Color Error Diffusion Halftoning

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

Abstract

Grayscale halftoning converts a continuous-tone image (e.g. 8 bits per pixel) to a lower resolution (e.g. 1 bit per pixel) for printing or display. Grayscale halftoning by error diffusion uses feedback to shape the quantization noise into high frequencies where the human visual system (HVS) is least sensitive. In color halftoning, the application of grayscale error diffusion methods to the individual colorant planes fails to exploit the HVS response to color noise. Ideally the quantization error must be diffused to frequencies and colors, to which the HVS is least sensitive. Further it is desirable for the color quantization to take place in a perceptual space so that the colorant vector selected as the output color is perceptually closest to the color vector being quantized. This article discusses the design principles of color error diffusion that differentiate it from grayscale error diffusion.

Index terms: color halftoning, color quantization, human visual system models, image display

N. Damera-Venkata is with the Imaging Systems Laboratory, Hewlett-Packard Labs, Palo Alto, CA 94304 USA (e-mail: damera@exch.hpl.hp.com). B. L. Evans and V. Monga are with the Embedded Signal Processing Laboratory, Dept. of Electrical and Computer Engineering, The University of Texas, Austin, TX 78712 USA (e-mail: bevans@ece.utexas.edu, vishal@ece.utexas.edu).

This research was supported by a US National Science Foundation CAREER Award under Grant MIP-9702707.

I. Introduction

Color error diffusion is a high-quality method for color rendering of *continuous-tone* digital color images on devices with limited color palettes such as low-cost displays and printers. For display applications the input colorant space is a triplet of red, green, and blue (RGB) values, and the choice of output levels (i.e., the color palette) is a design parameter. For printing applications the input colorant space is a quadruple of cyan, magenta, yellow, and black (CMYK) values and the output levels are fixed. For example, for a bi-level CMYK printer there are 16 possible output colors. In this article we use an input RGB color space and discuss binary error diffusion. This allows us to concentrate the exposition to the essential properties of color error diffusion system design without having to focus on the issues of color palette design and device dependent nonlinear color transformations. However, most of the results can be easily extended to multilevel images and other color spaces.

The application of grayscale error diffusion methods to the individual colorant planes fails to exploit the human visual system response to color noise. Ideally the quantization error must be diffused to frequencies and colors, to which the human visual system is least sensitive. Further it is desirable for the color quantization to take place in a perceptual space (such as Lab) so that the colorant vector selected as the output color is perceptually the closest to the color vector being quantized. In this article we discuss each of the above two design principles of color error diffusion that differentiate it from grayscale error diffusion. We have implemented these design principles in a freely distributable digital halftoning toolbox in Matlab at

http://www.ece.utexas.edu/~bevans/projects/halftoning/toolbox/index.html

II. Analysis of color error diffusion

Fig. 1 shows the system block diagram for color error diffusion halftoning. In Fig. 1, each signal is vector-valued, e.g. a vector of RGB values. The quantization error $\mathbf{e}[\mathbf{m}]$ is fed back, filtered by the error filter $\tilde{\mathbf{h}}(\mathbf{m})$, and added to the input continuous-tone image $\mathbf{x}[\mathbf{m}]$. The net effect is to diffuse the quantization error frequencies and colors.

The error filter $\tilde{\mathbf{h}}(\mathbf{m})$ has matrix-valued coefficients, and we use a ~ to differentiate a

matrix from a vector. The matrix-valued sequence $\tilde{\mathbf{h}}(\cdot)$ is defined on a causal support set $\mathcal{S}$ with $(0,0) \notin \mathcal{S}$. If the vector color error diffusion halftoning system were applied separably to each color, then the coefficients of $\tilde{\mathbf{h}}(\mathbf{m})$ would be diagonal matrices. We will be handling the more general case that the coefficients of $\tilde{\mathbf{h}}(\mathbf{m})$ are not diagonal matrices to account for diffusing the quantization error to other colors. The error filter $\tilde{\mathbf{h}}(\mathbf{m})$ is 2-D multifilter, and its output is defined by a matrix-vector convolution:

$$\left[\tilde{\mathbf{h}} \star \mathbf{e}\right](\mathbf{m}) = \sum_{\mathbf{k} \in \mathcal{S}} \tilde{\mathbf{h}}(\mathbf{k})\mathbf{e}(\mathbf{m}-\mathbf{k}) \tag{1}$$

In the z-domain, the matrix-vector convolution becomes a linear transformation

$$Z\left[\tilde{\mathbf{h}} \star \mathbf{e}\right](\mathbf{z}) = \tilde{\mathbf{H}}(\mathbf{z})\mathbf{E}(\mathbf{z}) \tag{2}$$

A. Modeling color error diffusion

The sole nonlinearity in the vector color error diffusion system in Fig. 1 is the quantizer. The quantizer of Fig. 1 can be linearized by modeling its effect on the input signal (the original image) and by modeling its injection of quantization noise into the system. The quantizer effect on the input signal is modeled by a constant linear transformation (gain) denoted by a matrix $\tilde{\mathbf{K}}$. This matrix gain is applied only to the signal components of the quantizer input. The injection of quantization noise is modeled as spatially-varying additive noise $\mathbf{n}(\mathbf{m})$. The additive noise is applied only to the noise components of the quantizer input [1]. The noise components of the quantizer input are those that are uncorrelated with the input signal. We refer to this matrix gain plus additive noise model for the quantizer as simply a matrix gain model.

Figs. 2 and 3 show the signal and noise paths, respectively, of the vector color error diffusion system linearized by the matrix gain model for the quantizer. This matrix gain model is a generalization of modeling the quantizer in a sigma-delta modulator [2] and grayscale error diffusion [3], [4]. Correlation among the signal color planes is represented by the off-diagonal terms in the matrix $\tilde{\mathbf{K}}$. Matrix $\tilde{\mathbf{K}}$ is chosen to minimize the error in approximating the quantizer with a linear transformation, in the linear minimum mean squared error sense [1]:

$$\tilde{\mathbf{K}} = \arg\min_{\tilde{\mathbf{A}}} E[\| \mathbf{b}(\mathbf{m}) - \tilde{\mathbf{A}}\ \mathbf{u}(\mathbf{m}) \|^2] \tag{3}$$

The solution to (3) when $\mathbf{b}(\mathbf{m})$ and $\mathbf{u}(\mathbf{m})$ are wide sense stationary processes is [1]

$$\tilde{\mathbf{K}} = \tilde{\mathbf{C}}_{\mathbf{bu}}\tilde{\mathbf{C}}_{\mathbf{uu}}^{-1} \tag{4}$$

where $\tilde{\mathbf{C}}_{\mathbf{bu}}$ and $\tilde{\mathbf{C}}_{\mathbf{uu}}$ are covariance matrices. As a direct consequence of this modeling [5], the noise process $\mathbf{n}(\mathbf{m})$ due to the signal approximation error is uncorrelated with the signal component of the quantizer output $\mathbf{u}_s(\cdot)$. Color error diffusion is analyzed by assuming a matrix gain of $\tilde{\mathbf{K}}$ for the signal path and a matrix gain of $\tilde{\mathbf{I}}$ (identity matrix) for the noise path. This corresponds to using the estimator to estimate signal components in the output of the quantizer from signal components at its input, and assuming an uncorrelated noise injection to model the noise [1]. Analyzing Figs. 2 and 3 in the frequency domain using z-transforms yields the relationships [1]:

$$\mathbf{B}_s(\mathbf{z}) = \tilde{\mathbf{K}}[\tilde{\mathbf{I}} + \tilde{\mathbf{H}}(\mathbf{z})(\tilde{\mathbf{K}} - \tilde{\mathbf{I}})]^{-1}\mathbf{X}(\mathbf{z}) \tag{5}$$

$$\mathbf{B}_n(\mathbf{z}) = [\tilde{\mathbf{I}} - \tilde{\mathbf{H}}(\mathbf{z})]\mathbf{N}(\mathbf{z}) \tag{6}$$

Here, $\mathbf{B}_s(\mathbf{z})$ and $\mathbf{B}_n(\mathbf{z})$ are the z-transforms of the signal and noise components, respectively, of the vector error diffused system output. The overall system response is

$$\mathbf{B}(\mathbf{z}) = \mathbf{B}_s(\mathbf{z}) + \mathbf{B}_n(\mathbf{z}) \tag{7}$$

Equations (5) and (6) predict the linear frequency distortion (sharpening/blurring) and noise shaping effects of vector color error diffusion, respectively. The model predictions are validated both quantitatively and qualitatively in [1] on natural images and are shown to be effective in representing these "linear" effects. Nonlinear effects such as limit cycles are not modeled.

B. Linear signal frequency distortion cancellation

The matrix gain model predicts that the linear distortion suffered by the color input image is given by (5). This means that if one prefilters the input color image by using the filter

$$\tilde{\mathbf{G}}(\mathbf{z}) = [\tilde{\mathbf{I}} + \tilde{\mathbf{H}}(\mathbf{z})(\tilde{\mathbf{K}} - \tilde{\mathbf{I}})]\tilde{\mathbf{K}}^{-1} \tag{8}$$

then the resulting halftone should exhibit no signal frequency distortion with respect to the original color image. In fact, it is shown in [1] that applying this pre-filter is equivalent

to halftoning using the block diagram of Fig. 4, which feeds a linear transformation $\tilde{\mathbf{L}} = \tilde{\mathbf{K}}^{-1} - \tilde{\mathbf{I}}$ of the input image into the quantizer input. Thus, linear frequency distortion due to the use of any error filter may be cancelled by using the block diagram of Fig. 4. The feasibility of signal distortion cancellation is important in the design of color error filters since it allows one to design a color error filter for optimum noise shaping without requiring it to compensate for signal frequency distortion it introduces. Distortion elimination removes the sharpening effect of an error filter globally. Desired sharpness may be introduced using a pre-filter prior to halftoning.

III. Error filter design for optimum noise shaping

Kolpatzik and Bouman [6] designed optimum color error filters by designing optimum scalar error filters for the luminance and chrominance channels of the color error without imposing constraints on the solution. Damera-Venkata and Evans [1] generalized this result to allow matrix-valued error filters capable of shaping luminance error into the chrominance channels where the color HVS is less sensitive. A constraint on the error filter is introduced to ensure that the errors are bounded for all inputs that vary continuously over the range of the outputs. Adaptive matrix-valued filters that attempt to minimize local mean squared error have also been developed [7].

A. Formulation

Based on the matrix gain model in Section II-A, the net noise component of the output is

$$\mathbf{b}_n(\mathbf{m}) = \left([\tilde{\mathbf{I}} - \tilde{\mathbf{h}}] \star \mathbf{n} \right) (\mathbf{m}) \tag{9}$$

The next step is to quantify the impact of the noise component on the human visual system. We form an objective function J to measure the average visually weighted noise energy in the output halftone. For the color human visual system model, we use a linear spatially-invariant matrix-valued model denoted by the matrix-valued filter function $\tilde{\mathbf{v}}(\cdot)$. We define a constraint set $\mathcal{C}$ to ensure that all the quantization error (represented in a device independent RGB space) is diffused and that the system is stable [8]. Thus, the color error diffusion system $(\tilde{\mathbf{h}}(\cdot), \tilde{\mathbf{v}}(\cdot))$ for a given vision model $\tilde{\mathbf{v}}(\cdot)$ may be solved for an

optimum filter $\tilde{\mathbf{h}}_{opt}(\cdot)$

$$\tilde{\mathbf{h}}_{opt}(\cdot) = \arg \min_{\tilde{\mathbf{h}}(\cdot) \in \mathcal{C}} J \tag{10}$$

where

$$J = E\left[\| \left(\tilde{\mathbf{v}} \star \left[\tilde{\mathbf{I}} - \tilde{\mathbf{h}} \right] \star \mathbf{n} \right) (\mathbf{m}) \|^2 \right] \tag{11}$$

and

$$\mathcal{C} = \left\{ \tilde{\mathbf{h}}(\mathbf{i}),\ \mathbf{i} \in \mathcal{S} \mid \sum_{\mathbf{i}} \tilde{\mathbf{h}}(\mathbf{i}) \mathbf{1} = \mathbf{1}, \tilde{\mathbf{h}}(\mathbf{i}) \geq \tilde{\mathbf{0}} \right\} \tag{12}$$

The necessary and sufficient conditions for an optimum solution to (10) is [1]

$$\sum_{\mathbf{k}} \tilde{\mathbf{v}}^T(\mathbf{k}) \tilde{\mathbf{r}}_{(\tilde{\mathbf{v}} \star \mathbf{n})\mathbf{n}}(-\mathbf{i} - \mathbf{k}) = \sum_{\mathbf{p}} \sum_{\mathbf{q}} \sum_{\mathbf{s}} \tilde{\mathbf{v}}^T(\mathbf{s}) \tilde{\mathbf{v}}(\mathbf{q}) \tilde{\mathbf{h}}(\mathbf{p}) \tilde{\mathbf{r}}_{\mathbf{nn}}(\mathbf{p} + \mathbf{q} - \mathbf{s} - \mathbf{i}) \tag{13}$$

These equations may be regarded as a generalization of the Yule-Walker equations [9] from linear prediction theory to the matrix case, with a generalized linear spatially-invariant weighting. The above set of generalized Yule-Walker equations may be solved for the optimal filter subject to the constraints of (12) to obtain the optimum color error filter coefficients. A white noise image is used as an approximation to the uncorrelated noise image $\mathbf{n}(\mathbf{m})$ [1], [6].

The required autocorrelation matrices are approximated as

$$\tilde{\mathbf{r}}_{\mathbf{nn}}(\mathbf{k}) = E\left[\mathbf{n}(\mathbf{m}) \mathbf{n}^T(\mathbf{m} + \mathbf{k}) \right] \approx \delta(\mathbf{k}) \tag{14}$$

$$\tilde{\mathbf{r}}_{(\tilde{\mathbf{v}} \star \mathbf{n})\mathbf{n}}(\mathbf{k}) = E\left[\left(\sum_{\mathbf{t}} \tilde{\mathbf{v}}(\mathbf{t}) \mathbf{n}(\mathbf{m} - \mathbf{t}) \right) \mathbf{n}^T(\mathbf{m} + \mathbf{k}) \right] \approx \sum_{\mathbf{t}} \tilde{\mathbf{v}}(\mathbf{t}) \delta(\mathbf{k} + \mathbf{t}) = \tilde{\mathbf{v}}(-\mathbf{k}) \tag{15}$$

B. A linear color model for the human visual system

As mentioned in Section III-A, we use a linear spatially-invariant matrix-valued model $\tilde{\mathbf{v}}(\cdot)$ for the color human visual system model in the error filter optimization given by (10). The linear spatially invariant human visual system color model is based on the work of Poirson and Wandell [10]. The Poirson and Wandell model operates in the device independent XYZ space and consists of

1. A linear transformation $\tilde{\mathbf{T}}_{XYZ \rightarrow \text{Opponent}}$ from XYZ into an opponent space consisting of a luminance (Black-White) channel, a Red-Green channel, and a Blue-Yellow channel.

2. Separable spatial filtering on each channel using a different spatial filter on each channel. Thus the luminance channel is filtered less aggressively than the chrominance channels. This operation may be regarded as a matrix convolution in the frequency domain by a filter with diagonal matrix-valued coefficients $\tilde{\mathbf{d}}(\cdot)$.

Since we are starting in RGB space, we will need an additional transformation $\tilde{\mathbf{T}}_{RGB\to XYZ}$ to transform from a linear RGB space into XYZ. Hence, $\tilde{\mathbf{v}}(\mathbf{m})$ is computed as

$$\tilde{\mathbf{v}}(\mathbf{m}) = \tilde{\mathbf{d}}(\mathbf{m})\ \tilde{\mathbf{T}}_{XYZ\to \text{Opponent}}\ \tilde{\mathbf{T}}_{RGB\to XYZ} \tag{16}$$

The parameter specifications for the model including the shapes of the luminance and chrominance spatial filters are given in [11].

Monga, Geisler and Evans [12] generalized this linear color model as a linear transformation $\tilde{\mathbf{T}}$ to a desired color space, which is not necessarily the opponent representation [10] but any one that satisfies pattern color separability, followed by appropriate spatial filtering in each channel. This generalization provides a platform for evaluation of different HVS models in perceptual meaning and error filter quality obtained by minimizing (11). Based on this framework, they evaluate four color spaces [12] in which to optimize matrix-valued error filters. The objective measure used for evaluation is the noise shaping gain of the optimal filter over the Floyd-Steinberg filter in decibels [1]:

$$\text{Noise Shaping Gain} = 10\log_{10}\left(\frac{J_{fs}}{J_{opt}}\right) \tag{17}$$

Here, J refers to the value of the objective function given by (11). They also performed a subjective assessment procedure that evaluates the halftones based on a paired comparison task as described in [12]. The results of the subjective test corroborate the objective measures. The color spaces in order of increasing quality are (1) YIQ space, (2) YUV space, (3) opponent color space [10], [11], and (4) linearized CIELab color space [13]. These color spaces in conjunction with appropriate spatial filters as described in [12] form a unique HVS model. The color HVS model based on transformation to the linearized CIELab [13] color space, spatial filters for the luminance frequency response due to Nasanen and Sullivan [14], and the chrominance frequency response as given by Kolpatzik and Bouman [6] yields the best halftones. The subjective test is available online at

http://www.ece.utexas.edu/~vishal/cgi-bin/test.html

C. Results

Fig. 5(a) shows the original *toucan* image. Fig. 5(b) shows a halftone generated by applying Floyd-Steinberg error diffusion separably. The green color impulses on the red toucan are easily visible on a color monitor. Fig. 5(e) and Fig. 5(g) show the green and blue planes of the Floyd-Steinberg halftone, respectively. The color impulses on the body of the red toucan are clearly visible in the green plane. Fig. 5(c) shows a halftone generated by applying an optimum matrix-valued error filter. The green color impulses are eliminated. Fig. 5(f) and Fig. 5(h) show the green and blue planes of the optimum halftone. The green channel (which contributes greatly to luminance) does not show spurious color impulses. However, since the error is shaped into the blue-yellow channel, the blue channel of the optimum halftone has several artifacts that are not easily visible in the optimum color halftone. The signal frequency distortion produced by the optimal error filter can be cancelled using the distortion cancellation method described in Section II-B. Fig. 5(d) shows the optimum halftone after distortion cancellation. Notice that the sharpening effect of the optimum error filter has been undone resulting in an *undistorted* halftone.

IV. Quantization based on perceptual criteria

The use of the mean squared error (MSE) criterion in a colorant space is equivalent to uniform, separable, scalar quantization. The visual quantization error may be further reduced by performing the quantization according to perceptual criteria. Such methods typically aim to minimize colorimetric error, luminance variations, or a combination of the two.

A. Calorimetric quantization in error diffusion

Hanieshi *et al.* [15] suggested the use of the XYZ and Lab spaces to perform quantization and error diffusion. In this case the rendering gamut is no longer a cube. The MSE criterion in the XYZ or Lab space is used to make a decision on the best color to output. The quantization error is a vector in the XYZ space is diffused using an error filter. (Lab space is not suitable for diffusing errors due to its nonlinear variation with intensity.) This method performs better than separable quantization but suffers from boundary artifacts [15] [16] such as the "smear artifact" and the "slow response artifact" at color boundaries

due to accumulated errors from neighboring pixels pushing quantizer input colors outside the gamut. This causes a longer lag in cancelling the errors. This effect may be reduced by clipping large color errors [15] [17] or by using a hybrid scalar-vector quantization method called semi-vector quantization [16]. This method is based on the observation that when errors in colorant space are small vector quantization does not produce the smear artifact. When large colorant space errors are detected, scalar quantization is used to avoid potential smearing. First, the colorants where the colorant space error exceeds a preset threshold are determined and quantized with scalar quantization. This restricts the possible output colors from which a color must be chosen using vector quantization in device independent color space.

B. Minimization of luminance variation

This class of color error diffusion algorithms is based on the observation that luminance fluctuation in the halftone appears as visible graininess and must be minimized. For example if a mid-tone were printed with black and white dots only then there would be a large luminance variation from the black to the white. If however the tone were printed using only Cyan, Magenta and Yellow dots with minimized overlap, then there would be less white space and the overall luminance variation will be much lower resulting in reduced graininess. For inkjet printing, the overprinting of Cyan and Magenta dots is to be minimized especially in the highlights, since the luminance variation due to the composite blue dots is usually objectionable.

Klassen, Eschbach and Bharat [18] distort the colorimetric color space of vector error diffusion using penalty functions so that at a given input tone level, favored colors with low luminance variation are more likely to be chosen. Lau, Arce, and Gallagher [19] control the overlap of individual color separations using an influence matrix to transform the quantizer input, prior to quantization. Shaked *et al.* [20] restrict the number of allowable output colors for a given input color in the color space according to a Minimum Brightness Variation Criterion (MBVC) to four allowable colors (the minimum number of colors required to reproduce an input color in the average color sense), which forms a Minimum Brightness Variation Quadruple. Such a restriction allows each tone to be printed with four possible colors that produce the least luminance variation.

Threshold modulation can be incorporated into the luminance variation minimization framework to produce homogeneous color dot distributions [21], [22]. Eschbach [21] uses cross-separation threshold imprint functions derived from the input and the past outputs and thresholds to achieve homogeneous color highlights and shadows. The imprint biases a threshold in the neighborhood of a minority pixel to discourage clumping and promote regular dot spacing. Using appropriate correlated imprint functions also minimizes inter-colorplane luminance variation. Levien [22] first weights the quantizer error using multichannel weights and then lowpass filters the result to derive a quantizer threshold modulation that minimizes luminance variation due to dot overlap. Further, an output dependent threshold modulation (which is equal to the difference between the expected inter-minority pixel distance $d_{avg} = 1/g^2$ at graylevel g [23] and the actual distance to the nearest minority pixel) is used to produce a pleasing distribution of minority pixels.

C. Results

Fig. 6(a) shows a halftone image generated using vector error diffusion in the device independent XYZ space [15]. Artifacts are noticeable at color boundaries, especially on the yellow toucan. Fig. 6(b) shows a halftone image generated using vector error diffusion in the device independent XYZ space with artifact reduction by semi-vector quantization [16]. The boundary artifacts are significantly reduced. Fig. 6(c) shows a halftone generated using the MBVC quantization as described in [20]. Fig. 6(d) and Fig. 6(e) show magnified views of the MBVC halftone and the Floyd-Steinberg halftone respectively. The MBVC halftone exhibits much smoother color with significantly reduced objectionable color variation.

V. Conclusion

This article reviews some of the recent analysis and design methods for bi-level color error diffusion halftoning systems. The key ideas in the design of these systems are color noise shaping and color quantization according to perceptual criteria.

Printing is performed in a subtractive (where certain wavelengths of light are subtracted or absorbed and the reflected light is viewed) CMY color space. For printing applications a fourth color, K (black) is introduced since the composite of the CMY colorants produces

a dark brown in practice. To print an RGB image, first a printer/media dependent non-linear color transformation from RGB to CMYK is made. The transformation from device independent XYZ to CMYK is non-linear and device dependent.

This paper has focused on the essential design principles of color error diffusion halftoning systems using the RGB space for convenience (transformation to XYZ is accomplished by a linear transformation of gamma uncorrected data). These principles may be applied to the CMYK color space without loss of generality. For example, the methods of Section III could be applied to diffuse error optimally in device independent XYZ color space in combination with XYZ colorimetric quantization method. In this case the nonlinear device dependent transformation from XYZ to CMYK must be computed based on physical measurement of color patches. The MBVC error diffusion developed in RGB space could be extended to CMYK by defining appropriate vector quantizers that depend on the relative luminance of the CMYK colorants. We have omitted a discussion of multi-level color error diffusion halftoning and color palette design methods due to space constraints. A good source of references on these topics is [24].

References

[1] N. Damera-Venkata and B. L. Evans, "Design and analysis of vector color error diffusion halftoning systems," *IEEE Trans. on Image Processing*, vol. 10, pp. 1552–1565, Oct. 2001.

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

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

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

[5] H. Stark and J. W. Woods, *Probability, Random Processes and Estimation Theory for Engineers.* Englewood Cliffs, NJ: Prentice-Hall, 1986.

[6] B. Kolpatzik and C. Bouman, "Optimized error diffusion for high quality image display," *Journal of Electronic Imaging*, vol. 1, pp. 277–292, Jan. 1992.

[7] L. Akarun, Y. Yardimici, and A. E. Cetin, "Adaptive methods for dithering color images," *IEEE Trans. on Image Processing*, vol. 6, pp. 950–955, July 1997.

[8] Z. Fan, "Stability analysis of error diffusion," *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, vol. 5, pp. 321–324, Apr. 1993.

[9] T. K. Moon and W. C. Stirling, *Mathematical Methods and Algorithms for Signal Processing.* New Jersey: Prentice-Hall International, Inc., 2000.

[10] A. B. Poirson and B. A. Wandell, "Appearance of colored patterns: Pattern-color separability," *Journal of Opt. Soc. Amer. A*, vol. 10, pp. 2458–2470, Dec. 1993.

[11] X. Zhang and B. A. Wandell, "A spatial extension of cielab for digital color image reproduction," *SID Digest of Technical Papers*, pp. 731–734, 1996.

[12] V. Monga, W. S. Geisler, and B. L. Evans, "Linear, color separable, human visual system models for vector error diffusion halftoning," *IEEE Signal Processing Letters*, Apr. 2003.

[13] T. J. Flohr, B. W. Kolpatzik, R. Balasubramanian, D. A. Carrara, C. A. Bouman, and J. P. Allebach, "Model based color image quantization," *Proc. SPIE Human Vision, Visual Proc. and Digital Display IV*, vol. 1913, pp. 270–281, 1993.

[14] J. Sullivan, L. Ray, and R. Miller, "Design of minimum visual modulation halftone patterns," *IEEE Trans. on Systems, Man, and Cybernetics*, vol. 21, pp. 33–38, Jan. 1991.

[15] H. Haneishi, T. Suzuki, N. Shimonyama, and Y. Miyake, "Color digital halftoning taking colorimetric color reproduction into account," *Journal of Electronic Imaging*, vol. 5, pp. 97–106, Jan. 1996.

[16] Z. Fan and S. Harrington, "Improved quantization methods in color error diffusion," *Journal of Electronic Imaging*, vol. 8, pp. 430–437, Oct. 1999.

[17] C. Kim, I. Kweon, and Y. Seo, "Color and printer models for color halftoning," *Journal of Electronic Imaging*, pp. 166–180, June 1997.

[18] R. Klassen, R. Eschbach, and K. Bharat, "Vector error diffusion in a distorted color space," *Proc. IS&T Conf.*, 1994.

[19] D. L. Lau, G. R. Arce, and N. C. Gallagher, "Digital color halftoning with generalized error diffusion and multichannel green-noise masks," *IEEE Trans. on Image Processing*, vol. 9, pp. 923–935, May 2000.

[20] D.Shaked, N. Arad, A.Fitzhugh, and I. Sobel, "Color diffusion: Error-diffusion for color halftones," *HP Labs Technical Report, HPL-96-128R1*, 1996.

[21] R. Eschbach, "Pseudo-vector error diffusion using cross separation threshold imprints," *Proc. IS&T Sym. on Electronic Imaging*, pp. 321–323, 1999.

[22] R. Levien, "Practical issues in color inkjet halftoning," *Proc. SPIE Color Imaging: Processing, Hardcopy and Applications VIII*, pp. 537–541, Jan. 2003.

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

[24] G. Sharma and H. Trussell, "Digital color imaging," *IEEE Trans. on Image Processing*, vol. 6, pp. 901–932, July 1997.

Niranjan Damera-Venkata received the B.E. degree in Electronics and Communication Engineering from the University of Madras in Madras, India, in July 1997, and the M.S. and Ph.D. degrees in Electrical Engineering from The University of Texas at Austin in May 1999 and December 2000, respectively. Since July 2000, he has been with the Imaging Systems Laboratory at Hewlett-Packard Laboratories in Palo Alto, CA, where he is currently a research engineer. His research interests include multimedia understanding, digital halftoning, secure printing and digital camera image processing. Dr. Damera-Venkata is a member of Sigma Xi.

Brian L. Evans received his BSEECS degree from the Rose-Hulman Institute of Technology in May 1987, and his MSEE and Ph.D. degrees from the Georgia Institute of Technology in December 1998 and September 1993, respectively. From 1993 to 1996, he was a post-doctoral researcher at the University of California, Berkeley. From 1996 to 2000, he was an Assistant Professor in the Dept. of Electrical and Computer Engineering at The University of Texas at Austin. He is currently an Associate Professor of ECE at UT Austin. His research interests include the design and real-time software implementation of high-quality halftoning for desktop printers, smart image acquisition for digital still cameras, 3-D sonar imaging systems, and ADSL/VDSL transceivers. Dr. Evans has published over 100 refereed conference and journal papers. Dr. Evans is a member of the Design and Implementation of Signal Processing Systems Technical Committee of the IEEE Signal Processing Society, and a recipient of a 1997 National Science Foundation CAREER Award.

Vishal Monga received the B. Tech. degree in Electrical Engineering from the Indian Institute of Technology, Guwahati, India, in May 2001, and the MSEE degree from The University of Texas at Austin in May 2003. He is currently a PhDEE student at UT Austin. His research interests lie in the intersection of signal processing and information theory. He received a 2002-2003 Texas Telecommunications Engineering Consortium Graduate Fellowship from UT Austin.

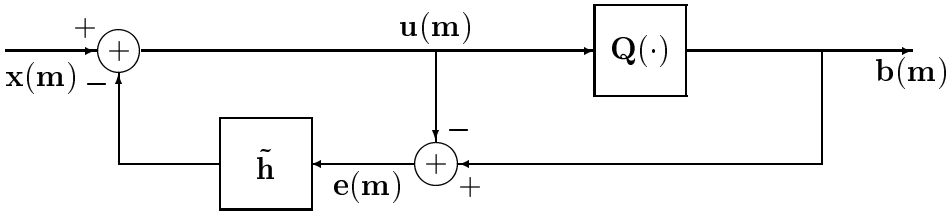

Fig. 1. System block diagram for vector color error diffusion halftoning. Here, $\tilde{\mathbf{h}}$ represents a fixed 2-D FIR error filter with scalar or matrix valued coefficients. The vector $\mathbf{m}$ represents the 2-D index (m_1, m_2). The quantizer $\mathbf{Q}(\cdot)$ performs scalar or vector quantization.

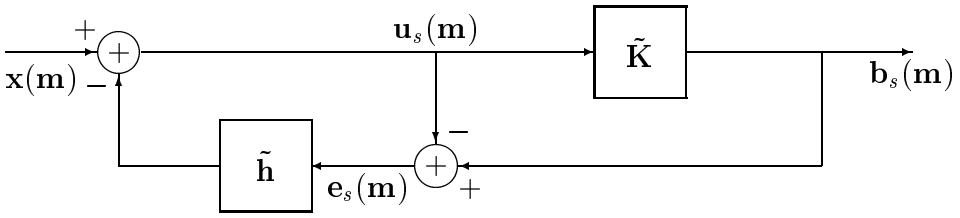

Fig. 2. System block diagram for the signal path obtained by substituting the matrix gain model for quantizer. Here, $\mathbf{z}_s(\mathbf{m})$ represents the signal component of the variable $\mathbf{z}(\mathbf{m})$. The matrix gain $\tilde{\mathbf{K}}$ represents a linear transformation of the signal component $\mathbf{u}_s(\mathbf{m})$ of $\mathbf{u}(\mathbf{m})$.

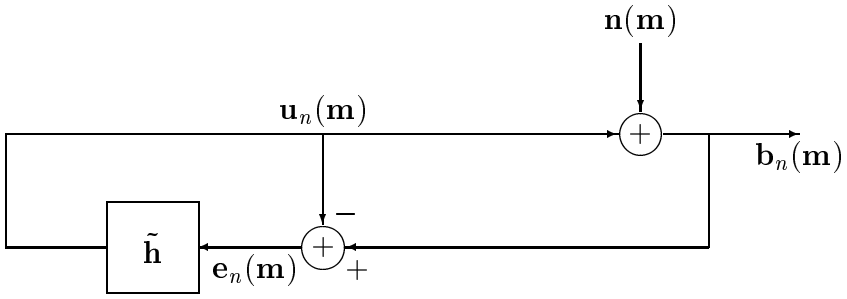

Fig. 3. System block diagram for the noise path obtained by replacing the quantizer with an additive noise source. The noise injection $\mathbf{n}(\mathbf{m})$ is a noise process that is uncorrelated with $\mathbf{u}_s(\mathbf{m})$ in Fig. 2. Here, $\mathbf{z}_n(\mathbf{m})$ represents the noise component of the variable $\mathbf{z}(\mathbf{m})$. $\mathbf{u}_n(\mathbf{m})$ represents the noise component of $\mathbf{u}(\mathbf{m})$.

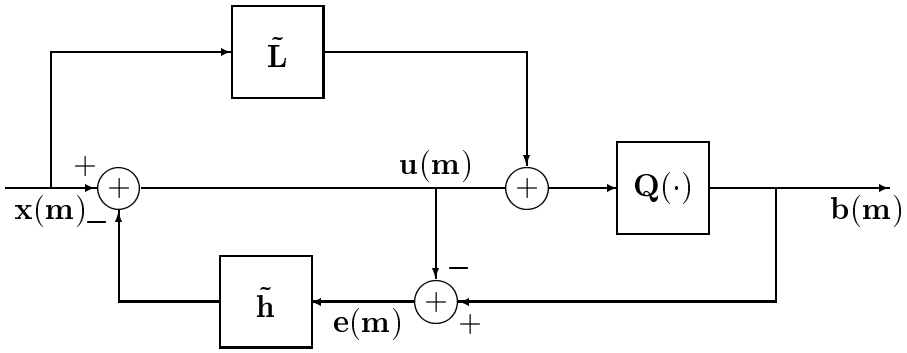

Fig. 4. System block diagram that modifies the vector color error diffusion halftoning system in Fig. 1. $\tilde{\mathbf{L}}$ represents a constant linear transformation (matrix gain) that controls the amount of sharpening in the color planes.

(a) Original image

(b) Floyd-Steinberg error filter

(c) Optimum error filter

(d) Signal distortion cancellation

(e) Green plane of Floyd-Steinberg halftone

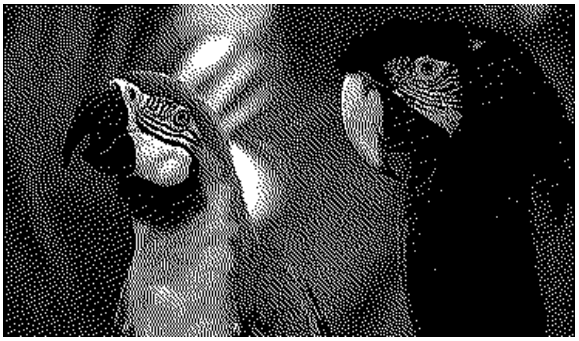

(f) Green plane of optimum halftone

(g) Blue plane of Floyd-Steinberg halftone

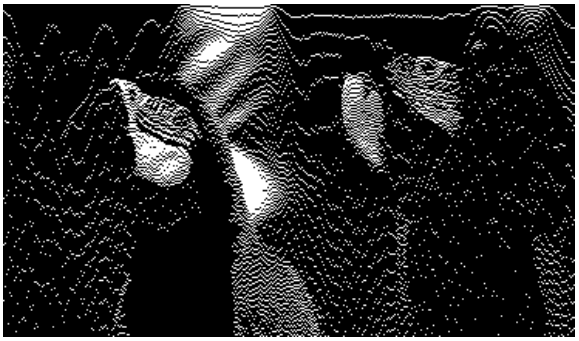

(h) Blue plane of optimum halftone

Fig. 5. Signal and noise shaping in color error diffusion.

(a) Vector error diffusion in XYZ space

(b) Boundary artifact reduction

(c) MBVC error diffusion

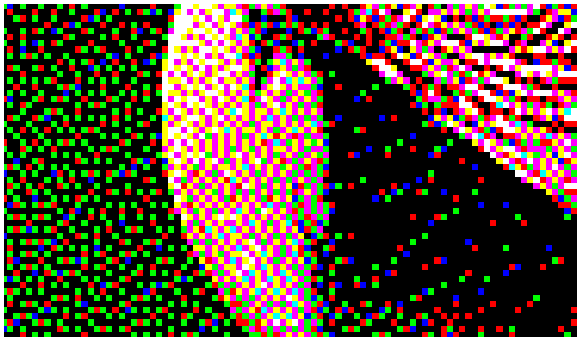

(d) Detail of MBVC

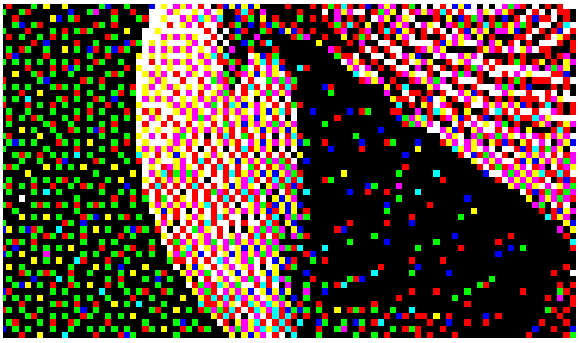

(e) Detail of Floyd-Steinberg halftone

Fig. 6. Color quantization in color error diffusion. Minimum Brightness Variation Criterion (MBVC) error diffusion uses separable Floyd-Steinberg error filters.