work, which results mostly in smoother looking images, with less noise and texture appearance.

Conclusion

We examined the design of a universal color palette in a visually uniform color space. The palette is generated using vector quantization methods, but also allows fast quantization, due to its sequential structure. In combination with our visually optimized error diffusion algorithm, we obtained a significant quality improvement over basic error diffusion with simple MSE quantization, operating separately on RGB coordinates.

Acknowledgements

This work was supported by an NEC Faculty Fellowship.

References

1. R. Balasubramanian, J. P. Allebach, "A New Approach to Palette Selection for Color Images," *J. Imaging Technol.*, v.17, no.6, 1991, pp.284 - 290.
2. R. S. Gentile, E. Walowit, J. P. Allebach, "Quantization and multi-level Halftoning of Color Images for Near Original Image Quality," *J. Opt. Soc. Am. A.*, vol. 7, no.6, pp.1019-1026, 1990.
3. M. T. Orchard, C. A. Bouman, "Color Quantization of Images," *IEEE Trans. Signal Processing*, vol. 39, no. 12, pp.2677 -2690, 1991.
4. G. Braudaway, "A Procedure for Optimum Choice of a Small Number of Colors from a Large Color Palette for Color Imaging," IBM internal report, RC 11367(\#51227), Sept. 16, 1985.
5. G. Goertzel and G. R. Thompson, "'Halftoning' Techniques for Displaying Images with a Limited Color Palette," EI West 1990, Pasadena, CA, pp.102-108.
6. D. Venable, J. Stinehour, P. Roetling, "Selection and Use of Small Color Sets for Pictorial Display," SPSE 43rd Annual Meeting, Rochester, NY, pp. 90 - 92 (1990).
7. B. W. Kolpatzik, C. A. Bouman, "Optimized Error Diffusion for Image Display," *J. Electronic Imaging*, vol.1, no.3, pp.277-292, 1992.
8. R. Balasubramanian, C. Bouman, and J. Allebach, "Sequential Scalar Quantization of Vectors: An Analysis'" to be submitted to the *IEEE Trans. on Image Processing*.
9. R. Balasubramanian, C. Bouman, and J. Allebach, "Sequential Scalar Quantization of Color Images," to be submitted to the *Journal of Electronic Imaging*.
10. A. Gersho, R. M. Gray, *Vector Quantization and Signal Compression*, Kluwer Academic Publishers, 1992.
11. R. W. Floyd and L. Steinberg, "An Adaptive Algorithm for Spatial Greyscale," Proc. SID, vol.17, no.2, pp.75-77, 1976.

Reprinted from the *Proceedings of IS&T's 46th Annual Conference*, May 9-14, 1993, Cambridge, Massachusetts.

Output Dependent Feedback in Error Diffusion Halftoning

Raph Levien, Levien Instrument Co., McDowell, Virginia

Abstract

Error diffusion halftoning is extended by modulating the threshold by an output dependent feedback term. The output dependent feedback can control the halftone patterns and texture with a great degree of power and precision, and a minimum of computational expense. Two particular types of output dependent feedback are discussed in detail: hysteresis and nearest dot distance. A hysteresis term clusters like colored pixels together, forming a coarser halftone less susceptible to degradation in the output marking stage. The hysteresis factor allows precise control of the coarseness. A nearest dot distance term, based on the distance from the pixel under consideration to the nearest previously generated dot, improves the visual texture of highlight tones, eliminating the wormy textures typical of error diffusion halftones. With a nearest dot distance term, highlight tones closely approximate placing dots at the centers of imaginary tightly packed circles. In the transition zone between highlights and midtones, the nearest dot distance term gradually becomes insignificant and is dominated by the underlying error diffusion pattern. These two terms can be used separately or together. In addition to these specific applications, it is felt that output dependent feedback offers a powerful and general technique for controlling the error diffusion halftoning process.

Background

The traditional model of error diffusion halftoning[8] is essentially a feedback loop, as shown in Figure 1. A threshold comparator quantizes the image, resulting in an error, which is then filtered and summed into the loop as negative feedback. In this way, the halftone realizes the goal of minimizing the grayscale error. In this paper, we demonstrate how to improve the quality by adding an additional feedback term based directly on the output halftone, rather than just the error between the output halftone and the input image. The augmented error diffusion process is shown in Figure 2. The additional term is handled as a form of threshold modulation. Knox and Eschbach[2] have analyzed the question of threshold modulation in general.

In this paper, we apply the technique of output dependent feedback to address two classic problems in error diffusion, namely the wormy texture in highlights and shadows, and the production of overly fine-grained texture, which is subject to degradation in theprinting process.

Texture in Highlights and Shadows

Error diffusion is plagued by a "wormy" texture in highlights and shadows. Ulichney[4] characterizes this mathematically as anisotropy in the frequency spectra of the resulting halftone, and provides a number of techniques for reducing the effect, primarily the introduction of random noise to the threshold in the error diffusion process. Knox[9] analyzes this introduction of random noise and finds it to be equivalent to the introduction of filtered random noise to the image. Even with the added random noise, the resulting halftones still exhibit visually disturbing non-uniform texture.

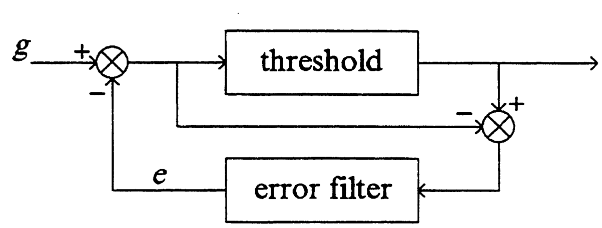

Figure 1. Standard error diffusion.

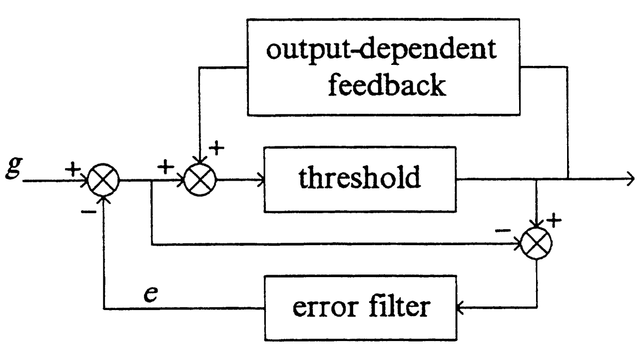

Figure 2. Error diffusion with output dependent feedback.

Other techniques[1,5,6,7] use more more complicated methods to make the texture more uniform, typically employing many iterations to produce the halftone. Such techniques are very computationally intensive, and may have disadvantages such as distorting the modulation tansfer function (MTF).

One example of a highly uniform texture is the placement of single-pixel dots in a hexagonal grid, as in Figure 3. This pattern is suitable for a constant, flat shade of gray, but not for gradations or detailed images. We will generalize this pattern first to gradations of highlight tones, then to images with arbitrary distributions of grayscale and detail.

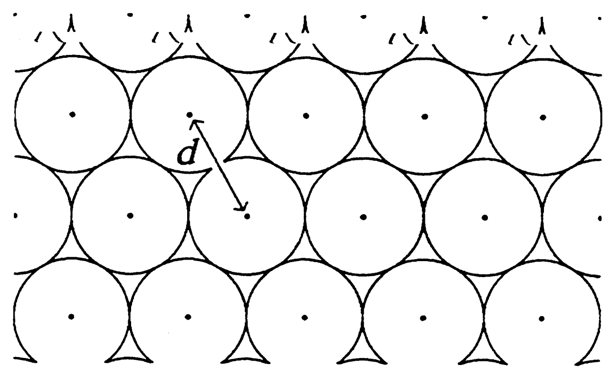

Figure 3. A hex grid of dots with surrounding circles.

To generalize to gradations of highlight tones, we first make the observation that in a hex grid with a spacing between dots of d, the dots correspond to the centers of tightly packed circles of diameter d. Thus, one algorithm for generating a hex grid of dots is to tightly pack circles of constant diameter on the plane, and draw the dots at the centers of the circles. Then, the algorithm is extended to pack circles of arbitrary diameters. The expected result, as in Figure 4, is a gradation of highlight tones with visually neutral texture. It may be interesting to note that the process of packing circles on the plane is analogous to the process observed in nature of collision based placement of, for example, the seeds in a sunflower or the placement of individual berries of a multi-berry fruit such as a raspberry[3].

In bilevel digital halftones, the positions of the dots are constrained to integer grid coordinates, thus only approximating the ideal packed circles of the previous paragraph. We can generate such an approximate halftone in raster scan order by applying the following steps for each pixel under consideration.

Consider the previously generated dots in the neighborhood of the present pixel under consideration. Find the distance to the nearest such dot.

If the distance exceeds d, place a dot at the present pixel position.

A particularly efficient implementation uses finite differences to compute the distances, and uses serpentine scan order[4,9] rather than raster scan. Such an implementation was used to produce Figure 4.

One problem remains, namely how to calculate the value of d for a given gray shade. An approximate answer is to set d equal to the reciprocal of the square root of the gray shade. (Note that this formula becomes singular when the gray value is white, i.e. zero. The singularity is avoided by simply never placing a dot at a pixel where the input gray value is zero, and thus skipping the computation).

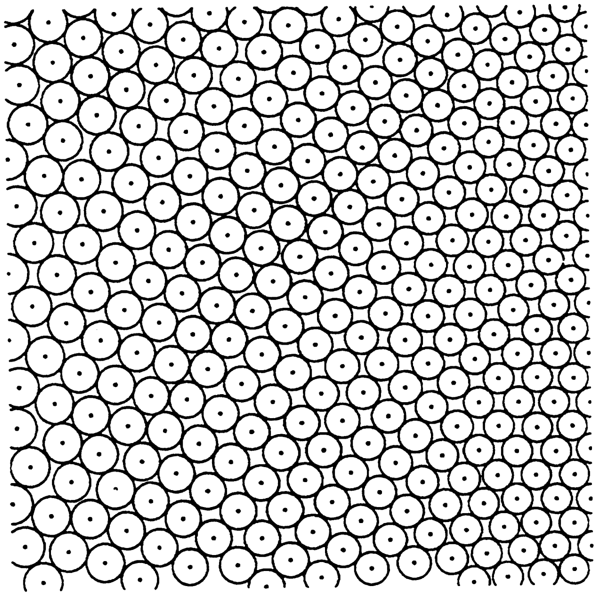

Figure 4. Dots with surrounding circles of different radii.

A more accurate solution to the value of *d*, however, is obtained by combining the packing of circles with the traditional error diffusion technique. Note particularly the similarity of the comparison in the second step with the thresholding of the image in the error diffusion process. We replace both of these steps with a single comparison that includes both a distance term and a grayscale error term. In this way, the grayscale error is used as feedback to iteratively solve for the correct value of *d* needed to reproduce the image grayscale.

The complete algorithm is given below:

Repeat for each pixel in serpentine scan order:
- Let g be the gray value of the input pixel.
- If $g = 0$ (i.e. white) generate a white pixel and skip to the next pixel.
- Let e be the value of the error accumulated to this pixel.
- Let d be the distance to the nearest previously generated dot.
- Test the value of $g + e + (d^2 - 1/g) \bullet c_1$.
- If positive, mark a dot at the present pixel.
- Calculate the error between $g + e$ and the output pixel.
- Filter and distribute to upcoming pixels.

Discussion of the Algorithm

There are a few points worth considering in the above algorithm. First, the value of d^2 was used in the threshold calculation, rather than simply d. The reason was to make the d-based term dominate for highlights (say 0-10%), and gradually diminish as the grayscale increases. By the time the gray value has reached 25% or so, the error diffusion term should dominate, producing textures similar to the traditional error diffusion technque. Between the highlights and quarter tones, there is a smooth transition, avoiding a visually noticable splice between the packed circle texture and the error diffusion texture[1].

Second, we use an initial guess of $d^2 = 1/g$, which is then refined by the error diffusion process, rather than start with an initial guess of, say, zero and letting the error diffusion bear all of the burden. The grayscale error need only compensate for the difference between the correct value and the initial guess, thus minimizing the resulting grayscale error.

Third, we have used a constant factor of c_1 to weight the distance term relative to the error term. A typical value is 0.01, as determined by experiment.

Fourth, the algorithm as shown treats only highlight dots. Shadow dots can be treated using a symmetrical arrangement. Often in practice, however, textures in highlights are more important, largely due to the dot gain in the printing device.

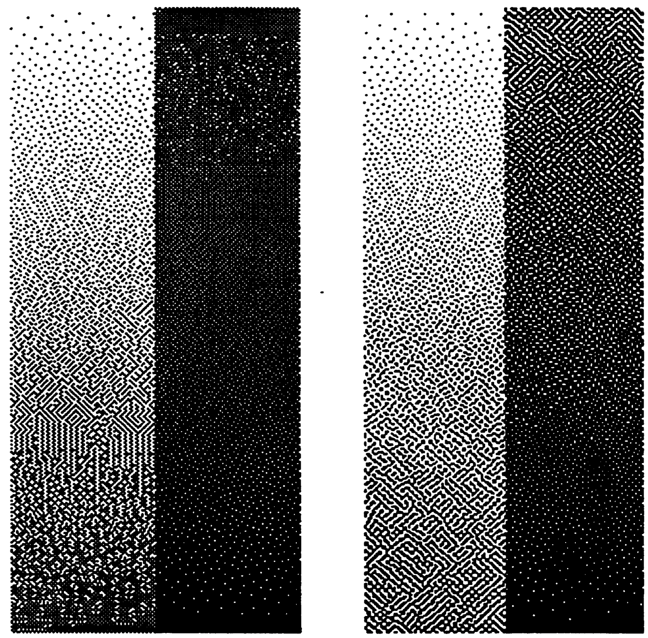

Figure 5. (a) a grayscale ramp (b) with hysteresis term..

Figure 5a shows a grayscale ramp halftoned with the above algorithm. The error filtering used for Figure 5a is the deterministic two weight error diffusion of Ulichney[4], on a serpentine raster.

Coarseness of Texture

Most error diffusion techniques are geared toward producing halftone textures with energy predominantly in the very high spatial frequencies, or in other words, "blue noise." Under the assumption of an ideal printing device, this leads to the highest quality halftones. However, under non-ideal printing conditions, the extremely fine textures cannot be accurately reproduced. A notorious example is the "broken checkerboard" pattern that appears in many halftone patterns in 45-55% gray areas[10]. This problem becomes much more severe as the resolution improves. For example, standard error diffusion produces extremely poor results on 600 dpi laser printers[11].

Our approach[12] to this problem is to deliberately coarsen the halftone textures, in effect shifting the spatial frequencies from "blue" to "green" noise. We accomplish this by adding an additional output dependent feedback term, namely the sum of the values of the two previously generated immediate neighbors of the pixel under consideration. The weight of this term is called the hysteresis constant h. The larger the value of h, the coarser the resulting textures.

The hysteresis term works by biasing the threshold to produce pixels of the same value as the previously generated neighboring pixels. The grayscale error builds up over a larger area until it becomes great enough to overcome the effects of the hysteresis term. Thus, the regions of like-valued pixels become larger and the resulting textures are coarser.

The optimum value of h can be determined by experiment, but can also be understood intuitively. The visually objectionable textures in the resulting halftone come from two sources, namely the textures inherent in the digital halftone and the textures arising from the non-ideality of the printer. For $h = 0$, the textures from the digital halftone are minimized, but the printer produces very bad patterning, perceived as graininess. As h increases, the halftone textures gradually become more objectionable, while the texturing from the printer rapidly diminishes. For 600 dpi laser printers, the optimum value of h is approximately 0.75.

For comparison, Figure 5b shows the same grayscale ramp as Figure 5a, but with the hysteresis term added ($h = 0.4$). An example 600 dpi halftone using both the hysteresis and nearest dot distance terms is shown in Figure 6.

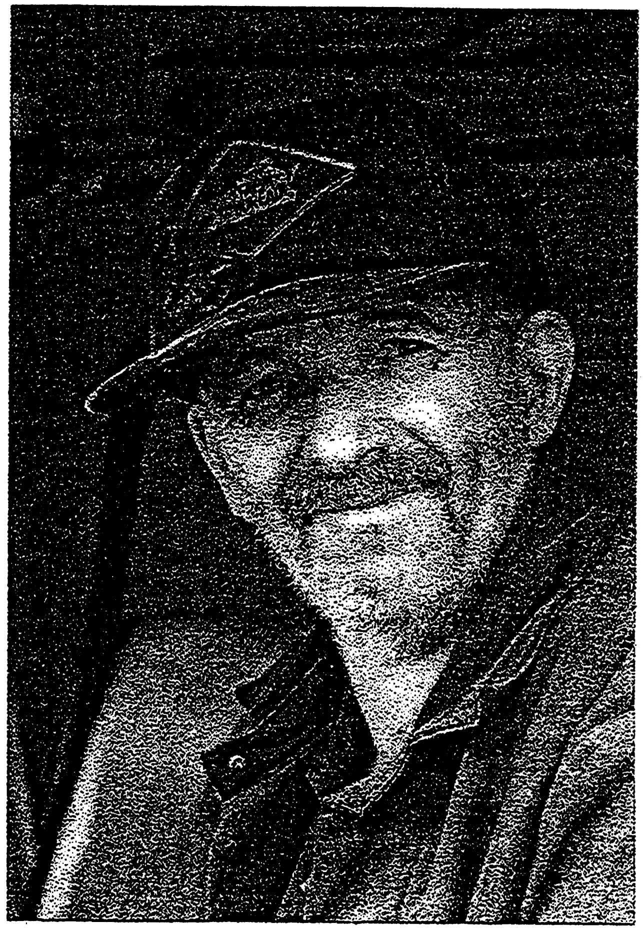

Fig. 6. A halftone screened with nearest distance and hysteresis terms. The hysteresis constant is 0.875.

Conclusion

The use of output dependent feedback allows a great deal of control over the textures of error diffusion halftones, with minimal compuational cost. In this paper, we have shown how to improve error diffusion halftones in two ways, first by making the highlight textures visually uniform, and second by adjusting the coarseness of the textures for optimum reproduction on non-ideal printing devices. In addition to these applications, we feel that output dependent feedback will prove to be a powerful tool to assert control over the textures produced by error diffusion processes.

References

1. D. A. Carrara, M. Analoui, and J. P. Allebach, Recent Progress in Digital Halftoning, *Proc. IS&T Eighth International Congress on Advances in Non-Impact Printing Technologies,* pp. 265-270, Oct. 1992.
2. K. Knox and R. Eschbach, Analysis of Threshold Modulation in Error Diffusion, *Proc. IS&T Eighth International Congress on Advances in Non-Impact Printing Technologies,* pp. 280-282, Oct. 1992.
3. D. Fowler, P. Prusinkiewicz, and J. Battjes, A Collision-based Model of Spiral Phyllotaxis, *SIGGRAPH '92 Conf. Proc.,* pp. 361-368, Jul. 1992.
4. R. Ulichney, *Digital Halftoning,* The MIT Press, 1987.
5. R. Eschbach and R. Hauck, A 2-D Pulse Density Modulation by Iteration for Halftoning, *Opt. Comm,* Vol. **62,** pp. 300-304, June 1987.
6. R. Eschbach, Pulse-Density Modulation on Rastered Media: Combining Pulse-Density Modulation and Error Diffusion, *J. Opt. Soc. Am. A,* Vol. **7,** pp. 708-716, April 1990.
7. T. N. Pappas, D.L. Neuhoff, and N. Seshadri, Least-Squares Model-Based Halftoning, *Proceedings of the 1992 SPIE/IS&T Symposium on Electronic Imaging Science and Technology,* San Jose, CA, February 9-14, 1992.
8. R. Floyd and L. Steinberg, An Adaptive Algorithm for Spatial Grey Scale, *SID International Symposium Digest of Technical Papers,* Vol. VI, pp. 36-37, 1975.
9. K. Knox, Error Diffusion: A Theoretical View, *Proceedings of the 1993 IS&T/SPIE Symposium on Electronic Imaging Science and Technology,* San Jose, CA, Jan 31-Feb 4, 1993.
10. C. Rosenberg, Measurement Based Verification of an Electrophotographic Printer Dot Model for Halftone Algorithm Tone Correction, *Proc. IS&T Eighth International Congress on Advances in Non-Impact Printing Technologies,* pp. 286-291, Oct. 1992.
11. Q. Lin, Halftone Image Quality Analysis Based on a Human Vision Model, *Proceedings of the 1993 IS&T/SPIE Symposium on Electronic Imaging Science and Technology,* San Jose, CA, Jan. 31-Feb. 4, 1993.
12. R. Levien, *Photographic Image Reproduction Device Using Digital Halftoning to Screen Images Allowing Adjustable Coarseness,* U.S. Patent 5,055,942, October 8, 1991.

Reprinted from the *Proceedings of IS&T's 46th Annual Conference,* May 9-14, 1993, Cambridge, Massachusetts.