Adaptive Threshold Modulation for Error Diffusion Halftoning

Niranjan Damera-Venkata, *Member, IEEE,* and Brian L. Evans, *Senior Member, IEEE*

***Abstract*—Grayscale digital image halftoning quantizes each pixel to one bit. In error diffusion halftoning, the quantization error at each pixel is filtered and fed back to the input in order to diffuse the quantization error among the neighboring grayscale pixels. Error diffusion introduces nonlinear distortion (directional artifacts), linear distortion (sharpening), and additive noise. Threshold modulation, which alters the quantizer input, has been previously used to reduce either directional artifacts or linear distortion. This paper presents an adaptive threshold modulation framework to improve halftone quality by optimizing error diffusion parameters in the least squares sense. The framework models the quantizer implicitly, so a wide variety of quantizers may be used. Based on the framework, we derive adaptive algorithms to optimize 1) edge enhancement halftoning and 2) green noise halftoning. In edge enhancement halftoning, we minimize linear distortion by controlling the sharpening control parameter. We may also break up directional artifacts by replacing the thresholding quantizer with a deterministic bit flipping (DBF) quantizer. For green noise halftoning, we optimize the hysteresis coefficients.**

***Index Terms*—Adaptive quantization, halftoning, limit cycles, raster image processing.**

I. Introduction

DIGITAL image halftoning quantizes a grayscale image to one bit per pixel for display and printing on binary devices. In halftoning by error diffusion [1], the quantization error is linearly filtered and fed back to the input in order to diffuse the quantization error among neighboring grayscale pixels, as shown in Fig. 1(a). Traditionally, the error filter has a finite impulse response (FIR) and the quantizer is a thresholding device with a fixed threshold at mid-gray. Error diffusion degrades the original image by nonlinear distortion (limit cycles), linear distortion (sharpening), and additive noise. The additive noise is shaped to be either highpass (i.e., blue noise) or bandpass (i.e., green noise). For conventional error diffusion, as shown in Fig. 1(a), the shape of the additive noise is highpass. The frequency distortion depends primarily on the error filter being used. Limit cycles appear as directional artifacts [2], and are common in sigma–delta modulation methods such as error diffusion [3].

Manuscript received April 2, 1999; revised August 11, 2000. This work was supported by a U.S. National Science Foundation CAREER Award under Grant MIP-9702707. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Jan P. Allebach.

N. Damera-Venkata was with The University of Texas, Austin, TX 78712. He is now with the Halftoning and Image Processing Group, Hewlett-Packard Laboratories, Palo Alto, CA 94304 USA.

B. L. Evans is with the Embedded Signal Processing Laboratory, The University of Texas, Austin, TX 94304 USA (e-mail: bevans@ece.utexas.edu).

Publisher Item Identifier S 1057-7149(01)00102-6.

Threshold modulation alters the quantizer input, e.g., to reduce directional artifacts or linear distortion. Fig. 1(b) and (c) shows two examples of threshold modulation—edge enhancement error diffusion and green noise halftoning. In green noise halftoning [4], [5], a filtered version of the output is added to the input of the quantizer. This approach clusters halftone dots so that the halftone would be more robust to ink spread and dot gain when printed. Edge enhancement error diffusion is explained next.

The Floyd–Steinberg error filter, which is shown in Fig. 2(a), is a simple nonseparable filter with four dyadic coefficients. Floyd–Steinberg halftones exhibit modest sharpening with respect to the original grayscale image. The longer Jarvis [6], [7] and Stucki [8] error filters, which are shown in Fig. 2(c) and (d), exhibit significant sharpening of the original image.

In edge enhancement error diffusion, Eschbach and Knox [9] modify conventional error diffusion to adjust halftone sharpness, as shown in Fig. 1(b). Their threshold modulation method scales the image by a constant L and adds the result to the quantizer input. As L increases, the sharpness of the resulting halftone increases. In a global sense, one value of L exists that minimizes sharpening, assuming that the image is wide sense stationary and the input and output of the quantizer are jointly wide sense stationary [10]. Smaller values of L would cause blurring, and larger values would cause sharpening, with respect to the original grayscale image. Hence, L can be set to reduce linear distortion.

Kite *et al.* [10], [11] develop a formula for the globally optimal value of L that causes the signal components to be rendered in the halftone without sharpening when using a thresholding quantizer. The quantizer is modeled as a linear gain plus uncorrelated noise. If the gain value is chosen to be the linear minimum mean square error (LMMSE) estimator of the quantizer output [10], [12], then the error is guaranteed to be uncorrelated with the quantizer input. Since the model linearizes the quantizer, halftoning may be analyzed using linear system theory. The linear gain value affects signal shaping in error diffusion and the additive uncorrelated noise affects the noise shaping [10], [12]. The linear gain value does not significantly affect the noise-shaping behavior of error diffusion [3], [10]. This approach assumes that the input and output of the quantizer are jointly wide sense stationary stochastic processes. Since we must approximate statistical averages with sample averages, computing the LMMSE estimator is computationally intensive.

In the most general case, the optimal value of L for sharpness compensation depends on

1) error filter coefficients $\mathcal{H}$;

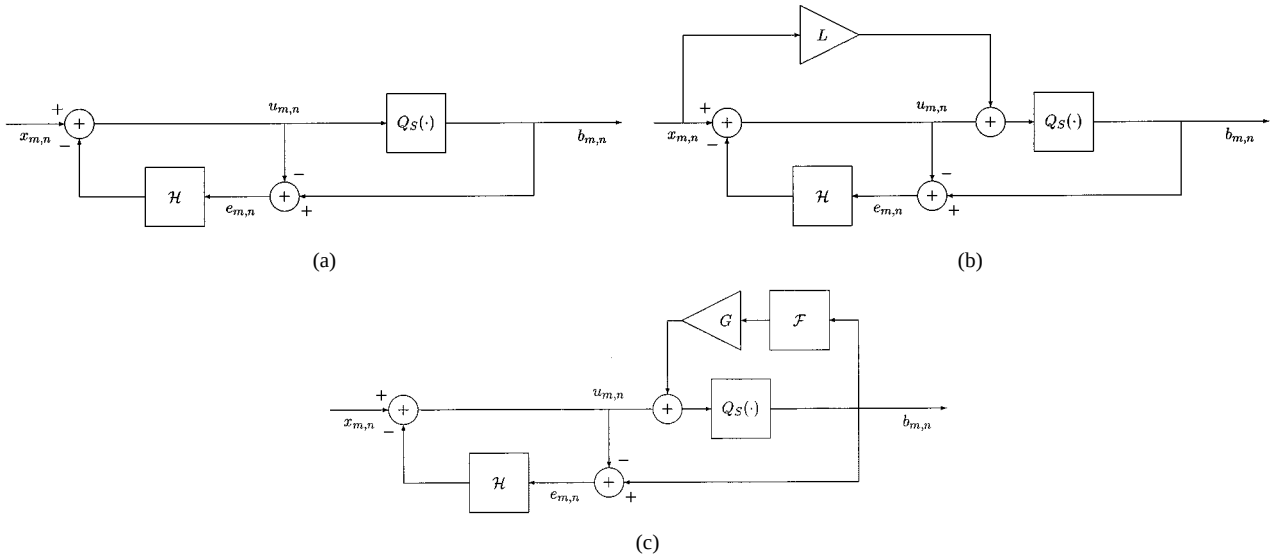

Fig. 1. Block diagrams for various forms of error diffusion. (a) Error diffusion halftoning with a standard thresholding quantizer. (b) Edge enhancement error diffusion halftoning with a standard thresholding quantizer and scalar gain L. (c) Green noise error diffusion halftoning with standard thresholding quantizer, hysteresis 2–D FIR filter $\mathcal{F}$, and scalar gain G.

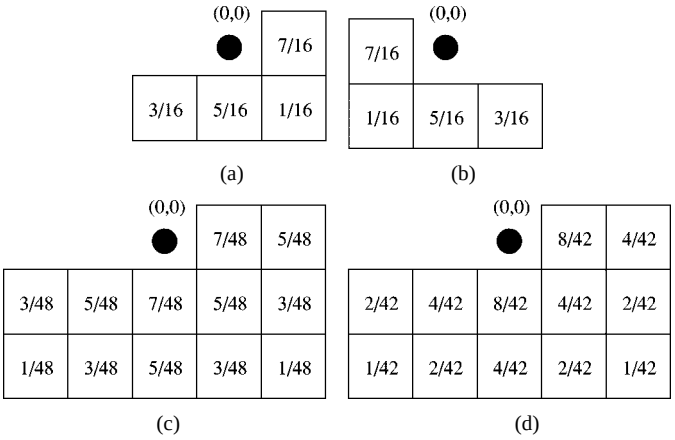

Fig. 2. Common error filters for error diffusion. (a) Floyd–Steinberg (raster), (b) Floyd–Steinberg (serpentine), (c) Jarvis (raster), and (d) Stucki (raster).

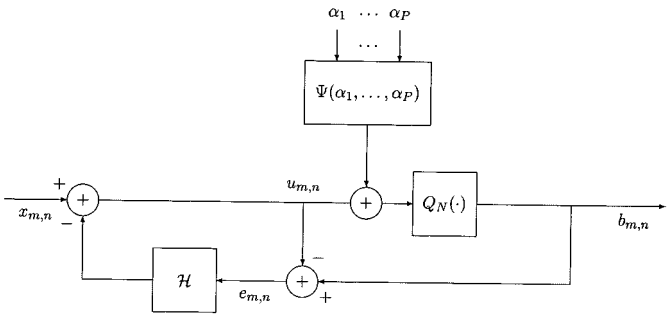

Fig. 3. Generalized error diffusion with threshold modulation.

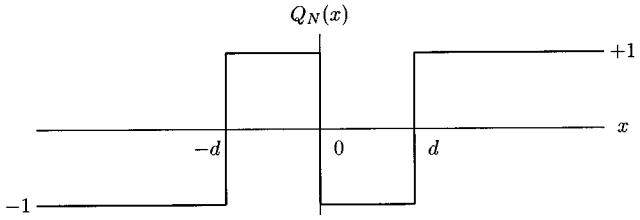

Fig. 4. Deterministic bit flipping quantizer.

Fig. 5. Original grayscale *Lena* image.

2) quantizer function $Q(\cdot)$;
3) input grayscale image.

This suggests a low complexity spatially adaptive algorithm for estimating the optimal value for L to give a what-you-see-is-what-you-get (WYSIWYG) halftone. We define a WYSIWYG halftone as a halftone that preserves the average sharpness of the original grayscale image. This paper develops a framework for spatially adaptive algorithms using adaptive threshold modulation. We show that the adaptive algorithm developed in this paper converges *in the mean* to the optimal value of L if the input and output of the quantizer are jointly wide sense stationary (WSS), and the input image is WSS. In a nonstationary environment, the algorithm tracks local variations in the input image.

Three different approaches modify Floyd–Steinberg error diffusion to reduce directional artifacts [6]–[8]. The first approach

uses longer error filters such as Jarvis [6], [7] and Stucki [8] filters. The second approach uses nonraster scans [13], [14]. A serpentine scan, which scans odd rows from left to right and even rows from right to left, reduces horizontal artifacts. However, it does not reduce artifacts from other directions and may add artifacts not seen in a raster scan [15]. The third approach uses threshold modulation. One type of threshold modulation adds dither (low-amplitude noise) to the quantizer input [15], [16] to break up limit cycles. Adding dither [17], however, adds noise in the halftone. It also increases computational complexity because the pseudo-random numbers either have to be generated on-line or stored in a long array. In a second type of threshold modulation, Wong [18] designs an adaptive algorithm to minimize the quantization error using a weighted mean-squared error (MSE). These halftones have better visual quality than those generated by adding dither.

In this paper, we present a framework for the on-line least squares optimization of error diffusion parameters to improve halftone quality. Because the framework uses an *implicit* LMMSE estimator for the quantizer function, a wide variety of quantizers can be used. Based on the framework, we derive adaptive algorithms to optimize two threshold modulation methods for error diffusion: 1) edge enhancement halftoning and 2) green noise halftoning. For green noise halftoning, the algorithm optimizes the hysteresis coefficients for optimal distribution of dots of a specified size. In this case, the algorithm is shown to converge when the input and output of the quantizer are jointly WSS. For edge enhancement halftoning, the algorithm adapts L to minimize linear frequency distortion (sharpening) to obtain WYSIWYG halftones. We demonstrate the ability of the framework to handle different quantizers by using a thresholding quantizer and a deterministic bit flipping quantizer. The deterministic bit flipping quantizer, which is used in one-dimensional sigma–delta modulators [19], breaks up limit cycles. Using a deterministic bit flipping quantizer with adaptive sharpness control, we simultaneously break up direction artifacts and minimize frequency distortion.

Section II analyzes error diffusion halftoning. It shows that when the least mean squares (LMS) algorithm [20] is used to adapt the error filter to minimize a local MSE criterion, it does not optimize the threshold modulation parameters. Section III derives a general framework for optimizing threshold modulation parameters to minimize a local MSE criterion. The derivation shows that the LMMSE estimator for the quantizer is implicit, which enables a wide variety of quantizers to be used. Section IV optimizes parameters in edge enhancement error diffusion [9] for generating WYSIWYG halftones for the following cases:

1) standard quantizer function and fixed error filter;
2) standard quantizer function and adaptive error filter;
3) nonstandard quantizer function and fixed error filter.

The first case optimizes the sharpness control parameter in modified error diffusion [9] to compensate for linear frequency distortion. This section also presents a low-complexity WYSIWYG halftoning algorithm using a deterministic bit flipping quantizer to break up limit cycles. Section V shows that green noise halftoning is a special case of adaptive threshold modulation. We apply adaptive threshold modulation to optimize the hysteresis filter coefficients in green noise halftoning. Section VI concludes the paper. In the Appendix, we prove that in the case of edge enhancement halftoning (a.k.a., modified error diffusion halftoning) and green noise digital halftoning, our algorithms converge in the mean to the optimal solution, under suitable statistical assumptions about the input and output of the quantizer and the input process. Throughout the paper, we use the 512×512 grayscale *Lena* or *peppers* images to illustrate our algorithms. However, we validated the algorithms on ten test images obtained from the USC image database.

II. Error Diffusion

This section analyzes two extreme examples of error diffusion systems. Section II-A describes conventional error diffusion which uses a fixed error filter and a thresholding quantizer. Error diffusion degrades the original image by nonlinear distortion (directional artifacts), linear distortion (sharpening), and additive noise. Section II-B describes error diffusion using threshold modulation and an adaptive error filter. Section II-B also derives an LMS algorithm to adapt the error filter coefficients in an attempt to minimize the squared error between the input and the output. We show that this LMS approach does not consider the threshold modulation parameters in the update equations.

A. Error Diffusion with a Fixed Error Filter

Fig. 1(a) shows conventional error diffusion. We use $x_{m,n}$ to denote the graylevel of the input image at pixel (m, n), where $x_{m,n} \in [-1, 1]$. We use $b_{m,n}$ to represent the output halftone pixel, where $b_{m,n} \in \{-1, 1\}$, $u_{m,n}$ to denote the input to the quantizer, and $e_{m,n}$ to denote the quantization error. Here, 1 is interpreted as the absence of a printer dot and -1 is interpreted as the presence of a printer dot. $Q_S(\cdot)$ denotes the standard quantizer function given by

$$Q_S(x) = \begin{cases} +1, & x \geq 0 \\ -1, & x < 0. \end{cases} \tag{1}$$

The quantization error at location (m, n) is given by

$$e_{m,n} = b_{m,n} - u_{m,n}. \tag{2}$$

The linear map $\mathcal{H}$, a.k.a. the error filter, filters the previous quantization errors $e_{m,n} \in [-1, 1]$

$$\mathcal{H}e_{m,n} = \sum_{(k,l) \in S} h_{k,l} e_{m-k,n-l} \tag{3}$$

where $\mathcal{H}e_{m,n}$ is fed back to the input, and the set S defines the extent of the error filter coefficient mask. The mask is causal with respect to the image scan, and $(0, 0) \notin S$. Typical raster scan masks for the Floyd–Steinberg filter [1] and Jarvis filter [6] are shown in Fig. 2(a) and (c), respectively.

For serpentine scans using Floyd–Steinberg filters, the mask is shown Fig. 2(a) for odd rows and Fig. 2(b) for even rows.

(a) (b)

Fig. 6. (a) Jarvis and (b) Floyd–Steinberg halftones using a raster scan. Fig. 2 gives the error correction coefficients.

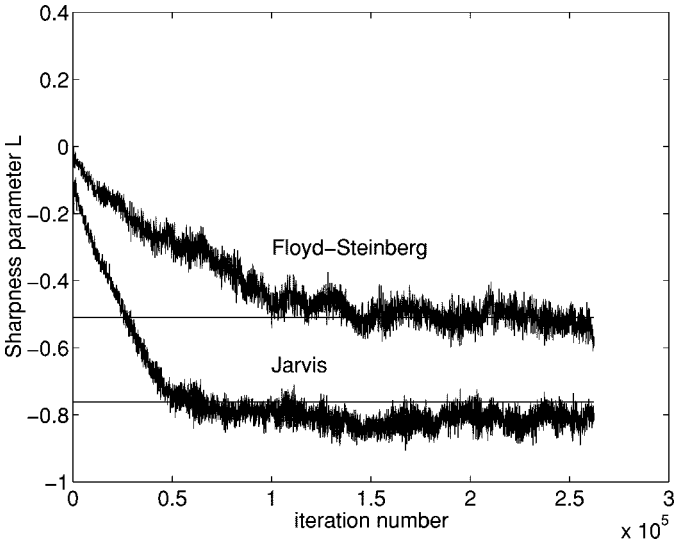

Fig. 7. Adaptation of L for WYSIWYG Jarvis and Floyd–Steinberg halftoning. The horizontal lines indicate the corresponding "optimal" values of L assuming staionary processes.

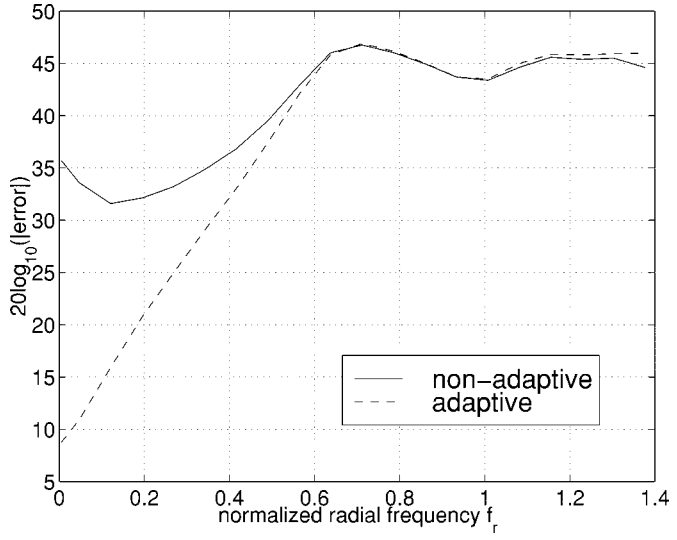

Fig. 8. Radially averaged error spectra on *Lena* image using Jarvis error filter.

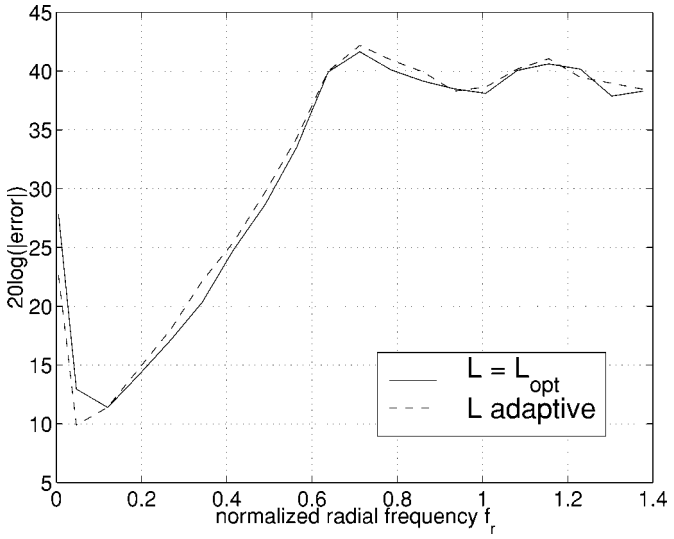

Fig. 9. Radially averaged error spectra using optimal L (assuming stationary processes) and adaptive L on a piecewise constant grayscale ramp. Note the low-frequency improvement of the adaptive method.

To ensure that all of the quantization error is diffused, $\mathcal{H}$ must satisfy the constraint [18]

$$\sum_{(k,l)\in S} h_{k,l} = 1. \tag{4}$$

The quantizer input $u_{m,n}$ and output $b_{m,n}$ are given by

$$u_{m,n} = x_{m,n} - \mathcal{H}e_{m,n} \tag{5}$$

$$b_{m,n} = Q_S(u_{m,n}). \tag{6}$$

B. Error Diffusion with Threshold Modulation and an Adaptive Error Filter

Fig. 3 generalizes error diffusion. $Q_N(\cdot)$ denotes an arbitrary quantizer function, where the subscript N indicates that it may be nonstandard. The linear map $\mathcal{H}^{(m,n)}$ changes at each

(a) (b)

Fig. 10. WYSIWYG error diffusion halftones. (a) Jarvis error filter and (b) Floyd–Steinberg error filter.

(a) (b)

(c)

Fig. 11. Performance of the adaptive algorithm on a mixed document. (a) Original grayscale image, (b) result of Jarvis filter, and (c) result of using adaptive L.

pixel in the image. The function $\Psi(\alpha_1, \cdots, \alpha_P)$ is a differentiable threshold modulating function that modulates the quantizer input. Its parameters $\alpha_1, \cdots, \alpha_P$ control the threshold modulating function

$$u_{m,n} = x_{m,n} - \mathcal{H}^{(m,n)} e_{m,n} \tag{7}$$

$$b_{m,n} = Q_N(u_{m,n} + \Psi(\alpha_1, \cdots, \alpha_P)) \tag{8}$$

$$e_{m,n} = b_{m,n} - u_{m,n}. \tag{9}$$

From (7) and (9), the squared error between the output and input is

$$c_{m,n} = (b_{m,n} - x_{m,n})^2 = \left[e_{m,n} - \mathcal{H}^{(m,n)} e_{m,n}\right]^2. \tag{10}$$

Wong [18] suggests the following approach to minimize the local mean squared error given by (10) by using an adaptive LMS algorithm, in which the weighting is omitted for simplicity

$$\frac{\partial c_{m,n}}{\partial h_{k,l}^{previous}} = -2\left[e_{m,n} - \sum_{(k,l)\in S} h_{k,l}^{previous} e_{m-k,n-l}\right] \times e_{m-k,n-l} \tag{11}$$

$$\hat{h}_{k,l}^{(m,n)} = h_{k,l}^{previous} - \mu \frac{\partial c_{m,n}}{\partial h_{k,l}^{previous}} \tag{12}$$

where μ controls the convergence rate of the algorithm. To satisfy (4)

$$h_{k,l}^{(m,n)} = \hat{h}_{k,l}^{(m,n)} + \gamma \tag{13}$$

where γ is a constant chosen to satisfy

$$\sum_{(k,l)\in S} h_{k,l}^{(m,n)} = 1. \tag{14}$$

Wong's derivation of the update equations does not consider (8). Hence, the parameters of the modulating function $\Psi(\alpha_1, \cdots, \alpha_P)$ are not optimized. The next section demonstrates how the parameters of the modulating function may be modified to minimize an MSE measure.

III. Generalized Adaptive Threshold Modulation

Fig. 3 shows generalized error diffusion. Using (8), the squared error between the output and input is

$$d_{m,n} = (b_{m,n} - x_{m,n})^2 = [Q_N(u_{m,n} + \Psi(\alpha_1, \cdots, \alpha_P)) - x_{m,n}]^2. \tag{15}$$

For any $i \in \{1, 2, \cdots, P\}$

$$\frac{\partial d_{m,n}}{\partial \alpha_i^{previous}} = 2[Q_N(u_{m,n} + \Psi(\alpha_1, \cdots, \alpha_P)) - x_{m,n}] \times \frac{\partial Q_N(u_{m,n} + \Psi(\alpha_1, \cdots, \alpha_P))}{\partial \alpha_i^{previous}} \tag{16}$$

where

$$\frac{\partial Q_N(u_{m,n} + \Psi(\alpha_1, \cdots, \alpha_P))}{\partial \alpha_i^{previous}} = \frac{\partial Q_N(x_{m,n} - \mathcal{H}^{(m,n)} e_{m,n} + \Psi(\alpha_1, \cdots, \alpha_P))}{\partial \alpha_i^{previous}}. \tag{17}$$

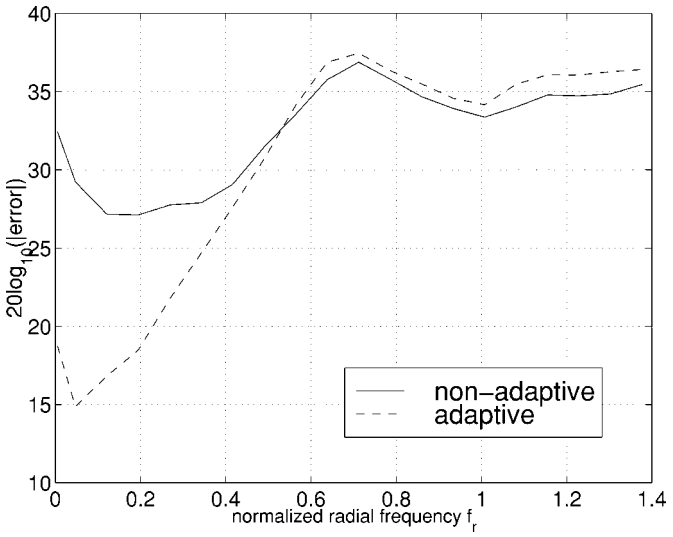

Fig. 12. Radially averaged error spectra on mixed document using Jarvis error filter.

Fig. 13. Adaptation of both edge sharpening parameter L and error filter $\mathcal{H}$.

Fig. 14. Low-complexity WYSIWYG halftoning with reduced artifacts generated by adapting sharpness parameter L and using a DBF quantizer.

For the purpose of computing the derivative, we use a linear minimum mean squares error (LMMSE) estimator for the quantizer.

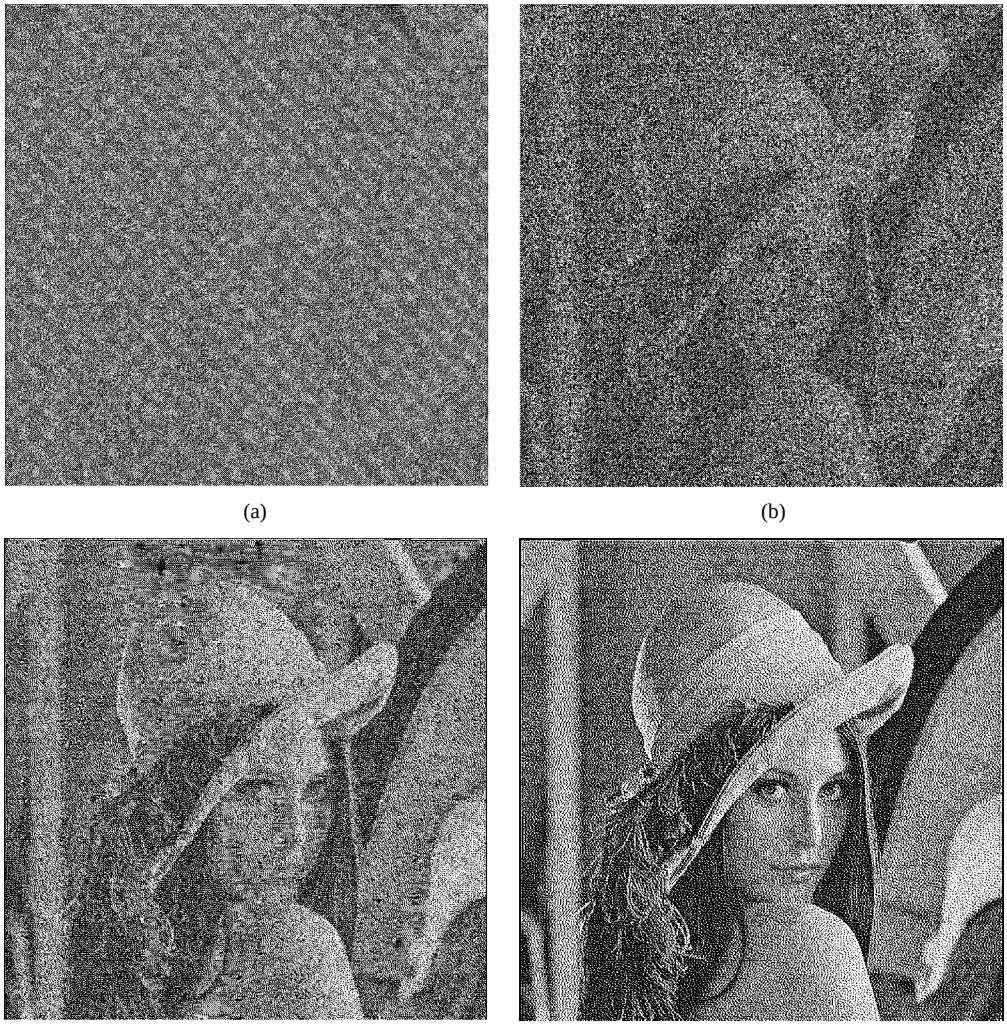

(a) (b) (c) (d)

Fig. 15. Comparison of error images in various halftoning schemes. (a) Error image for adaptive DBF halftoning, (b) error image for nonadaptive DBF halftoning, (c) error image for fixed Floyd–Steinberg halftoning, and (d) error image for fixed Jarvis halftoning.

We model the output of a scalar quantizer $Q(x)$ with input x as [21]

$$\hat{Q}(x) = Ax \tag{18}$$

where

$$A = \frac{\text{Cov}(x,\, Q(x))}{\sigma_x^2}. \tag{19}$$

By substituting (18) and (19) into (17)

$$\frac{\partial(A[x_{m,n} - \mathcal{H}^{(m,n)}e_{m,n} + \Psi(\alpha_1, \cdots, \alpha_P)])}{\partial\alpha_i^{previous}} = A\frac{\partial\Psi(\alpha_1, \cdots, \alpha_P)}{\partial\alpha_i^{previous}}. \tag{20}$$

The parameters of the modulating function are updated as follows:

$$\alpha_i^{new} = \alpha_i^{previous} - \lambda\frac{\partial d_{m,n}}{\partial\alpha_i^{previous}} \tag{21}$$

$$\frac{\partial d_{m,n}}{\partial\alpha_i^{previous}} = 2A[Q_N(u_{m,n} + \Psi(\alpha_1, \cdots, \alpha_P)) - x_{m,n}] \times \frac{\partial\Psi(\alpha_1, \cdots, \alpha_P)}{\partial\alpha_i^{previous}}. \tag{22}$$

The constant $2A$ may be absorbed into the convergence parameter λ. Thus, the update equations do not depend on computing the LMMSE estimator in (18). Next, we use this algorithm to optimize modified error diffusion [9] and green noise [4] halftoning.

IV. Adaptive Threshold Modulation in Modified Error Diffusion

This section develops low-complexity adaptive WYSIWYG halftoning methods. In a WYSIWYG halftoning technique, the sharpness of the halftone and the original grayscale image should be approximately the same. In other words we attempt to minimize linear frequency distortion in the halftone. A user controlled sharpness may be added as a preprocessing step before halftoning. We consider several cases in which the objective is to minimize linear frequency distortion with respect to

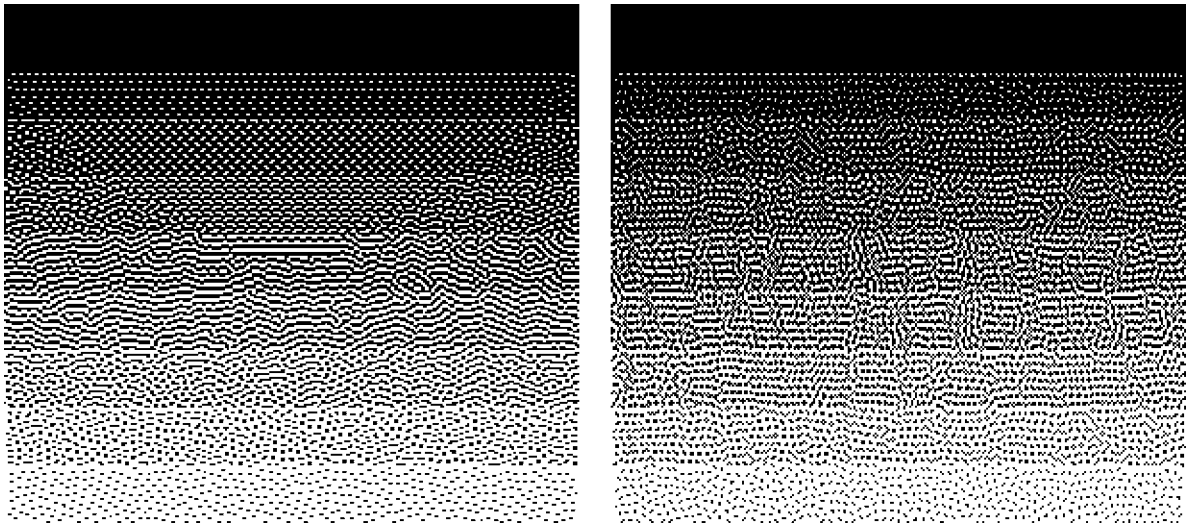

(a) (b)

Fig. 16. Green noise halftones of a grayscale piecewise constant ramp using a Stucki error filter with $G = 0.5$. (a) Floyd–Steinberg hysteresis filter and (b) adaptive hysteresis filter.

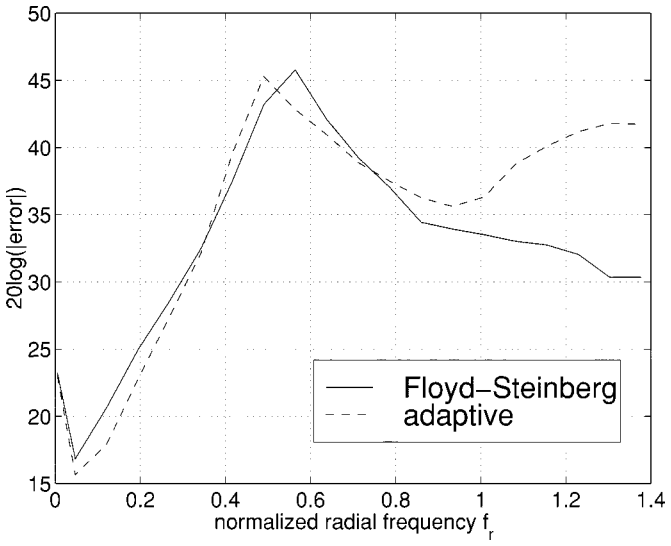

Fig. 17. Radially averaged error-spectra for fixed and adaptive hysteresis green noise halftoning.

the original image. Section IV-A fixes the error filter and uses a thresholding quantizer function $Q_S(\cdot)$. Section IV-B adapts the error filter but still uses the standard quantizer function $Q_S(\cdot)$. Section IV-C fixes the error filter but uses a nonstandard quantizer function $Q_N(\cdot)$. We compare our algorithms with traditional error diffusion schemes based on the correlation of the quantization error image [22] with the original image.

A. Adapting Sharpness for a Fixed Error Filter and a Thresholding Quantizer

Eschbach and Knox [9] show that the sharpness of a halftone may be changed by adding a fraction L of the input image to the quantizer input as in Fig. 1(b). We seek to find the optimal L that will preserve the average sharpness of the grayscale image. That is, we want the signal component of the halftone to be the same as the original image [3], [10]. The threshold modulating function, by inspection of Fig. 1(b), is given by

$$\Psi(L, x_{m,n}) = Lx_{m,n}. \tag{23}$$

By substituting (23) into (15) and differentiating the result

$$\frac{\partial d_{m,n}}{\partial L^{previous}} = [Q_S(x_{m,n} - \mathcal{H}e_{m,n} + L^{previous}x_{m,n}) - x_{m,n}]x_{m,n} \tag{24}$$

$$L^{new} = L^{previous} - \lambda \frac{\partial d_{m,n}}{\partial L^{previous}}. \tag{25}$$

The Appendix shows that in the case of a WSS stochastic process, if the input and the output of the quantizer are jointly WSS processes, and the parameter λ is appropriately chosen, then the algorithm introduced in this section converges in the mean to the globally optimal value of L that minimizes "linear distortion" between the input grayscale image and the output halftone. Here minimizing linear distortion means flattening the signal transfer function of the error diffusion system [3], [10], [12]. Since natural images are nonstationary in general, the algorithm need not converge to a particular value, but rather tracks slowly varying image features.

Fig. 7 shows the values of L over the entire image for Jarvis and Floyd–Steinberg halftones, on the *Lena* image, along with the optimal value of L, found by computing the LMMSE estimator A for the *Lena* image assuming stationary processes—see (39) in the Appendix. Fig. 8 shows a plot of the radially averaged [23] error spectrum obtained by using the fixed Jarvis filter with and without the adaptive algorithm. At the low and mid frequencies where quantization noise is small, the adaptive algorithm, which reduces the linear frequency distortion between the image and the signal component of the halftone, has lower error. At the high frequencies, however, the quantization noise dominates the error spectrum. The adaptive algorithm also introduces more adaptation noise when the parameter L is changed rapidly

(a) (b)

Fig. 18. Result of fixed and adaptive green noise halftoning on *Lena* image ($G = 0.5$). (a) Floyd–Steinberg hysteresis filter and (b) adaptive hysteresis filter.

to high frequency features. This noise is however buried in the shaped quantization noise present at high frequencies. Similar spectra were obtained with the Floyd–Steinberg error filter and on several other test images.

In practice, it is computationally expensive to compute the optimal L using (39) and (19) on an image by image basis, even if we assume stationary processes. The adaptive algorithm presented above overcomes this problem, since the LMMSE estimator is implicitly modeled. Fig. 9 shows the effect of using a fixed L computed using (19) and (39) on a piecewise constant grayscale ramp image (L_{opt} was -0.42 in this case) using the Jarvis error filter, and the effect of using the adaptive algorithm. The adaptive method performs better at the very low frequencies because it tracks changes in the piecewise constant input image.

Fig. 10 shows the results of the adaptive halftoning on the *Lena* image using raster scanning with fixed Jarvis and Floyd–Steinberg error filters, respectively. Fig. 6 shows the results of standard error diffusion using Jarvis and Floyd–Steinberg filters. A visual inspection shows that the WYSIWYG property has been obtained in Fig. 10 (since there is no visible sharpening with respect to the original image), while the halftones in Fig. 6 appear sharper than the original grayscale image shown in Fig. 5. These visual observations are further supported by using a measure of sharpness of the signal component of an error diffused halftone, which we describe in Section IV-E.

Fig. 12 shows the error spectrum for the Jarvis filter, on Fig. 11(a) which is a composite of a natural image with an image containing text. The algorithm is seen to adapt to the slowly varying structure of the natural image, but introduces noise while adapting over the rapidly varying text regions. Fig. 11(b) and (c) shows the resulting halftones. It may be desirable in practice to sharpen/enhance text by pre-sharpening or using a constant value of $L > 0$ over the text regions [9]. This is a limitation of the adaptive algorithm presented in this section.

B. Adapting Sharpness and the Error Filter Coefficients

We adapt both L and the error filter $\mathcal{H}^{(m,n)}$ simultaneously to remove directional artifacts and retain the WYSIWYG property. $\mathcal{H}^{(m,n)}$ is adapted according to (11) with $\mu = 0.05$ as suggested in [18]. Fig. 13 shows the resulting halftone. The initial guess for the error filter was the set of Floyd–Steinberg coefficients, and raster scanning was used. The resulting halftone retains the WYSIWYG property, while the directional artifacts of Floyd–Steinberg error diffusion are also minimized.

C. Adapting Sharpness and a Deterministic Bit Flipping Quantizer

Magrath and Sandler [19] introduce deterministic bit flipping (DBF) quantizers to reduce limit cycles in sigma–delta modulators. DBF is implemented with a modification to the basic quantizer function, as shown in Fig. 4. This quantizer is equivalent to using a standard quantizer, followed by deterministically flipping certain output bits when the quantizer input was in a predetermined range. DBF can be implemented at a much lower complexity than random dithering, and produces higher frequency noise [19]. This is crucial in halftoning applications because we can break up limit cycles by modifying the quantizer function without adding much visible noise. Such a nonstandard quantizer function has not been used previously in halftoning applications.

We may also combine the nonstandard DBF quantizer with the adaptive sharpness control scheme to produce WYSIWYG

halftones with no directional artifacts at a low computational complexity. The DBF quantizer function is

$$Q_N(x) = \begin{cases} 1, & d \le x \\ -1, & \text{otherwise.} \end{cases} \tag{26}$$

The value of d was chosen as the least value that eliminated limit cycles in a piecewise constant grayscale image with ten uniformly spaced graylevels between 0 and 1 (limit cycle behavior for graylevels between -1 and 0 is identical to the above case, with -1 replaced by 1). Based on this test, 0.2 was chosen as a suitable value of d for the DBF quantizer in error diffusion. Fig. 14 shows the output halftone using the DBF quantizer. We use raster scanning and a fixed Floyd–Steinberg error filter. Comparing Fig. 14 with the halftone produced by adapting $\mathcal{H}$ and adapting L in Fig. 13 shows that they are of comparable quality. Section IV-D shows that the complexity of the DBF algorithm is far lower than that of the adaptive error filter approach.

D. Complexity of Adaptive DBF vs. Adaptive Error Filter Techniques

The nonstandard DBF quantizer function $Q_N(\cdot)$ differs from the standard quantizer function $Q_S(\cdot)$ by only one comparison, because

$$Q_N(\theta) = \begin{cases} -Q_S(\theta), & |\theta| \le d \\ Q_S(\theta), & \text{otherwise.} \end{cases} \tag{27}$$

Thus, the DBF quantizer in (27) requires one extra comparison over the standard thresholding quantizer. Updating L according to (24) and (25) requires an extra three additions and three multiplications per pixel because $\mathcal{H}^{(m,n)}e_{m,n}$ is already computed as part of error diffusion. Using (11)–(14), the adaptive error filter method [18] with K filter coefficients requires $(3K+1)$ additions and $(2K+1)$ multiplications per pixel. The complexity of the adaptive DBF method over the adaptive error filter method [18] may be measured using the complexity ratio

$$C = \frac{\text{complexity of adaptive DBF}}{\text{complexity of adaptive error filter}}. \tag{28}$$

If the complexity of an addition is f times the complexity of a multiplication, $f \le 1$, then

$$C = \frac{3}{1 + K\left(\dfrac{3f+2}{f+1}\right)}. \tag{29}$$

When applying the two adaptive methods on Floyd–Steinberg halftoning using a conventional digital signal processor, $K = 1$, $f = 1$ and $C = 3/11$.

E. Cause of Sharpening in Error Diffusion

Knox defined the *error image* in error diffusion to be the matrix of quantization errors scaled and displayed as an image [22]. Kite [11] shows that the correlation of the *error image* with respect to the original is directly related to the frequency distortion produced by error diffusion. Fig. 15(a) shows the error image of the halftone generated by adaptive sharpness using a DBF quantizer, as proposed in this paper. Fig. 15(b) shows the error image of the halftone using a DBF quantizer, without adaptive sharpness control. Some of the correlated image components are visible in the error image. Fig. 15(c) and (d) shows Floyd–Steinberg and Jarvis halftones, respectively. The greater the correlation of the original image with the *error image*, the sharper the halftone. For Fig. 15(a)–(d), the correlations with respect to the original image were 0.0001, 0.14, 0.25, and 0.45, respectively. Also, the DBF quantizer successfully eliminates the directional artifacts of the Floyd–Steinberg halftoning. Thus, both frequency distortion and artifacts are minimized by the WYSIWYG adaptive DBF algorithm.

We validated all of the adaptive algorithms introduced in this section by testing the error images of 10 halftones. All error images had low correlation (<0.006) with respect to the original image. This correlation dropped by two orders of magnitude over the $L = 0$ (no sharpness control) case.

V. Optimal Green-Noise Digital Halftoning

Fig. 1(c) shows the setup for output-dependent feedback proposed by Levien [4]. The effect of adding a filtered version of the output of the quantizer input results in clustering of output pixels. Green noise makes printing devices, such as laser printers, much easier to predict. The benefits of green-noise halftoning are in printing processes with nonideal printing conditions [5]. Lau, Arce and Gallager [5] report that the quantization noise contains *intermediate frequency* between blue noise patterns and ordered dither patterns. They call it "green noise." The hysteresis constant G controls the size of the dot clusters in green noise digital halftones.

We use the theory developed in Section III to adapt the hysteresis filter coefficients $\mathcal{F}^{(m,n)}$. All algorithms involving the method shown in Fig. 1(c) need to use serpentine scanning to avoid strong diagonal artifacts. By analyzing Fig. 1(c), we derive the following equations governing green noise digital halftoning:

$$u_{m,n} = x_{m,n} - \mathcal{H}e_{m,n} \tag{30}$$

$$b_{m,n} = \mathcal{Q}_S\left[u_{m,n} + G\mathcal{F}^{(m,n)}b_{m,n}\right] \tag{31}$$

$$\mathcal{F}^{(m,n)}b_{m,n} = \sum_{(k,l)\in S_F} f_{k,l}\, b_{m-k,n-l}. \tag{32}$$

The quantization error and $u_{m,n}$ are computed as usual using (2) and (5), respectively.

By inspection of Fig. 1(c)

$$\Psi(G, \mathcal{F}^{(m,n)}, b_{m,n}) = G\mathcal{F}b_{m,n}. \tag{33}$$

Therefore, the adaptation equations for the hysteresis filter coefficients become

$$\frac{\partial d_{m,n}}{\partial f_{k,l}^{(previous)}} = Gb_{m-k,n-l} \times \left[Q_S\left(x_{m,n} - \mathcal{H}e_{m,n} + \mathcal{F}^{(previous)}b_{m,n}\right) - x_{m,n}\right]. \tag{34}$$

Note that the two sums in the quantizer function are computed anyway and do not add complexity to the adaptation process.

To form the final updates, we add a constraint to guarantee that the hysteresis filter coefficients are nonnegative and sum to one. This ensures that the hysteresis filter does not change the dot size.The set of filters $\mathcal{C}$ satisfying the constraints forms a closed convex set in $\mathcal{R}^{|S_F|}$ [24] where $|S_F|$ denotes the cardinality of S_F. The adaptation equation in this case is given by

$$\mathbf{f}^{(m,n)} = P_{\mathcal{C}}(\mathbf{f}^{previous} - \mu \nabla d_{m,n}) \tag{35}$$

where $P_{\mathcal{C}}(\cdot)$ is the Hilbert space projection onto the closed convex set $\mathcal{C}$, $\mathbf{f}^{(m,n)}$ denotes the vector of filter coefficients and the argument of $P_{\mathcal{C}}(\cdot)$ denotes the standard update (21) in vector notation.

The Appendix shows that the adaptive algorithm described in this section converges in the mean to the globally optimal solution under suitable statistical assumptions. The Appendix discusses a method to compute the orthogonal projection operator $P_{\mathcal{C}}(\cdot)$.

To illustrate the framework of Section III, we adapt a four-tap hysteresis filter. The error filter coefficients are the Stucki coefficients. The initial guess for the hysteresis filter coefficients are the Floyd–Steinberg coefficients. This corresponds to the 4-tap hysteresis filter, 12-tap error filter green noise scheme used in [5].

Fig. 16(a) shows the halftone obtained on a grayscale ramp by using the Floyd–Steinberg hysteresis coefficients, while Fig. 16(b) shows the halftone obtained by using the adaptive algorithm described in this section. In both cases the dot size was held constant by fixing $G = 0.5$ as suggested in [5]. The adaptive algorithm breaks up the directional artifacts obtained on using the Floyd–Steinberg hysteresis filter [5]. Fig. 17 shows the improvement in the error spectra of the adaptive green noise halftone at the lower frequencies. Fig. 18 shows the results of fixed and adaptive hysteresis error filter on the *Lena* image. The adaptive algorithm breaks up the long "worm" artifacts seen in the smooth regions of the image.

The Appendix shows that the above algorithm converges in the mean if the input and output of the quantizer are jointly wide sense stationary. This assumption need not hold true for natural images. However, the mean value of the iterate is a good value to use if fixed hysteresis coefficients are desired. In our simulations, we use $\mu = 0.005$ as the convergence parameter in our adaptive algorithms.

VI. Conclusion

In this paper, we introduce a general framework for adapting the parameters of a differentiable threshold modulating function to minimize a mean square error measure. Competing techniques that adapt error filter coefficients and inject noise by adding dither do not optimize the threshold modulating parameters and have higher implementation complexity. Based on the framework, we develop a low-complexity algorithm to minimize the linear distortion (sharpening) in error diffused halftones. The savings on a conventional digital signal processor are a factor of 3.67 for Floyd–Steinberg halftoning and a factor of 10.33 for Jarvis halftoning, when compared with adaptive error filter algorithms [18]. By preserving the sharpness of the grayscale image, a separate preprocessing method may be applied for customized image enhancement.

Using the framework, we optimize hysteresis coefficients in green noise halftoning, and the edge sharpening parameter in edge enhancement halftoning. By replacing a thresholding quantizer with a deterministic bit flipping quantizer, we break up directional artifacts caused by limit cycles. The framework may be further improved by incorporating human visual models in the objective function and by using variable step size adaptive algorithms.

Appendix

In the Appendix, we prove that the adaptive algorithms proposed in this paper converge in the mean when the input and output of the quantizer are jointly wide sense stationary, and the quantizer may be modeled using a linear gain model [3], [10]. The linear gain model depends on the accuracy of modeling the quantizer with a scalar gain for the signal component. This model has been validated in [10], [21], and accurately predicts linear effects in halftoning such as linear distortion (sharpening) and noise shaping.

A. Adaptive Modified Error Diffusion

Proposition 1: The optimal value of L, L^{opt} under the assumptions stated above is given by $L^{opt} = ((1-A)/A)$, where $A = (\mathrm{Cov}(u_{m,n}, b_{m,n})/\mathrm{Var}(u_{m,n}))$ is the LMMSE estimator for the quantizer output, a.k.a. the "linear gain" of the quantizer.

Proof: For the optimal result in the global sense we need to take expectations on both sides of (24) and set the result to zero. This means that the optimal solution L_{opt}, satisfies

$$\frac{\partial E[d_{m,n}]}{\partial L^{previous}} = E[[Q_S(x_{m,n} - \mathcal{H}e_{m,n} + L^{opt}x_{m,n}) - x_{m,n}]x_{m,n}] = 0. \tag{36}$$

Using the linear gain approximation for the quantizer function this becomes

$$E[[A(x_{m,n} - \mathcal{H}e_{m,n} + L^{opt}x_{m,n}) - x_{m,n}]x_{m,n}] = 0. \tag{37}$$

Also, since the optimal solution leads to an error image, uncorrelated with the input image [10], [22], we can make the following approximation based on the linear gain model [10], [12]:

$$E[\mathcal{H}e_{m,n}x_{m,n}] \approx 0. \tag{38}$$

By substituting (38) into (36) and solving for L^{opt}

$$L^{opt} = \frac{1-A}{A}. \tag{39}$$

This completes the proof.

In fact, this condition is the same as the *globally* optimal value of L derived by Kite *et al.* using frequency domain methods [10]. What is crucial, however, is that the optimal L, L^{opt} satisfies (36). We will use this fact to establish the following theorem.

Theorem 1: The adaptive algorithm proposed in Section IV-A converges in the mean to the globally optimal solution L^{opt} if $0 < \lambda < 2/(A\,E[x_{m,n}^2])$.

Proof: The update equation for the adaptive parameter L_p in the pth iteration/adaptation is given by (24) and (25), which can be written as

$$L_{p+1} = L_p - \lambda\{A[x_{m,n} - \mathcal{H}e_{m,n} + L_p x_{m,n}] - x_{m,n}\}x_{m,n} \tag{40}$$

where we have made use of the linear gain model, by introducing the parameter A. By rewriting (40)

$$L_{p+1} = L_p(1 - \lambda A x_{m,n}^2) - \lambda A(\mathcal{H}e_{m,n})x_{m,n} + \lambda A x_{m,n}^2 - \lambda x_{m,n}^2. \quad (41)$$

By subtracting L^{opt} from both sides of (41) and taking expectations

$$E[L_{p+1}] - L^{opt} = L_r(1 - \lambda A E[x_{m,n}^2]) - \lambda A E[(\mathcal{H}e_{m,n})x_{m,n}] + \lambda A E[x_{m,n}^2] - \lambda E[x_{m,n}^2] - L^{opt}. \quad (42)$$

Since L^{opt} satisfies (36)

$$-\lambda A E[(\mathcal{H}e_{m,n})x_{m,n}] + \lambda A E[x_{m,n}^2] - \lambda E[x_{m,n}^2] = \lambda A L^{opt} E[x_{m,n}^2]. \quad (43)$$

By substituting (43) into (42)

$$E[L_{p+1}] - L^{opt} = [1 - \lambda\ A\ E[x_{m,n}^2]][L_p - L^{opt}]. \quad (44)$$

Therefore

$$E[L_p] - L^{opt} = [1 - \lambda\ A\ E[x_{m,n}^2]]^p[L_0 - L^{opt}] \quad (45)$$

where L_0 is an arbitrary initial guess. Hence if $|1 - \lambda\ A\ E[x_{m,n}^2]| < 1$, or $0 < \lambda < 2/(A\ E[x_{m,n}^2])$, then the adaptive algorithm converges in the mean to the optimal value of L. This means that $E[L_p] \longrightarrow L^{opt}$ under the assumptions we have made at the beginning of this Appendix. This completes the proof of the theorem.

B. Adaptive Green Noise Error Diffusion

We denote the hysteresis filter at the pth iteration as a vector $\mathbf{f}_p$ and the output image pixels covered by the hysteresis filter mask at location (m, n) by the vector $\mathbf{b}_{m,n}$. Thus, the hysteresis filter output at location (m, n) and iteration p is given by $\mathbf{b}_{m,n}^T\mathbf{f}_p$.

Theorem 2: The adaptive algorithm proposed in Section V converges in the mean to a globally optimal solution $\mathbf{f}^{opt}$ if $0 < \mu < 2/(A\ G^2\ \mathrm{Tr}(\mathbf{K_b}))$ and the output vector $\mathbf{b}_{m,n}$ and the hysteresis filter vector $\mathbf{f}_p$ are statistically independent. $\mathrm{Tr}(\mathbf{K_b})$ refers to the trace of the autocorrelation matrix $\mathbf{K_b}$ which is defined as the sum of its diagonal elements.

This assumption is similar to the independence assumption made for conventional LMS adaptive filters [25], which says that the data (input) and the LMS weight vector are statistically independent.

Proof: The unconstrained optimal solution $\mathbf{f}^o$ must satisfy

$$\left.\frac{\partial E[d_{m,n}]}{\partial \mathbf{f}}\right|_{\mathbf{f}=\mathbf{f}^o} = GE[\mathbf{b}_{m,n}[A(x_{m,n} - \mathbf{h}^T\mathbf{e}_{m,n} + G\mathbf{b}_{m,n}^T\mathbf{f}^o) - x_{m,n}]] = 0. \quad (46)$$

However, we impose the requirement that the elements of the constrained optimal solution $\mathbf{f}^{opt}$ must be positive and sum to one. Since the constraint set $\mathcal{C}$ is a closed convex set in $\mathcal{R}^N$, where N is the dimension of $\mathbf{f}_p$, we can define a projection operator $P_{\mathcal{C}}$ onto $\mathcal{C}$ by

$$(\forall \mathbf{w} \in \mathcal{R}^N)\|P_{\mathcal{C}}(\mathbf{w}) - \mathbf{w}\| = \min_{\mathbf{v}\in\mathcal{C}} \|\mathbf{v} - \mathbf{w}\|. \quad (47)$$

The constrained optimal solution is characterized using the projection operator [26], [27] by

$$\mathbf{f}^{opt} = P_{\mathcal{C}}(\mathbf{f}^{opt} - \mu GE[\mathbf{b}_{m,n}[A(x_{m,n} - \mathbf{h}^T\mathbf{e}_{m,n} + G\mathbf{b}_{m,n}^T\mathbf{f}^{opt}) - x_{m,n}]]) \quad (48)$$

which may be rewritten as

$$\mathbf{f}^{opt} = P_{\mathcal{C}}((\mathbf{I} - \mu AG^2\mathbf{K_b})\mathbf{f}^{opt} + \mu[(A-1)GE[\mathbf{b}_{m,n}x_{m,n}] + AGE[\mathbf{b}_{m,n}\mathbf{h}^T\mathbf{e}_{m,n}]]) \quad (49)$$

where $\mathbf{K_b} = E[\mathbf{b}_{m,n}\mathbf{b}_{m,n}^T]$ and $\mathbf{I}$ is the $N \times N$ identity matrix. It follows from the development in Proposition 1 in [26] as well as [27] and the fact that the unconstrained solution satisfies (46), that the iteration

$$\mathbf{f}_{p+1} = P_{\mathcal{C}}((\mathbf{I} - \mu AG^2\mathbf{K_b})\mathbf{f}_p + \mu[(A-1)GE[\mathbf{b}_{m,n}x_{m,n}] + AGE[\mathbf{b}_{m,n}\mathbf{h}^T\mathbf{e}_{m,n}]]) \quad (50)$$

converges in mean to $\mathbf{f}^{opt}$, if $0 < \mu < 2/(A\ G^2\ \lambda_N)$ where λ_N is the maximum eigenvalue of the matrix $\mathbf{K_b}$. Since $\mathrm{Tr}(\mathbf{K_b}) > \lambda_N$, the theorem follows. The iteration converges to a global optimum because the Hessian of the objective function to be minimized over the convex set $\mathcal{C}$, is given by a positive semi-definite $AG^2\mathbf{K_b}$. $AG^2\mathbf{K_b}$ is positive semi-definite because $\mathbf{K_b}$ is an autocorrelation matrix.

Computation of the Projection Operator $P_{\mathcal{C}}$: To enforce that the iterates do not leave the constraint set $\mathcal{C}$, we introduce the auxiliary variable $\theta_{k,l}$ such that

$$f_{k,l} = \theta_{k,l}^2 \quad (51)$$

and

$$\sum_{(k,l)\in S_F} \theta_{k,l}^2 = 1. \quad (52)$$

In terms of $\theta_{k,l}$, the adaptation equation given by (21) becomes

$$\theta_{k,l}^{new} = \theta_{k,l}^{previous} - \mu\frac{\partial d_{m,n}}{\partial \theta_{k,l}^{previous}} \quad (53)$$

with

$$\frac{\partial d_{m,n}}{\partial \theta_{k,l}^{previous}} = \frac{\partial d_{m,n}}{\partial f_{k,l}^{previous}}\, 2\theta_{k,l}^{previous}. \quad (54)$$

To enforce (52), we normalize the update as follows:

$$\hat{\theta}_{k,l}^{new} = \frac{\theta_{k,l}^{new}}{\|\theta\|}. \quad (55)$$

The projected parameters $f_{k,l}$ satisfying the constraints are given by

$$f_{k,l}^{new} = (\hat{\theta}_{k,l}^{new})^2. \tag{56}$$

We use the operator that maps the iterate into the constraint set $\mathcal{C}$ as an approximation to the true projection operator $P_{\mathcal{C}}$.

REFERENCES

[1] R. Floyd and L. Steinberg, "An adaptive algorithm for spatial grayscale," *Proc. Soc. Image Display*, vol. 17, no. 2, pp. 75–77, 1976.

[2] Z. Fan and R. Eschbach, "Limit cycle behavior of error diffusion," in *Proc. IEEE Conf. Image Processing*, vol. 2, Nov. 1994, pp. 1041–1045.

[3] T. D. Kite, B. L. Evans, A. C. Bovik, and T. L. Sculley, "Digital halftoning as 2-D delta–sigma modulation," in *Proc. IEEE Conf. Image Processing*, vol. 1, Oct. 1997, pp. 799–802.

[4] R. Levien, "Output dependent feedback in error diffusion halftoning," *IS&T Imag. Sci. Technol.*, vol. 1, pp. 115–118, May 1993.

[5] D. L. Lau, G. R. Arce, and N. C. Gallagher, "Green-noise digital halftoning," *Proc. IEEE*, vol. 86, pp. 2424–2442, Dec. 1998.

[6] J. Jarvis, C. Judice, and W. Ninke, "A survey of techniques for the display of continuous tone pictures on bilevel displays," *Comput. Graph. Image Process.*, vol. 5, pp. 13–40, 1976.

[7] J. Jarvis and C. Roberts, "A new technique for displaying continuous tone images on a bilevel display," *IEEE Trans. Commun.*, pp. 891–898, Aug. 1976.

[8] P. Stucki, "MECCA—A multiple-error correcting computation algorithm for bilevel hardcopy reproduction," IBM Res. Lab., Zurich, Switzerland, Res. Rep. RZ1060, 1981.

[9] R. Eschbach and K. Knox, "Error-diffusion algorithm with edge enhancement," *J. Opt. Soc. Amer. A*, vol. 8, pp. 1844–1850, Dec. 1991.

[10] T. D. Kite, B. L. Evans, and A. C. Bovik, "Modeling and quality assessment of halftoning by error diffusion," *IEEE Trans. Image Processing*, vol. 9, pp. 909–922, May 2000.

[11] T. D. Kite, "Design and quality assessment of forward and inverse error-diffusion halftoning algorithms," Ph.D. dissertation, Dept. Elect. Comput. Eng., Univ. Texas, Austin, Aug. 1998.

[12] N. Damera-Venkata and B. L. Evans, "Design and analysis of vector color error diffusion halftoning systems," *IEEE Trans. Image Processing*, submitted for publication.

[13] I. Witten and R. Neal, "Using Peano curves for bilevel display of continuous-tone images," *IEEE Comput. Graph. Applicat.*, pp. 47–51, May 1982.

[14] R. Ulichney, "Dithering with blue noise," *Proc. IEEE*, vol. 76, pp. 56–79, Jan. 1988.

[15] K. Knox and R. Eschbach, "Threshold modulation in error diffusion," *J. Electron. Imag.*, vol. 2, pp. 185–192, July 1993.

[16] K. Knox, "Threshold modulation in error diffusion on nonstandard rasters," *Proc. SPIE*, vol. 2179, pp. 159–169, Feb. 1994.

[17] S. Norsworthy, R. Schreier, and G. Temes, Eds., *Delta–Sigma Data Converters*. New York: IEEE Press, 1997.

[18] P. Wong, "Adaptive error diffusion and its application in multiresolution rendering," *IEEE Trans. Image Processing*, vol. 5, pp. 1184–1196, July 1996.

[19] A. J. Magrath and M. B. Sandler, "A sigma–delta modulator topology with high linearity," *Proc. IEEE Int. Symp. Circuits Systems*, pp. 53–56, June 1997.

[20] B. Widrow and S. Stearns, *Adaptive Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1985.

[21] S. Ardalan and J. Paulos, "An analysis of nonlinear behavior in delta–sigma modulators," *IEEE Trans. Circuits Syst.*, vol. CAS-34, pp. 593–603, June 1987.

[22] K. Knox, "Error image in error diffusion," *Proc. SPIE*, vol. 1657, pp. 268–279, Feb. 1992.

[23] R. Ulichney, *Digital Halftoning*. Cambridge, MA: MIT Press, 1987.

[24] H. Stark and Y. Yang, *Vector Space Projections*. New York: Wiley, 1998.

[25] S. Haykin, *Adaptive Filter Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1996.

[26] P. L. Combettes and P. Bondon, "Adaptive linear filtering with convex constraints," *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, vol. 2, pp. 1372–1375, May 1995.

[27] ——, "Constrained adaptive filtering for dependent processes,", Tech. Rep., Elect. Eng. Dept./CUNY and LSS/CNRS, Oct. 1994.

Niranjan Damera-Venkata (M'00) received the B.S.E.E. degree from the University of Madras, Madras, India, in July 1997 and the M.S.E.E. and Ph.D.E.E. degrees from The University of Texas, Austin, in May 1999 and December 2000, respectively.

He is currently a Research Engineer with the Hewlett-Packard Research Laboratories, Palo Alto, CA. His research interests include document image processing, symbolic design and analysis tools, image and video quality assessment, and fast algorithms for image processing.

Dr. Damera-Venkata is a member of Sigma Xi. He won the 1998–1999 Texas Telecommunications Engineering Consortium Graduate Fellowship from The University of Texas.

Brian L. Evans (S'88–M'93–SM'97) received the B.S.E.E.C.S. degree from the Rose-Hulman Institute of Technology, Terre Haute, IN, in May 1987, and the M.S.E.E. and Ph.D. degrees from the Georgia Institute of Technology, Atlanta, in December 1988 and September 1993, respectively.

From 1993 to 1996, he was a Postdoctoral Researcher with the University of California, Berkeley, where he worked on electronic design automation for embedded systems as a Member of the Ptolemy Project. He is the Primary Architect of the Signals and Systems Pack for Mathematica, which has been on the market since October 1995. He is currently an Associate Professor with the Department of Electrical and Computer Engineering, The University of Texas, Austin. He is also the Director of the Embedded Signal Processing Laboratory within the Center for Vision and Image Sciences. His research interests include real-time embedded systems; signal, image and video processing systems; system-level design; symbolic computation; and filter design. He developed and currently teaches multidimensional digital signal processing, embedded software systems, and real-time digital signal processing laboratory.

Dr. Evans is an Associate Editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING, a member of the Design and Implementation of Signal Processing Systems Technical Committee of the IEEE Signal Processing Society, and the recipient of a 1997 National Science Foundation CAREER Award.