

A tracking-based general framework to image halftoning

ANAMITRA MAKUR AND M. R. RAGHUVVEER

Electrical Communication Engineering, Indian Institute of Science Bangalore 560012, India, amakur@ecc.iisc.ernet.in
Electrical Engineering, Rochester Institute of Technology, Rochester, NY 14623-0887, USA,
mrreec@ritvax.isc.rit.edu

Abstract

Existing approaches to digital halftoning of image are based primarily on thresholding. We propose a general framework for image halftoning where some function of the output halftone tracks another function of the input gray-tone. This approach is shown to unify most existing algorithms and to provide useful insights. Further, the new interpretation allows us to remedy problems in existing algorithms such as the error diffusion, and subsequently to achieve halftones having superior quality. The proposed method's very general nature is an advantage since it offers a wide choice of three filters and a update rule. An interesting product of this framework is that equally good, or better, halftones are possible to be obtained by thresholding a noise process instead of the image itself.

1. Introduction

An image, in its conventional sense, is a gray-level image that may contain various shades of gray from black to white. (An image that contains color is usually labeled *color image*; such image is not considered in this article.) Such an image is referred to as a *continuous-tone* or *gray-tone image*, and is seen in a photograph, on television screen, and more recently on computer monitors. A printed image in newspapers, books and magazines, on the other hand, does not belong to the earlier class. While a photographic paper or a television/computer screen is capable of producing various shades of gray, a printed paper is not. Since ink may only be put or not put on a particular location of the paper, essentially only black or white may be reproduced in a printing process. Therefore, a conversion method is used in any printing that converts a gray-tone image to a binary (black or white) image, often called a *two-tone image* or a *half-tone image* or simply a *halftone*. The process of conversion is known as *screening* in conventional printing, and as *halftoning* to the image processing community. A halftone, though a binary image, possesses the appearance of continuous shades of gray owing to the way the black and white dots are distributed. Hence the objective of halftoning is to create a halftone as much perceptually truthful to the gray-tone image as possible.

Traditional halftones were analog halftones generated using various screening methods. With the advent of digital image processing attention was turned to the development of digital halftoning methods which would provide a better perceived image quality than traditional halftone screening methods for a given dot density (printing resolution). This trend has been accompanied by an ever more dominating role of digital techniques in storage, transmission and presentation of visual information. As a result, today we find digital halftones in diverse fields, such as in a few printed materials (primarily some technical books and journals),

Work done at the Department of Electrical Communication Engineering, Indian Institute of Science, Bangalore

in computer printing utilities (e.g., PostScript), in computer displays (e.g., color image display on an IBM compatible), and in the voter identity cards of India. Digital halftoning is used extensively in multimedia utilities involving image representation for various display and printing devices.

2. Existing halftoning algorithms

Let the gray-tone image be $I(m, n)$, taking values between 0 (black) and 1 (white). A halftoning algorithm would generate a halftone $H(m, n)$ of the same size as the gray-tone image (one dot per pixel; halftones where the dot density is larger than pixel density is not considered in this paper). With a few exceptions, all proposed halftoning methods are based on *thresholding* (1 bit quantization). The input gray-tone image, $I(m, n)$, is compared with a threshold, $T(m, n)$, to decide the output as follows:

$$H(m, n) = \begin{cases} 0 & \text{if } I(m, n) < T(m, n) \\ 1 & \text{if } I(m, n) \geq T(m, n) \end{cases} \quad (1)$$

A description and evaluation of early halftoning methods may be found in the review by Stoffel and Moreland¹. We briefly present various types of halftoning algorithms below. Threshold based halftoning algorithms may be classified into two classes, *passive* and *active* halftoning². While in passive halftoning, the threshold $T(m, n)$ is independent of the gray-tone image, in active halftoning the threshold depends on $I(m, n)$.

2.1. Passive halftoning

If $T(m, n)$ is a constant, the simplest halftoning algorithm known as *fixed thresholding* results.

In some algorithms $T(m, n)$ is taken as a periodic signal. If the signal, called the *carrier* or the *screen function*, is designed to have slow variations, the algorithm is known as *electronic screening*. If the signal, called *dither*, is a pseudo-random noise pattern, the algorithm is referred as *ordered dithering*. While the former method produces clustered dots, the later method produces dispersed dots. Research is still being carried out on this type of algorithms towards choosing perceptually better dither signals, for example refer to the work of Mitsa and Parker in generating a blue noise dither³.

2.2. Active halftoning

If $T(m, n)$ depends on the local statistics of the gray-tone image, then *adaptive thresholding* is achieved. One example is the *constrained average thresholding*, where

$$T(m, n) = k_1 I_{av}(m, n) + k_2, \quad (2)$$

$I_{av}(m, n)$ is the average gray-tone intensity of a neighborhood of (m, n) , and k_1 and k_2 are constants.

In *error diffusion*, past quantization errors are subtracted from the threshold. The block schematic of the algorithm is shown in Figure 1. Here $e_1(m, n)$ is the error in converting a gray value $\hat{I}(m, n)$ to a binary value $H(m, n)$. This error passes through G , (a causal FIR) filter, to produce $e_2(m, n)$ which is then added to $I(m, n)$ to obtain $\hat{I}(m, n)$. Error diffusion derives its




FIG. 1. Error Diffusion

name from the fact that addition of filtered past error may be seen as diffusing the error to future pixels. $\hat{I}(m, n)$ is thresholded by a constant threshold 0.5 which may be viewed as thresholding $I(m, n)$ by an active threshold. Error diffusion algorithm is a very active area of research. Weissbach and Wyrowski, for example, worked on designing the filter⁴. Similar efforts towards optimal choice of the filter by Barnard may be found in⁵. Adapting the filter based on the gray value⁶ or on the local spectrum⁷ have also been attempted by Wong and Kumar respectively. Sullivan *et al* incorporated human visual model in the error diffusion algorithm⁸. Eschbach and Knox had a similar objective in providing edge enhancement in error diffusion halftones⁹. Their scheme essentially modifies the constant threshold of 0.5 to an adaptive threshold $0.5 + (\kappa - 1)I(m, n)$ for some constant κ .

2.3. Other algorithms

As mentioned earlier, a few recent halftoning algorithms do not conform to a single-pass thresholding of the gray-tone image. Peli proposed a recursive algorithm that works parallelly on the entire image, but requires a few iterations¹⁰. Such iterative algorithms have also been suggested by Kollias and Anastassiou¹¹, and by Kumar⁷. Kollias and Anastassiou have also introduced a neural network based halftoning algorithm which may be viewed as an iterative error diffusion algorithm with a non-causal filter¹². Scheermesser *et al*¹³ introduced a spectrum-based iterative algorithm which has been further refined by Mrusek *et al*¹⁴ belonging to the same group.

3. The general framework

In this paper we propose a framework for halftoning that is general enough to describe most existing techniques, yet subtle enough to provide insight into the halftoning process. A halftone should *look like* the gray-tone image. However, for equal dot density (one dot per pixel), a halftone is not expected to have all features the gray-tone image has. Instead, a halftone should be as close as possible to some function (such as, local average) of the gray-tone image. This function should select the features of the gray-tone that are desirable and possible for a halftone to reproduce. Let this function be $f_1(m, n)$.

Further, since a halftone contains high frequency noise while most gray-tone images don't, at least at high frequency a halftone need not resemble the gray-tone. Therefore, a function of the halftone should be used for comparison with the gray-tone as well. This function should select the perceptually relevant parts of a halftone. We denote this function by $f_2(m, n)$.

We therefore propose to have $f_h(m, n)$, a function of the output halftone, *track* as closely a possible $f_l(m, n)$, a function of the input gray-tone image. This tracking-based approach is natural from a control systems view point. Let $e(m, n)$ be the tracking error,

$$e(m, n) = f_h(m, n) - f_l(m, n) \quad (3)$$

The objective of the halftoning algorithm is to minimize the tracking error $e(m, n)$. Therefore the threshold $T(m, n)$ is updated based on $e(m, n)$ in such a fashion so that the tracking error is reduced. Now the thresholding is conventionally done on the gray-tone image $I(m, n)$. However, since in our scheme the threshold $T(m, n)$ is already a function of $I(m, n)$, it is *not necessary* that the gray-tone is to be thresholded. Instead, a noise process or a carrier signal may be thresholded. This may sound a bit far-fetched, since no known halftoning algorithm (to our knowledge) thresholds a signal unrelated to the gray-tone image. However, the claim may be justified as follows.

Since most of the gray-tone image energy is present in the lower frequencies, a halftone that looks like the gray-tone would have a similar low frequency spectrum. Consequently a halftone contains much more energy in the high frequency region than the gray-tone. The idea here is that high frequency noise is less bothersome to the eye while contributing to the breaking up of undesirable patterns. One way to introduce such energy is to add high frequency noise to the gray-tone and then to threshold it to produce a halftone. This is nothing but ordered dithering. Another way is to add a high frequency carrier, which leads to electronic screening. Therefore, adding noise (or, carrier) to the gray-tone image before thresholding with a constant threshold is common. This is equivalent to thresholding the noise (or, carrier) with an active threshold that depends on the gray-tone value. What we propose is just that; rather than thresholding the sum of the gray-tone and the noise (or, carrier), we threshold the noise (or, carrier) itself.

This approach to halftoning leads to the simple structure shown in figure 2. If we choose to threshold a random noise process X , the halftoning algorithm may be called the *noise thresholding algorithm*¹⁵. Alternatively we may choose to threshold the gray-tone image I as is conventionally done.

The *feedforward filter* defines a functional $f_l(m, n)$ on the set $\{I(m-i, n-j): (i, j) \in W_l\}$ with W_l being a set of pairs of integers including $(0, 0)$. Note that no causality requirement is necessary for W_l since the entire gray-tone image is available. The feedforward filter offers a choice about which part of the gray-tone image spectrum is more important to be retained.




FIG. 2. The Proposed Tracking Approach.

Typically, more importance is attached to the low frequency spectrum. This leads to a low pass feedforward filter which produces the local average of the gray values. The tracking approach then attempts to retain the local average by the halftoning process. Note that a feedforward filter is *not* essential; good halftones may be produced by simply tracking $I(m, n)$ by $f_H(m, n)$.

The *feedback filter* produces a functional $f_H(m, n)$ on the set $\{H(m-i, n-j); (i, j) \in W_H\}$. W_H is, again, a set of paired integers that satisfy the causality requirement for some scan order if only single-pass halftoning is considered. If, on the other hand, iterative halftoning is considered, W_H need not be causal^{7,12}. We shall assume single-pass halftoning for the time being. The feedback filter achieves on the halftone what the feedforward filter achieves on the gray-tone. Therefore, a natural choice is to have an identical filter for both the feedforward and the feedback path. This would imply that the halftoning system tries to produce a halftone having the part of the spectrum selected by the filter to be nearly identical to that of the gray-tone. The feedback filter involves little implementational complexity since it operates on a binary input.

While W_H has to be causal, it may or may not include $(0, 0)$. If W_H does not include $(0, 0)$ then the threshold updating block computes the threshold

$$T(m, n) = f_T(e(m, n)) \quad (4)$$

which is subtracted from the input value before quantizing it with the 1-bit quantizer. Without loss of generality we may assume the quantizer to have the decision threshold at 0.5. Hence, the halftone output is given by

$$H(m, n) = Q[Y - T(m, n)] = \begin{cases} 0 & \text{if } Y - T(m, n) < 0.5 \\ 1 & \text{if } Y - T(m, n) \geq 0.5 \end{cases} \quad (5)$$

where Y is the input to be thresholded, either $I(m, n)$ or X . The *threshold update function* f_T should be chosen such that $H(m, n)$ is made to track $I(m, n)$ in the sense of keeping $f_H(m, n)$ close to $f_I(m, n)$, or keeping $e(m, n)$ close to zero.

On the other hand, if W_H includes $(0, 0)$ then $f_H(m, n)$ cannot be computed unless $H(m, n)$ is known. It, however, is possible to decide whether $H(m, n)$ must be set to 0 or 1 for some (m, n) . In our proposed tracking approach the following are computed:

$$E_0 = |e(m, n)| \text{ with } H(m, n) = 0, \quad (6)$$

and

$$E_1 = |e(m, n)| \text{ with } H(m, n) = 1. \quad (7)$$

Now, $H(m, n)$ is set equal to 1 if $E_1 \leq E_0$, and is set to 0 otherwise. Thus, here too, $f_H(m, n)$ is closely following $f_I(m, n)$.

In the noise thresholding variation, the noise *shaping filter* is a filter which has no causality requirement since the noise process may be generated in advance. This filter may be used to produce any spectral shape of the random noise (colored noise). It may be 1-D or 2-D depending on the noise process and the desirable spectral properties. Use of blue noise is popular in halftoning³, and the same can be achieved by choosing a proper filter.

The proposed structure is very general in nature since it includes a wide choice of the feedforward filter, the feedback filter, the shaping filter, and the threshold update function, including *nonlinear* functions. We demonstrate it in the next section.

4. Realizing various systems

4.1. The feedback path and mean preservation

From control systems viewpoint it is known that better tracking would be achieved if the feedback path exists. We now discuss the role of the feedback path with reference to halftoning. Gray scale rendition is a desirable property of any halftoning algorithm. This is achieved if, for a constant gray-tone input, the mean of the resulting halftone is equal to this constant. We refer to this property as *mean preservation*.

In the conventional approach where $I(m, n)$ is thresholded, mean preservation may be achieved using the feedback path with the following choices:

$$\begin{aligned} f_f(m, n) &= I(m, n) \\ f_{ff}(m, n) &= H(m, n) \\ T(m, n) &= T(m-1, n) - e(m-1, n). \end{aligned} \quad (8)$$

Assuming the initial threshold $T(0, 0)$ to be zero, it follows that the updated threshold after m line-scanned samples would be

$$T(m, 0) = - \sum_{i=0}^{m-1} e(i, 0) \quad (9)$$

Therefore, the threshold value scaled by $\frac{1}{m}$ denotes the difference between the constant gray-tone and the mean of the halftone of the entire past. A large enough positive threshold T would force H to be 0, leading to a positive error e (since $I \geq 0$) and consequently to a smaller threshold value in the next instance. Thus, it is not possible for the threshold to grow either to a large positive or a large negative value. As a result, the mean of the halftone would always approach the gray-tone value. In fact, if the gray-tone value is a rational number $\frac{p}{q}$, then it may be shown that the above system would produce exactly p 1's out of the first q halftone samples (thereby preserving the mean), and this pattern would repeat thereafter. This example turns out to be the simplest form of error diffusion.

However, mean preservation is not possible without the feedback path. For any realistic choice of the feedforward filter, $f_f(m, n)$ is constant for a constant input. The error $e(m, n)$ is equal to $f_f(m, n)$ in absence of a feedback path. Consequently the threshold $T(m, n)$ is also a constant. Thus the halftoning algorithm reduces to fixed thresholding, which fails to preserve the mean.

This, however, is not the case in the noise thresholding approach. Both the *closed loop* (presence of the feedback path) and the *open loop* (absence of the feedback path) systems may achieve mean preservation. In an open loop system the threshold is set to that constant value which makes the expected value of the halftone equal to the gray-tone value. We shall assume a

mean preserving feedforward filter (magnitude response at zero frequency is unity). Let the constant gray-tone be i and the corresponding constant threshold be $T_o(i)$. Since $H(m, n)$ takes only values 0 or 1,

$$E\{H(m, n)\} = p(H(m, n) = 1) = \int_{T_o(i)}^{\infty} f_X(X) dX \quad (10)$$

where $f_X(X)$ is the probability density function (pdf) of the noise process. Therefore,

$$T_o(i) = F_X^{-1}(1-i) \quad (11)$$

where $F_X(X)$ is the cumulative distribution function of the noise process. For an uniform pdf from 0 to 1, it simplifies to $T_o(i) = 1-i$. For a Gaussian pdf with mean $\frac{1}{2}$ and variance 1, it simplifies to

$$T_o(i) = \begin{cases} \frac{1}{2} + \text{erf}^{-1}(1-2i) & \text{if } 0 \leq i \leq \frac{1}{2} \\ \frac{1}{2} - \text{erf}^{-1}(2i-1) & \text{if } \frac{1}{2} \leq i \leq 1 \end{cases} \quad (12)$$

and for a symmetric triangular pdf extending from 0 to 1, the threshold is given by

$$T_o(i) = \begin{cases} 1 - \sqrt{\frac{1}{2}} & \text{if } 0 \leq i \leq \frac{1}{2} \\ \sqrt{\frac{1-i}{2}} & \text{if } \frac{1}{2} \leq i \leq 1 \end{cases} \quad (13)$$

It is possible for open loop noise thresholding to retain the mean only because the noise is effectively being added to the gray-tone. In a similar way, mean preservation is possible if a carrier is thresholded.

In the closed loop approach the threshold $T(m, n)$ is updated based on the error $e(m, n)$. Mean is preserved if the following threshold update function is chosen:

$$T(m, n) = T_o(I(m, n)) - e(m, n). \quad (14)$$

This in fact is again a corrective measure by feeding back (diffusing) the error. We shall show that for the uniform pdf case, this system has the desirable property of mean preservation. Let the constant gray-tone be i . Then

$$T(m, n) = 1 - i - e(m, n) = 1 - I(m, n) = T_o(I(m, n)) \quad (15)$$

where $I(m, n) = i + e(m, n)$. Let the expected value of $I(m, n)$ be μ . Then it follows that the mean of the halftone would be μ since such a threshold is being used. It also follows that the expected value of $e(m, n)$ is $\mu - i$. Since $e(m, n) = i - H(m, n)$,

$$E\{e(m, n)\} = \mu - i = i - E\{H(m, n)\} = i - \mu = 0 \quad (16)$$

and the mean is preserved.

To conclude, we note that the feedback path is essential for acceptable halftoning if the gray-tone is thresholded in the proposed tracking system. For noise thresholding, both open loop and closed loop approaches are possible in principle. The open loop approach, however,

fails to produce reliable halftone for some given image because there is too much high frequency noise. The closed loop approach is expected to produce much better halftones due to the presence of feedback.

4.2. Error diffusion revisited

The invention of the error diffusion method is regarded as a breakthrough in digital halftoning. References^{4,5,8,16} include analysis of the error diffusion process and a discussion of some of its advantages. For a halftone of a specified dot density, the error diffusion method produces better quality than methods invented before. It can be shown that error diffusion preserves the mean⁸. However, the error diffusion method has two drawbacks. First, it suffers from *avalanche* or *worm* artifacts in uniform regions. Second, edge-blurring results from the error diffusion. Investigators have looked into modifications to overcome these defects.

Since $\hat{I}(m, n) = I(m, n) + e_2(m, n)$ is quantized in error diffusion, it is equivalent to quantizing $I(m, n)$ with a threshold of $T(m, n) = -e_2(m, n)$ in the proposed tracking approach (refer to figure 2). In z -domain, $e_2(m, n)$ is given by

$$\mathbf{e}_2(z_1, z_2) = G(z_1, z_2) \mathbf{e}_1(z_1, z_2) \quad (17)$$

where bold letters denote the z -transformed signals, and $G(z_1, z_2)$ is the filter of figure 1. Since the error is

$$\begin{aligned} \mathbf{e}_1(z_1, z_2) &= \hat{\mathbf{I}}(z_1, z_2) - \mathbf{H}(z_1, z_2) \\ &= \mathbf{I}(z_1, z_2) - \mathbf{H}(z_1, z_2) + \mathbf{e}_2(z_1, z_2) \\ &= \mathbf{I}(z_1, z_2) - \mathbf{H}(z_1, z_2) + G(z_1, z_2)\mathbf{e}_1(z_1, z_2) \end{aligned} \quad (18)$$

we can express it as follows:

$$\mathbf{e}_1(z_1, z_2) = \frac{\mathbf{I}(z_1, z_2) - \mathbf{H}(z_1, z_2)}{1 - G(z_1, z_2)} \quad (19)$$

Consequently, the equivalent threshold in the conventional tracking approach is given by

$$\mathbf{T}(z_1, z_2) = -\frac{G(z_1, z_2)}{1 - G(z_1, z_2)} \mathbf{I}(z_1, z_2) + \frac{G(z_1, z_2)}{1 - G(z_1, z_2)} \mathbf{H}(z_1, z_2). \quad (20)$$

Thus, the error diffusion algorithm is shown to be a special case with

$$\begin{aligned} f_I(m, n) &= \frac{G}{1 - G} * I(m, n), \\ f_H(m, n) &= \frac{G}{1 - G} * H(m, n), \\ T(m, n) &= -e(m, n). \end{aligned} \quad (21)$$

The above interpretation of the error diffusion algorithm using the proposed framework allows us to better understand the algorithm. The error diffusion algorithm uses an FIR filter $G(z_1, z_2)$, sum of whose impulse-response coefficients is 1 (mean preservation). Therefore, $G/(1 - G)$ becomes an IIR filter with poles on the unit bicircle. Indeed, we attribute the edge-blurring ob-

served with error diffusion to the above fact, namely, the use of unstable IIR filters. The proposed tracking framework readily provides a remedy, too, which is discussed in the next section.

4.3. Other algorithms as special cases

Almost all digital halftoning algorithms that we considered can be accommodated in the proposed tracking based halftoning framework. We illustrate below a few examples, and provide explanations why these algorithms perform the way they are.

Fixed thresholding of section 2.1 is easily achieved by setting $T(m, n)$ to a constant and thresholding the gray-tone. Since no updating takes place, no tracking is achieved. Consequently the resulting halftones are unacceptable.

Electronic screening of section 2.1 is achieved by setting $f_f(m, n) = I(m, n)$, $f_{ff}(m, n) = 0$, $T(m, n) = e(m, n)$, and thresholding the screen function. This being an open loop tracking, the scheme produces good gray-scale rendition (mean preservation) but poor detail rendition. Ordered dithering of section 2.1 is also achieved similarly by setting $f_f(m, n) = I(m, n)$, $f_{ff}(m, n) = 0$, $T(m, n) = e(m, n)$, and thresholding a pseudo-random noise. Similar performance is observed here.

Adaptive thresholding schemes of section 2.2 are special cases of the tracking framework, too. In constrained average thresholding the feedforward filter computes the local average such that $f_f(m, n) = I_{av}(m, n)$, $f_{ff}(m, n) = 0$, $T(m, n) = k_1 e(m, n) + k_2$, and the gray-tone is thresholded. Absence of the feedback path when the gray-tone is thresholded, results in failure to preserve the mean. However, since the threshold is updated based on the local average, details are preserved.

Error diffusion has already been shown as a special case in the earlier section. Most variations of the error diffusion scheme can also be accommodated, hence interpreted in a unified manner, in the proposed framework. For example, the edge enhanced error diffusion⁹ may be obtained with $f_f(m, n) = (G/(1-G)) * I(m, n) + (\kappa - 1)I(m, n)$, $f_{ff}(m, n) = (G/(1-G)) * H(m, n)$, $T(m, n) = -e(m, n)$ and thresholding the gray-tone. For $\kappa > 1$, the magnitude response of the feedforward filter is equally raised by an amount $(\kappa - 1)$. Since the conventional error diffusion filter is low-pass, such a κ allows high-frequency components to pass through the feedforward filter. Consequently the edges, being high frequency in nature, are enhanced.

Some iterative halftoning algorithms of section 2.3, described in^{7,10,11,12} can be captured in the tracking framework if we extend it to include iterations. This would mean the signals obtained in one iteration, such as $H(m, n)$ and $e(m, n)$, are used in the next iteration.

5. Results and conclusion

Our studies have shown that by varying the feedforward and feedback filters and the threshold update algorithm, it is possible to generate halftones that provide better tone quality while retaining edges better than the error diffusion and other approaches. It is also observed that the closed loop noise thresholding algorithm produces halftones which bear greater fidelity to the gray-tone than the error diffusion and other conventional approaches. We have included two



FIG 3a. Original Lenna.



FIG 3b. Error Diffusion.

such examples here. Since the error diffusion algorithm provides a useful benchmark, the results are compared with the error diffusion halftones.

Figure 3 shows the popular face of Lenna. Figure 3a is the gray-tone (actually, it is a 300 dpi halftone that retains practically all features of the gray-tone since it has 16 dots/pixel, 16 times the dot density of other halftones). Figures 3b to 3d are 1 dot/pixel halftones printed at 75 dpi. Figure 3b is the halftone generated by the error diffusion scheme. Figure 3c shows a halftone generated by the proposed tracking approach. The gray-tone is thresholded, no feedforward filter is used, a lowpass filter is used in the feedback path, and a nonlinear threshold update algorithm is chosen. The details are



FIG. 3c. Conventional Tracking.



FIG. 3d. Noise Thresholding

$$\begin{aligned}
 f_i(m, n) &= I(m, n) \\
 f_{ii}(m, n) &= G * H(m, n) \\
 T(m, n) &= -\text{sign}[e(m, n)] \cdot \alpha \cdot |e(m, n)|^\beta
 \end{aligned} \tag{22}$$

where G is the popular causal FIR 3×5 error diffusion filter. Its impulse response coefficients are

$$\begin{bmatrix}
 0 & 0 & 0 & 0.15 & 0.1 \\
 0.06 & 0.1 & 0.15 & 0.1 & 0.06 \\
 0.03 & 0.06 & 0.1 & 0.06 & 0.03
 \end{bmatrix} \tag{23}$$

where 0 denotes the (0, 0)th coefficient. α and β are two parameters controlling the amount of nonlinearity. Figure 3d shows another halftone using the proposed algorithm. A 1-D random noise is thresholded, a high-pass shaping filter is used, both the feedforward and the feedback filters are identical lowpass filters, and a linear threshold update rule is chosen. The details are shaping filter = $[-0.0156 \ 0.0938 \ -0.2344 \ 0.3125 \ -0.2344 \ 0.0938 \ -0.0156]$

$$\begin{aligned}
 f_i(m, n) &= G * I(m, n) \\
 f_{ii}(m, n) &= G * H(m, n) \\
 T(m, n) &= T_n(I(m, n)) - e(m, n)
 \end{aligned} \tag{24}$$

where G is the same filter as before.

Comparing figure 3c with figure 3b, it is obvious that the halftone produced by the tracking method has greater contrast. Related to the above is the better reproduction of tone variation as can be observed in sweeping across the hat. In doing so, it can also be observed that the ridges and edges of the gray-tone are better preserved by the tracking method. Also note that an isolated bright spot behind the hat (on the left, where feathers are beginning; it actually is the result of light reflecting from a standing straw) that is present in the original image, appears in the tracking halftone but not in the error diffusion halftone. We may conclude that using G (tracking) or a stable filter, instead of $G/(1 - G)$ (error diffusion) which is unstable, reduces edge-blurring associated with the error diffusion algorithm; as an added incentive the worms are less in figure 3c. Comparing figure 3d with figure 3b, one finds that the tracking method, although noise-based, produces less noisy halftone than the error diffusion method, while providing a good reproduction of features. Between the two tracking halftones, figure 3c is marginally better than figure 3d.

Figure 4 shows M31, the Andromeda galaxy. Figure 4a is the gray-tone (300 dpi halftone at 9 dots/pixel). Figures 4b to 4d are 1 dot/pixel halftones printed at 100 dpi. Figure 4b is again the error diffusion halftone. Figure 4c is the halftone generated by the tracking method similar to the one used for figure 3c except a linear threshold update rule is used ($\beta = 1$). Similarly, figure 4d uses a noise thresholding method identical to the one used for figure 3d.

Comparing figure 4c with figure 4b, the halftone produced by the tracking method has a better detail rendition since more stars are visible in figure 4c than in figure 4b. Comparing figure 4d with figure 4b, similar observations to figure 3 may again be made. Further, if we were to



FIG. 4a. Original Andromeda Galaxy




FIG. 4b. Error Diffusion.

look at the central white region in the figures, the tracking halftone provides more faithful rendering in terms of gray level as well as the region's boundary.

We have shown that the tracking based framework not only unifies existing halftoning strategies but also provides insight into the role of each constituent. Further, this framework is capable of generating new and efficient halftoning algorithms, such as the noise thresholding scheme. We have also shown that using this framework, better halftones are possible to achieve. The tracking halftones have higher contrast, sharper edges, greater fidelity to the gray-tone, and are less noisy when compared to error diffusion halftones. The tracking approaches are simple




FIG. 4c. Conventional Tracking.



FIG. 4d. Noise Thresholding.

and easy to implement in real time. This framework also points to the potential existence of newer and better halftoning algorithms using various nonlinear choices of the filters and the threshold update rule.

References

- 1 STOFFEL, J. C. AND MORELAND, J. F., A survey of electronic techniques for pictorial image representation, *IEEE Trans. Comm.*, vol. 29, **12**, pp. 1989–1925, 1981.
- 2 WOLF, E., *Progress in Optics*, vol. 33, Elsevier, 1994.
- 3 MITSU, T. AND PARKER, K. J., Digital halftoning technique using a blue-noise mask, *J. Opt. Soc. Am. A*, vol. 9, **11**, pp. 1920–1929, 1992.
- 4 WEISSBACH, S. AND WYROWSKI, F., Error diffusion procedure theory and applications in optical signal processing, *Applied Optics*, vol. 31, **14**, pp. 2518–2534, 1992.
- 5 BARNARD, E., Optimal error diffusion for computer-generated holograms, *J. Opt. Soc. Am. A*, vol. 5, no. 11, pp. 1803–1817, 1988.
- 6 WONG, P. W., Error diffusion with dynamically adjusted kernel, *Proc. ICASSP*, vol. 5, pp. 113–116, 1994.
- 7 KUMAR, A., *Some improvements in error diffusion halftoning of images*, M.E. Thesis, Dept. of ECE, Indian Inst of Science, 1996.
- 8 SULLIVAN, J., MILLER, R. AND PIOS, G., Image halftoning using a visual model in error diffusion, *J. Opt. Soc. Am. A*, vol. 10, **8**, pp. 1714–1724, August 1993.
- 9 ESCHBACH, R. AND KNOX, K. T., Error-diffusion algorithm with edge enhancement, *J. Opt. Soc. Am. A*, vol. 8, **12**, pp. 1844–1850, 1991.
- 10 PELI, E., Multiresolution, error-convergence halftone algorithm, *J. Opt. Soc. Am. A*, vol. 8, **4**, pp. 625–636, 1991.
- 11 KOLLIAS, S. AND ANASTASSIOU, D., A progressive scheme for digital image halftoning, coding of halftones and reconstruction, *IEEE Trans. Sig. Proc.*, vol. 40, 1992.
- 12 KOLLIAS, S. AND ANASTASSIOU, D., A unified neural network approach to digital image halftoning, *IEEE Trans. Sig. Proc.*, vol. 39, **4**, pp. 980–984, 1991.
- 13 SCHEERMESSE, T., BROJA, M. AND BRYNGDAHL, O., Adaptation of spectral constraints to electronically halftoned pictures, *J. Opt. Soc. Am. A*, vol. 10, **3**, pp. 412–417, 1993.
- 14 MRUSEK, R., BROJA, M. AND BRYNGDAHL, O., Halftoning by carrier and spectrum control, *Opt. Commun.*, vol. 75, pp. 375–380, 1990.
- 15 MAKUR, A. AND RAGHUVEER, M. R., Digital image halftoning by noise thresholding, *Proc. ICASSP*, pp. 2535–2538, 1995.
- 16 ULICHNEY, R. A., Dithering with blue noise, *Proc. IEEE*, vol. 76, pp. 56–79, 1988.