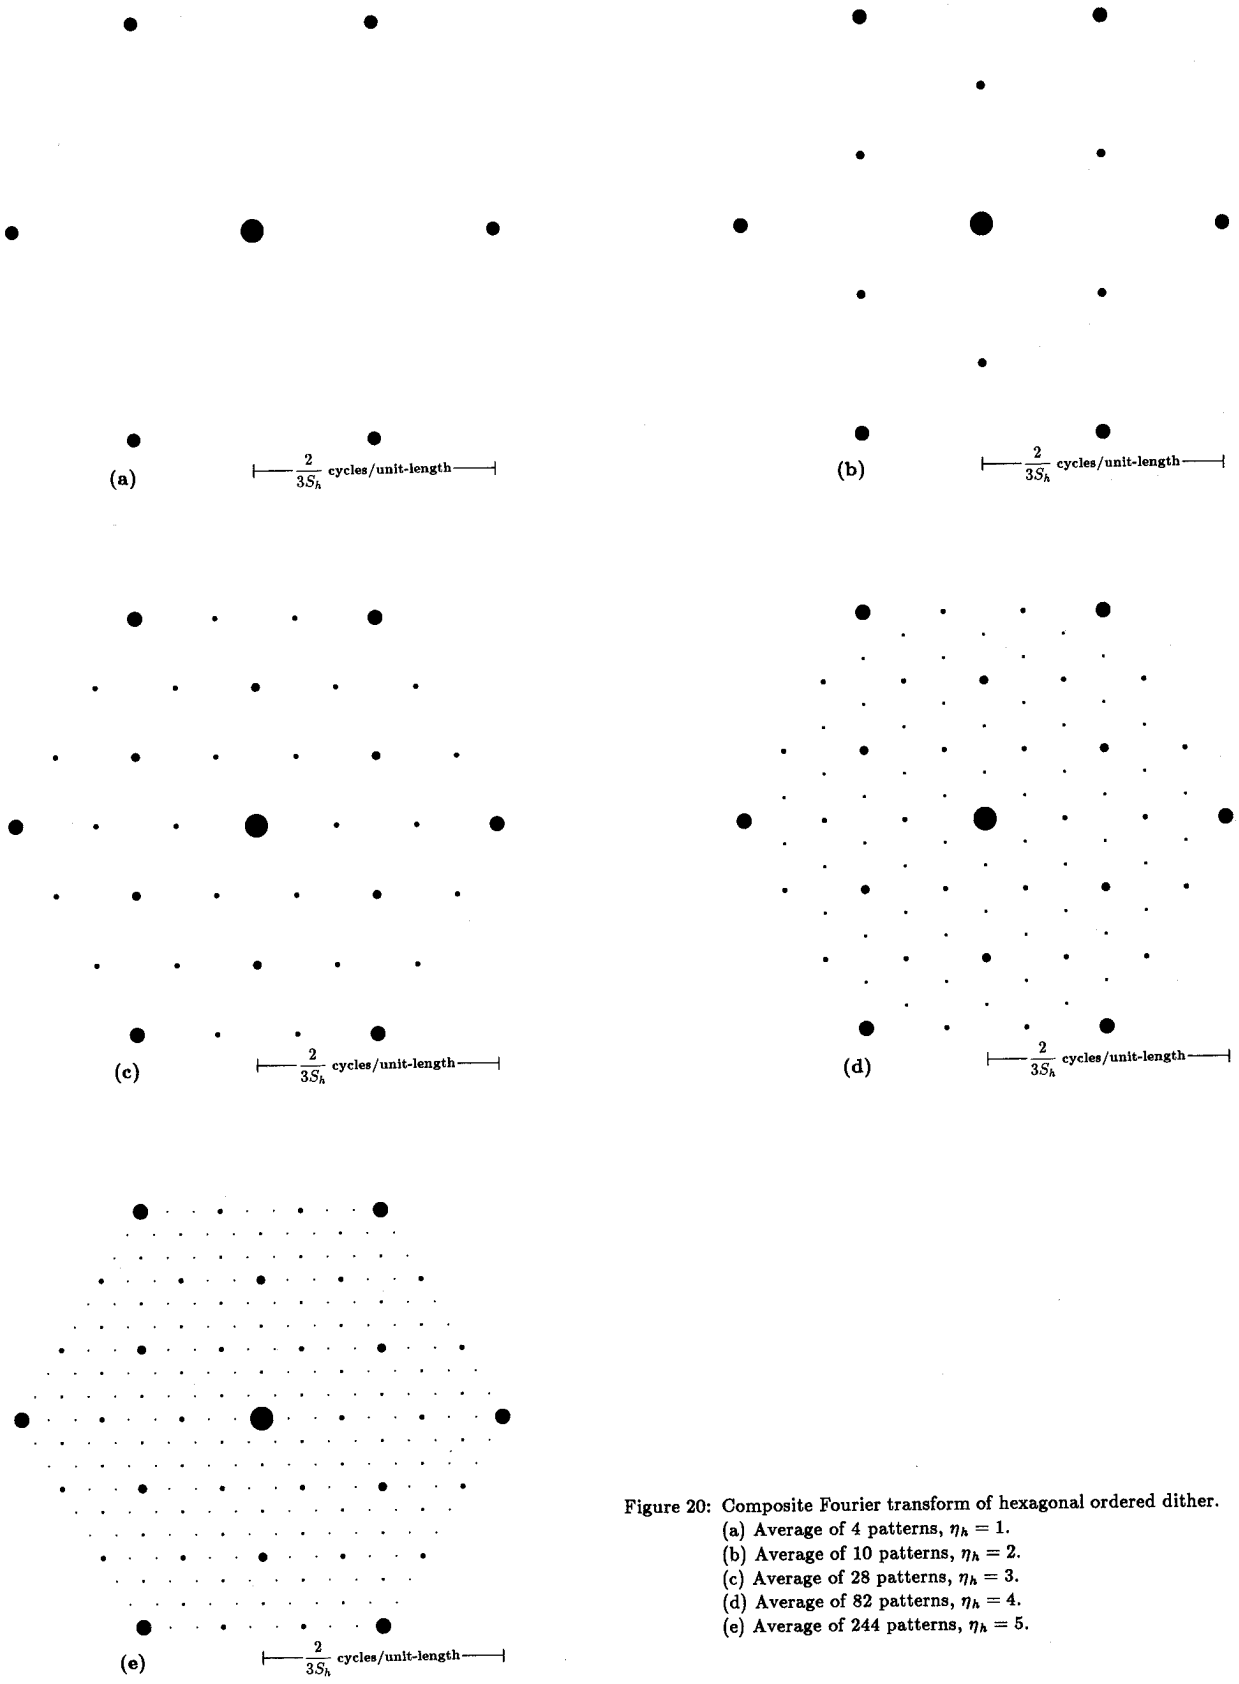

Figure 20: Composite Fourier transform of hexagonal ordered dither.
(a) Average of 4 patterns, $\eta_h = 1$.
(b) Average of 10 patterns, $\eta_h = 2$.
(c) Average of 28 patterns, $\eta_h = 3$.
(d) Average of 82 patterns, $\eta_h = 4$.
(e) Average of 244 patterns, $\eta_h = 5$.

The rectangular and hexagonal sample periods, S_r and S_h, used to define the scales in each exposure plot are constrained by the equal-pixel-area condition, $S_r L_r = S_h L_h$, and the aspect ratios used, $\alpha_r = 1$ and $\alpha_h = \frac{2}{\sqrt{3}}$. They are thus related as

$$\frac{S_r}{S_h} = \sqrt{\frac{\sqrt{3}}{2}}.$$

4.3.1 Analysis

The exposure plots of figures 19 and 20 beautifully illustrate the results of the method of recursive tessellation.

For $\eta = 1$, the only location of energy in the frequency domain besides the zero frequency term is at the "corners" of the baseband. This corresponds to the highest frequency which the grid is capable of accommodating. For the rectangular grid, it is the checkerboard pattern. Hexagonal grids have two "highest frequency" patterns which are negatives of one another. The patterns occur when either 1 or 2 of the 3 elements of a first order threshold array are "on".

For every order, η, these high frequency corners are second only to the zero frequency term in magnitude. This is exactly what one would expect from a good dispersed-dot ordered dither algorithm.

As η, and thus the number of elements in the threshold period, increases, the additional frequency terms in the Fourier transform assume positions in between those of smaller η, with proportionally less energy associated with them. The size of the baseband, however, always remains fixed.

Observe that families of frequency coefficients of the same magnitude are arranged exactly as the families of threshold points within the stages of the method of recursive tessellation! Families of larger dots correspond to earlier stages.

Also, the gray-level-number vs. low-frequency-texture tradeoff demonstrated in the examples of section 3 can be explained in the frequency domain. As the number of gray levels increases, so does the number of coefficients. As coefficients get closer to the zero frequency location, textures of lower frequency appear in the image.

On high spatial resolution displays, it is the set of frequency terms closest to the zero frequency center that determine whether or not any patterns will be perceived. Redrawing the data given by Taylor [18], in polar form would reveal a somewhat square shaped frequency threshold plot with the cusps oriented along the horizontal and vertical axes. Any nonzero frequency components inside this perceptual mask would be seen as a periodic pattern; components outside this area would be spatially integrated to yield a sensation of a solid average gray value.

So, in observing the exposure plots for rectangular grids (Figure 19) for any even ordered period, η, the *same* susceptibility of the human visual system to perceive low frequency periods exists for the next higher odd ordered period, $\eta + 1$, which can display twice as many gray levels! Therefore, from this examination in the frequency domain it can be argued that only odd rectangular ordered dither arrays should used.

The superior radial symmetry of hexagonal grids is also evident in the frequency domain. Energy is distributed more isotropically. While the corners of the rectangular Fourier transforms are slightly further from the zero frequency term than the corners of the hexagonal Fourier transforms, the hexagonal transforms have higher frequency components along both the horizontal and vertical directions where high frequencies are most needed due to the visual system's acute sensitivity there. The superposition of a square and hexagon of equal area in Figure 21 demonstrates these geometric characteristics.

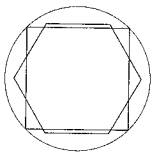

Figure 21: Superposition of square and regular hexagonal basebands of equal area corresponding to equal sample densities. Hexagonal grids support higher horizontal and vertical frequencies, but the highest frequencies are accommodated on rectangular grids.

5 Concluding Remarks

In this paper, dispersed-dot ordered dither has been generalized for application on both rectangular and hexagonal grids. The first 8 rectangular and the first 5 hexagonal threshold array period sizes, or orders, have been completely examined. The set of generated examples serve as a guide for evaluating the tradeoff between gray level number capability and low frequency texture.

The nature of the patterns resulting from halftoning uniform planes of constant gray level were evaluated in the frequency domain for regular rectangular and regular hexagonal grids of the first kind with equal sample densities. Insight was gained by displaying the frequency data in the form of "exposure plots". It was shown that the rectangular DFT can be used to find the Fourier transform of any periodic pattern of impulses on any periodic grid when the number of elements in the period is a perfect square.

The usual arguments in favor of hexagonal grids over rectangular girds are: (1) For circular dots of size just big enough to ensure complete coverage, dot overlap is reduced by 13.4%; less overlap reduces tone scale error. (2) For circularly bandlimited images, alias-free sampling can occur at a lower density. It has also been shown [4] that images represented on semiregular hexagonal grids are much more immune to variation in grid aspect ratio than are rectangular grids with respect circular dot covering efficiency.

The advantages of hexagonal displays can be realized in most cases without converting all storage formats in an imaging system to hexagonal form. For the case where a rectangularly sampled continuous-tone image is to be produced on binary displays of lower resolution, subsampling could be performed in a hexagonal fashion before halftoning. The hardware modification needed to convert a rectangular display to a semi-regular hexagonal one is quite simple; an alternate line offset of one half sample period is all that is required.

References

[1] Jarvis, J.F., 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*, vol. 5, pp. 13–40, (1976).

[2] Roetling, P.G., "Binary approximation of continuous-tone images", *Photographic Science and Engineering*, vol. 21, pp. 60–65, (1977).

[3] Stoffel, J.C. and J.F. Moreland, "A survey of electronic techniques for pictorial reproduction", *IEEE Tran. Commun.*, vol. 29, 1898–1925, (1981).

[4] Ulichney, R.A., *Digital Halftoning*, Cambridge: The MIT Press, (1987).

[5] Eschbach, R. and R. Hauck, "Binarization using a two-dimentional pulse-density modulation", *J. Opt. Soc. of Am.*, vol. 4, no. 10, pp. 1873–1878, (1987).

[6] Floyd, R.W., and L. Steinberg, "An adaptive algorithm for spatial greyscale", *Proc. SID*, vol. 17/2, pp. 75–77, (1976).

[7] Stucki, P., "MECCA -a multiple-error correcting computation algorithm for bilevel image hardcopy reproduction", Research Report RZ1060, IBM Research Laboratory, Zurich, Switzerland, (1981).

[8] Stevenson, R.L. and G.R. Arce, "Binary display of hexagonally sampled continuous-tone images", *J. Opt. Soc. Am. A*, vol. 2, no. 7, pp. 1009–1013, (1985).

[9] Ulichney, R.A., "Dithering with Blue Noise", *Proc. IEEE*, vol. 76, no, 1, Jan. (1988).

[10] Limb J.O. "Design of dither waveforms for quantized visual signals", *Bell Sys. Tech. J.*, Sep., pp. 2555–2582, (1969).

[11] Lippel, B. and M. Kurland, "The effect of dither on luminance quantization of pictures", *IEEE Trans. Comm.*, vol. COM-19, no. 6, pp. 879–888, (1971).

[12] Lippel, B. "Two and three-dimensional ordered dither in bi-level picture displays", *Proc. SID*, vol. 17/2, pp. 115–121, (1976).

[13] Bayer, B.E., "An optimum method for two level rendition of continuous-tone pictures", *Proc. IEEE Int. Conf. Commun., Conference Record*, pp. (26-11)–(26-15), (1973).

[14] Ulichney, R.A., "Generalized ordered dither", M.I.T., ATRP-T-51. *also* Digital Equipment Corporation, DEC-TR-412, (1985).

[15] Gardos,T., G.R. Arce, and J.P. Allebach, "The optimal ordered dither cell for the binary representation of continuous tone images on a hexagonal lattice", *Proc. of the 1986 Conf. on Inforamtion Sciences and Systems*, Princeton Univ., March (1986).

[16] Campbell, F.W., J.J. Kulikowski, J. Levinson, "The effect of orientation on the visual resolution of gratings", *J. Physiology London*, vol. 187, pp. 427–436, (1966).

[17] Higgins, G.C. and K. Stultz, "Visual acuity as measured with various orientations of a parallel-line test object", *J. Opt. Soc. Am.*, vol. 38, no. 9, pp. 756–758, (1948).

[18] Taylor, M.M., "Visual discrimination and orientation", *J. Opt. Soc. Am.*, vol. 53, June, pp. 763–765, (1963).

[19] Kermisch, D. and P.G. Roetling, "Fourier spectrum of halftone images", *J. Opt. Soc. Am.*, vol. 65, pp. 716–723, (1975).

[20] Allebach, J.P. and B. Liu, "Analysis of halftone dot profile and aliasing in the discrete binary representation on images", *J. Opt. Soc. Am.*, vol. 67, pp. 1147–1154, (1977).

[21] Petersen, D.P. and D. Middleton, "Sampling and reconstruction of wave-number-limited functions in N-dimensional Euclidean spaces", *Information and control*, vol. 5, pp. 279–323, (1962).

[22] Mersereau, R.M., "The processing of hexagonally sampled two-dimensional signals", *Proc. IEEE*, vol. 67, no. 6, pp. 930–949, (1979).

[23] Dudgeon, D.E. and R.M. Mersereau, *Multidimensional Digital Signal Processing*, pp. 39–41, Englewood Cliffs, NJ: Prentice-Hall, (1984).

[24] Bracewell, R.N., *The Fourier Transform and Its Application*, New York: McGraw-Hill, (1978).