Journal of Electronic Imaging 9(2), 151–158 (April 2000).

Halftone postprocessing for improved rendition of highlights and shadows

Clayton Brian Atkins[a)]
Hewlett-Packard Company
Hewlett-Packard Laboratories
Imaging Technology Department
Palo Alto, California 94304-1126
E-mail: cbatkins@hpl.hp.com

Jan P. Allebach
Charles A. Bouman
Purdue University
School of Electrical and Computer Engineering
Electronic Imaging Systems Laboratory
West Lafayette, Indiana 47907-1285

Abstract. *Many binary halftoning algorithms tend to render extreme tones (i.e., very light or very dark tones) with objectionable dot distributions. To alleviate this artifact, we introduce a halftone postprocessing algorithm called the Springs algorithm. The objective of Springs is to rearrange minority pixels in affected regions for a smoother, more attractive rendition. In this paper, we describe the Springs algorithm, and we show results which demonstrate its effectiveness. The heart of this algorithm is a simple dot-rearrangement heuristic which results in a more isotropic dot distribution. The approach is to treat any well-isolated dot as if it were connected to neighboring dots by springs, and to move it to a location where the energy in the springs is a minimum. Applied to the whole image, this could degrade halftone appearance. However, Springs only moves dots in selected regions of the image. Pixels that are not minority pixels are not moved at all. Moreover, dot rearrangement is disabled on and around detected edges, since it could otherwise render those edges soft and diffuse.* © 2000 SPIE and IS&T.
[S1017-9909(00)00102-1]

1 Introduction

Digital halftoning is the process of transforming a ''continuous-tone'' digital image into one which may be rendered by a device having a relatively limited palette.[1] In the case of binary halftoning, the intended output device can only produce two colors—usually black and white. This case is especially important because of its relevance in the field of printing.

One drawback of many common binary halftoning algorithms is that they often render extreme tones (i.e., very light and very dark tones) with objectionable dot distributions. One such algorithm is error diffusion,[2] which is known to produce a strongly directional ''worming'' artifact. Another example is screening using traditional dispersed-dot dither arrays (see Ref. 3, for example), which produce artificially periodic patterns.

What makes these behaviors so objectionable is that they suggest the presence of texture in regions where there should not be any texture. When it comes to printing binary halftone images, this is especially problematic in the case of light tones, where the minority pixels—black dots against the white background[4]—tend to be more visible. This is generally less of a problem in the dark tones, since the visibility of individual minority pixels—white dots against the black background—can be significantly reduced because of factors collectively referred to as ''dot gain.''

A variety of novel halftoning methods have addressed this problem. Included among them are methods for modified error diffusion,[5,6] dither array generation[7–10] two-dimensional pulse-density modulation,[11] and halftoning based on tree coding,[12] which specifically address the problem of how minority pixels are arranged in regions of extreme tone. These techniques work by controlling the spacing among minority pixels in order to avoid the artifacts described above. Despite the satisfactory performance of many of these methods, error diffusion (as it was originally proposed) and traditional dither arrays are still in wide use.

In this paper we introduce a halftone postprocessing technique called the *Springs* algorithm, which has been developed to suppress the above-described artifacts. The objective of *Springs* is to rearrange minority pixels in affected regions, in order to give a smoother, more attractive rendition.

An important aspect of *Springs* is that it is separate and apart from whatever halftoning algorithm was used to generate the input halftone image. Moreover, *Springs* works without any knowledge of the continuous-tone image from which the input halftone was generated. This means that *Springs* is a halftone postprocessing algorithm strictly, and it distinguishes *Springs* from halftoning algorithms (or half-

[a)]This work was performed while Clayton Brian Atkins was with Purdue University.

Paper 99006 received Feb. 25, 1999; revised manuscript received Nov. 22, 1999; accepted for publication Jan. 31, 2000.
1017-9909/2000/$15.00 © 2000 SPIE and IS&T.

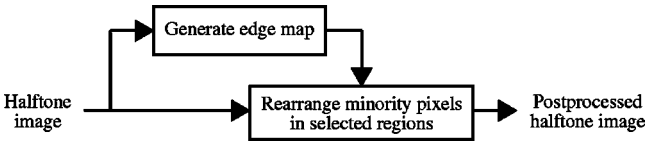


Fig. 1 Structure of the *Springs* algorithm.

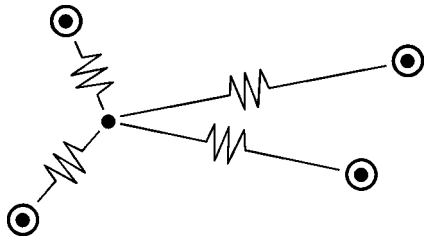


Fig. 2 Central dot connected to neighboring dots by springs. (Neighboring dots are inside circles.)

tone optimization algorithms) that effectively perform halftone postprocessing such as in Refs. 13–15, and 16. *Springs* does share some similarity with the algorithm of Eschbach and Hauck described in Ref. 13, since both involve halftone postprocessing based on treating minority pixels as if they were connected to neighboring minority pixels by imaginary springs. The difference is that the algorithm of Ref. 13 uses knowledge of the continuous-tone input image, while *Springs* does not.

To our knowledge, there is very little prior published work in the area of halftone postprocessing for improved rendition of extreme tones. One exception is Ref. 17, where we describe an algorithm which is very similar to the one presented here. Other work in the area of halftone postprocessing relates to the important but different field of digital document restoration.[18,19]

The remainder of this paper is organized as follows. In the next section, we describe the general structure of the *Springs* algorithm. Then, in Sec. 3, we give a detailed description of the dot-rearrangement procedure which forms the basis of *Springs*. In Sec. 4, we describe the scheme that is used to detect edge regions in the halftone, where dot rearrangement is not allowed. Finally, in Sec. 5, we show some experimental results; and in Sec. 6, we give some conclusions.

2 The *Springs* Algorithm

The *Springs* algorithm works by first identifying regions where postprocessing may be warranted, and then rearranging minority pixels in those regions for a more homogeneous distribution. The two stages of this process are depicted in Fig. 1.

Note that this involves moving black dots in light regions, and it also involves moving white dots in dark regions. In this paper, we will use the term ''dot'' to refer to a minority pixel in a generic sense. There is little room for confusion here, since in regions of extreme tone, the meaning of the term *minority pixel* is unambiguous; and we will see in Sec. 3 that we do not even move pixels that are not minority pixels.

The movement of a dot amounts to the minimization of a cost functional. This cost functional represents the energy stored in virtual springs that connect the dot to some neighboring dots, as depicted in Fig. 2. To minimize this functional, we allow the central dot to move, while the neighboring dots remain fixed. Intuitively this minimization has the effect of relaxing the dot distribution, which results in a more homogeneous texture. In order to evaluate the functional, we assume that the ''rest length'' of each of the springs is equal to the mean of the distances between the dot (at its starting location) and its neighbors. This means that the form of the functional actually adapts to the initial dot distribution.

Prior to the dot-rearrangement stage, *Springs* generates a binary map identifying edge regions, where dot rearrangement will not be allowed. This is important since subjecting edges to the dot-rearrangement process could render them soft and diffuse.

3 Dot Rearrangement

The dot-rearrangement stage of the *Springs* algorithm is iterative. Each iteration consists of one raster-ordered pass through the lattice of pixel locations in the image. In the turn of each pixel location, we consider the pixel there as a candidate for being moved; and we move it, provided that certain conditions are met. These conditions are listed in the following paragraph. The form of the cost functional depends upon the position of the dot relative to the positions of N neighboring dots. An important aspect of the algorithm is the procedure by which we choose those neighbors.

In order for a pixel to be moved, it must satisfy three criteria. First, it must be spatially separated from any edges detected in the halftone. These regions are identified by a binary edge map. (We will describe how the edge map is generated in the next section.) Second, the pixel must not be adjacent to any other pixel of the same color. This is important for the preservation of delicate but continuous structures like thin lines and serifs in text. Finally, the mean of the distances between the pixel and its N neighboring pixels of the same color must be greater than some threshold value. We use a threshold value of three pixel units; so that loosely speaking, postprocessing is disabled in any region having dot coverage less than 3.4%, or greater than 96.6%. This is a requirement since the rearrangement process does not improve dot distributions that are too densely packed. Note that these last two rules effectively ensure that we will not move pixels that are not minority pixels.

3.1 *Cost Functional Minimization*

The dot movement is implemented as a sequence of short traversals which terminates at a local minimum of the cost functional. In the $(k+1)$th traversal, the dot is moved from its location $n^{(k)}$ to the adjacent location $n^{(k+1)}$ which decreases the cost the most. (The set of adjacent locations is comprised of the eight nearest neighbor pixels.) Any location which has no adjacent location that decreases the cost is a local minimum. The only restriction to this process is that the dot may not be moved into a location having another dot adjacent to it. We describe our motivation for this restriction below.

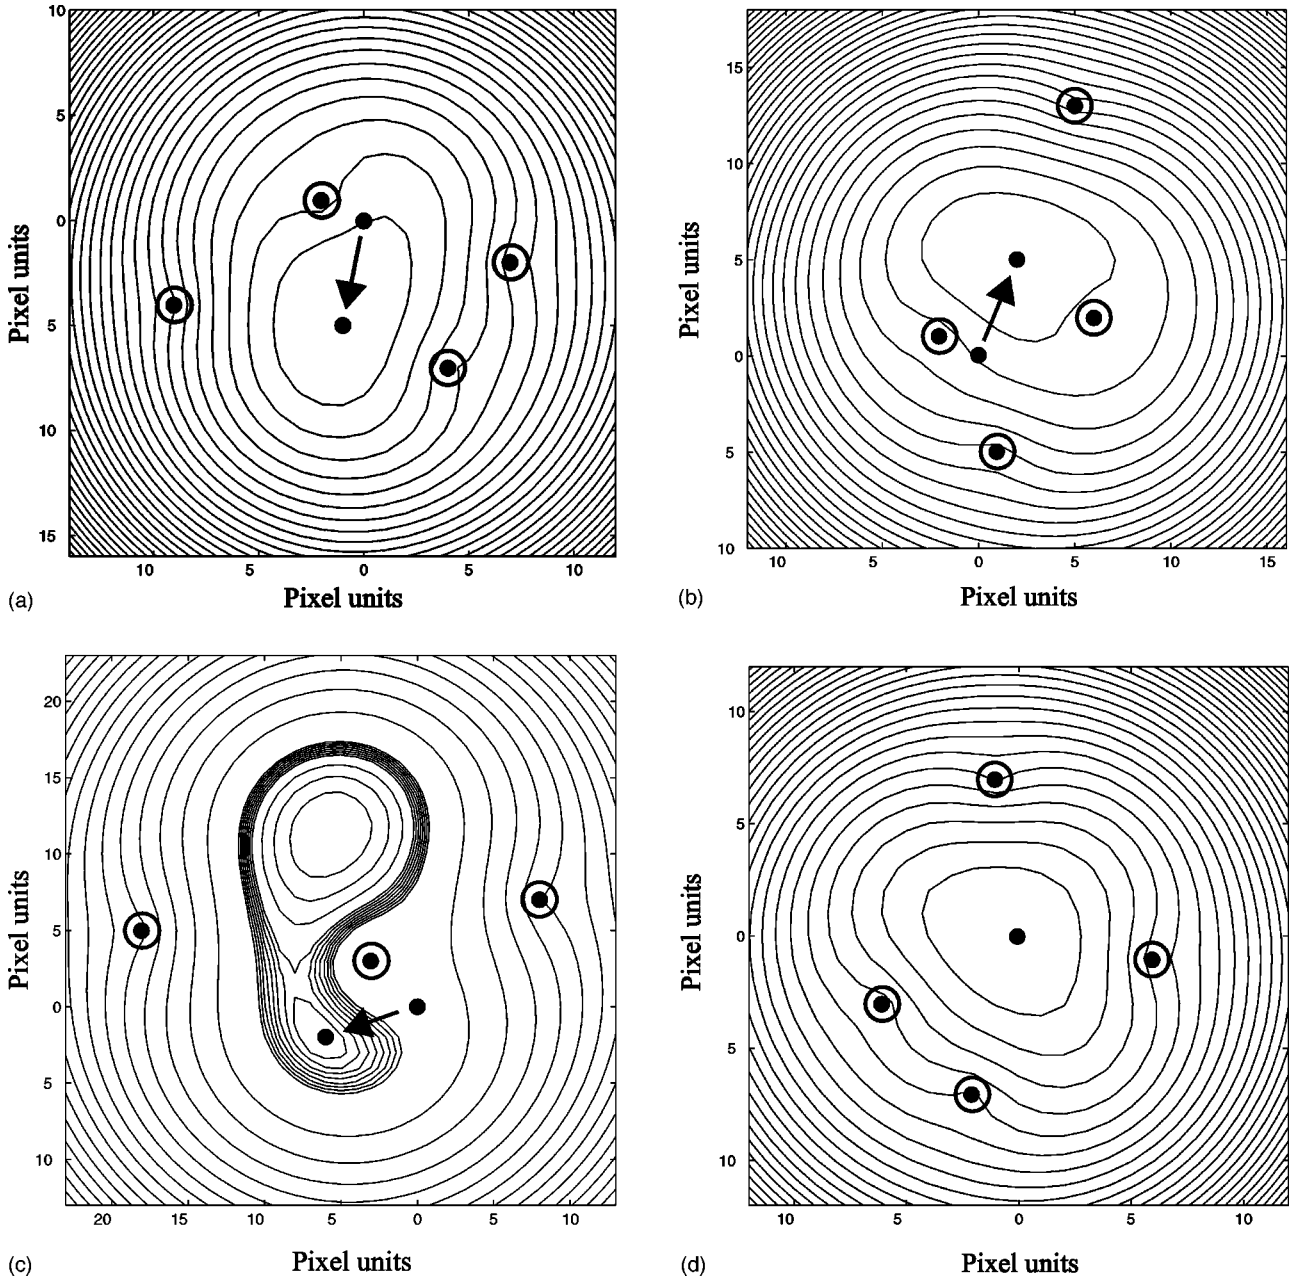


Fig. 3 Example dot-movement scenarios. Curves represent level sets of the cost functional, which is defined by the relative positions of neighboring dots. (Neighboring dots are inside circles.)

To evaluate the cost functional, it is only necessary to add up the energy stored in imaginary springs which connect the dot to its neighbors. If the dot is at location n, then the total energy $e(n)$ is

$$e(n)=\sum_{i=1}^{N} e_i(n), \tag{1}$$

where $e_i(n)$ is the energy in the spring connected to the ith neighboring dot. This is computed by squaring the magnitude of its displacement. Specifically

$$e_i(n)=|d_i(n)-r|^2, \tag{2}$$

where $d_i(n)$ is the distance between n and the location n_i of the ith dot

$$d_i(n)=\|n-n_i\| \tag{3}$$

and r is the rest length of the spring.

We assume that the rest length r is the same for each of the springs. It is computed as the mean of the distances between the initial dot location and the neighboring dots; specifically, r is computed as

$$r=\frac{1}{N}\sum_{i=1}^{N} \|n^{(0)}-n_i\|, \tag{4}$$

where $n^{(0)}$ is the initial dot location.

Figure 3 shows some example dot-movement scenarios extracted from the processing of actual halftone images.

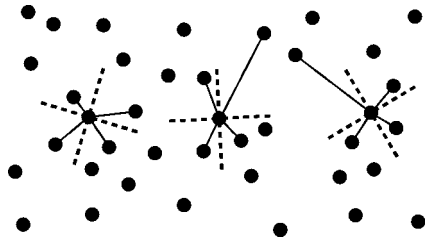


Fig. 4 One neighboring dot is found in each sector. Boundaries of sectors are oriented randomly.

Each contour represents a set of points in the plane where the cost functional has a fixed value. Note how the form of the cost functional depends upon the locations of the neighboring dots, which are inside circles. We have viewed several example dot-movement scenarios. For the most part, the cost functionals are approximately convex, with minor nonconvexities occuring around the locations of the neighboring dots, as shown in Figs. 3(a) and 3(b). But generally, the cost functionals can have considerably more irregular shapes. Moreover, they can have multiple local minima, as shown in Fig. 3(c). In any case, since the minimization is greedy, the dot being moved always ends up at the minimum of the ''watershed basin'' that contains its initial location. Of course, if the dot being moved has its initial location at a local minimum, then it is not moved at all, as shown in Fig. 3(d).

3.2 *Finding Neighboring Dots*

It is critical to choose the N neighboring dots so that the resulting dot distribution is homogeneous. For this, we choose one neighboring dot from each of N separate sectors subtending equal angles, as illustrated in Fig. 4. This has the effect of forcing dots to mingle. Otherwise, isolated groups of at least $N+1$ dots could permanently congregate in self-stable ''cliques.'' (We actually observed this in early experiments, when we were using a different neighbor-selection procedure.) To further encourage isotropic dot distributions, we orient the boundaries of the sectors randomly for each dot. This too is illustrated in Fig. 4. We tried several other methods for choosing the neighboring dots; but the procedure described here gives the best results.

Note that since we find the neighboring dots in separate sectors, there may be other neighboring dots which are close to the dot being moved, but which are not included in the cost functional. Practically, the only potentially bad consequence of this would be that the dot could end up adjacent to another dot. This would actually be a serious problem, since an adjacent pair of minority pixels in a region of extreme tone is surprisingly visible. To prevent this, we enforce the rule that the dot is not allowed to move into any location that has another dot as a neighbor.

Finally, an important parameter is the number N of neighboring dots which are chosen. We have used a value of $N=4$ to obtain our most homogeneous, isotropic dot distributions.

We found that two iterations of the dot rearrangement are sufficient for good results. More iterations will increase the smoothness and homogeneity of the dot distribution. However, this may not be worth the extra computation.

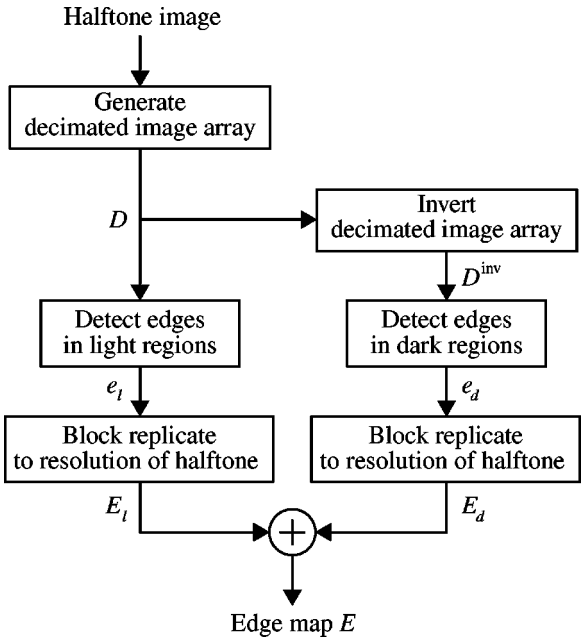


Fig. 5 Edge map generation process.

4 Generating the Edge Map

Before the dot-rearrangement stage, we generate an edge map E that assigns the value 1 to pixel locations close to edges, and the value 0 to all other pixel locations. It is used to disable postprocessing around edges, which can be rendered diffuse and soft by the dot rearrangement process. Other approaches to finding edges in binary images have been proposed.[20] However, the approach that we describe here is specifically geared toward finding edges in regions of extreme tone. An overview of our edge map generation process is illustrated in Fig. 5.

We obtain E as the logical ''or'' of two separate edge maps: one identifying edges in light regions, and the other identifying edges in dark regions. Formally, for pixel location n

$$E(n)=\begin{cases}1 & \text{if } E_l(n)=1 \text{ or } E_d(n)=1\\ 0 & \text{else}\end{cases}, \tag{5}$$

where E_l and E_d are edge maps for the light and dark regions, respectively.

The edge maps E_l and E_d are obtained from lower-resolution edge maps e_l and e_d using block replication. For this reason, the high-resolution edge map E often has a blocky appearance.

To generate e_l and e_d, we use a lower-resolution version of the input halftone called the decimated image array. By convention, we define each element in this array to be the number of black dots in an $(L\times L)$ pixel block of the halftone image, where L is the subsampling factor (see Fig. 6). This means that the decimated image array takes values in $\{0,...,L^2\}$.

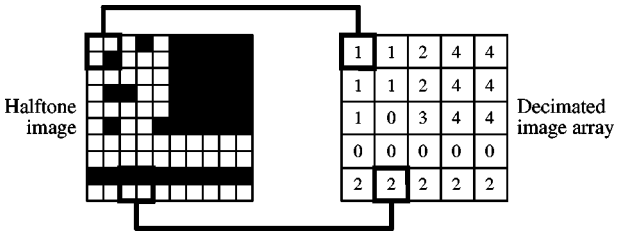


Fig. 6 Construction of the decimated image array with a subsampling factor of $L=2$.

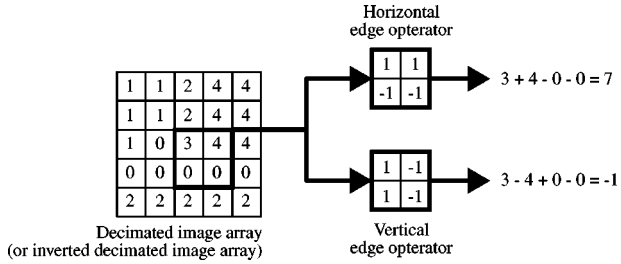


Fig. 7 Edge-detection operators.

To construct e_l, we assign a 1 to any location in the decimated image array which is part of a 2×2 block where either a horizontal or vertical edge was detected, and a 0 to all other locations. To detect edges in a block, we apply edge operators to the block, each of which yields a scalar output as illustrated in Fig. 7. The magnitudes of the scalars are compared to a locally adaptive threshold τ; and if one of them is greater than the threshold, then we say that an edge is present. The threshold τ is computed as

$$\tau(\sigma)=K_1\sigma+K_2, \tag{6}$$

where σ is the sum of the values in the 2×2 block, and K_1 and K_2 are non-negative constants. For the block outlined in bold in Fig. 7, for example, the threshold τ would take the value of $K_1(3+4+0+0)+K_2$, or $K_1\cdot7+K_2$.

To construct e_d, the procedure is essentially the same. The only difference is that instead of using the decimated image array, we use the inverted decimated image array, which is computed as

$$D^{\text{inv}}=L^2-D, \tag{7}$$

where D and D^{inv} denote the decimated image array and the inverted decimated image array, respectively.

Note from Eq. (6) how the edge-detection threshold τ depends linearly on σ. In the case of generating e_l, for example, the result of this is that the edge detection process is most sensitive to edges in the lightest regions; and it gets less sensitive as the tone decreases in lightness toward the midtones. (An analogous situation exists for the case of generating e_d.) The adaptive threshold always performs at least as well as a fixed threshold would, since that is simply the same as constraining K_1 to equal 0. In a majority of our experiments, the best result achievable by allowing both K_1 and K_2 to vary is about the same as the best result achievable by fixing K_1 to be 0. However, in some cases, our adaptive-thresholding approach can perform appreciably better.

We found that the best way to obtain values for K_1 and K_2 is to experiment with ranges of values for these variables, and then to pick the pair that yields the best result. Certainly, it would have been preferable to find a fixed pair (K_1, K_2) that is suitable for all images. But short of this, we do recommend choosing values for K_1 from [0.0, 0.6], and for K_2 from [2.0, 8.0].

An important parameter in the edge detection process is the block size L for the decimated image array. If L is too large, then the resulting edge map is too coarse. However, if L is small, then the dynamic range of the decimated image array will also be small. This can degrade the edge detection by limiting the range of values that the edge operators can possibly yield. For the results printed in this paper, we found that a value of $L=8$ yields an edge map that identifies edges adequately without being too coarse. Note that the results in this paper are printed at 100 dpi. For images rendered at higher resolutions, it may be possible to increase L without making the edge map excessively coarse.

5 Results

We tested the *Springs* algorithm by applying it to images acquired from various sources and halftoned using both error diffusion and screening. In this section, we first show the results of applying *Springs* to three different images. Then we demonstrate some shortcomings of the *Springs* algorithm. All images in this paper are reproduced at 100 dots per inch (dpi).

Figures 8, 9, and 10 show results generated using three different images. The image for Fig. 8 was scanned from a magazine and halftoned by error diffusion; the image for Fig. 9 was taken from a photo CD library and halftoned by error diffusion; and the image for Fig. 10 was created in Photoshop and screened using a 128×128 dispersed-dot dither array. Here, by ''error diffusion,'' we refer to the instantiation introduced in Ref. 2; and to design the dither array for screening, we used the algorithm described in Ref. 7. The *Springs* algorithm could equally well have been applied to halftones generated by other versions of error diffusion (see, for example, Refs. 21, 5, 22, or 4), or by screening using dither arrays generated from other methods (see, for example, Refs. 23–25, or 9).

In each figure, image (a) shows the input halftone, and image (b) shows the output of the *Springs* algorithm obtained using the parameters shown in Table 1. Note that in the light regions and in the dark regions, the minority pixels have been redistributed for a smoother, more homogeneous appearance. Image (c) in each figure shows the edge map, with black representing the regions where dot rearrangement was not allowed due to the detection of edge structures. To see why this is important, refer to image (d), which shows the result of rearranging dots while ignoring the edge segmentation. It is clear that if the edge segmentation is not used, then major edges in highlight and shadow regions—which are of critical importance when it comes to image quality—can be rendered soft and diffuse.

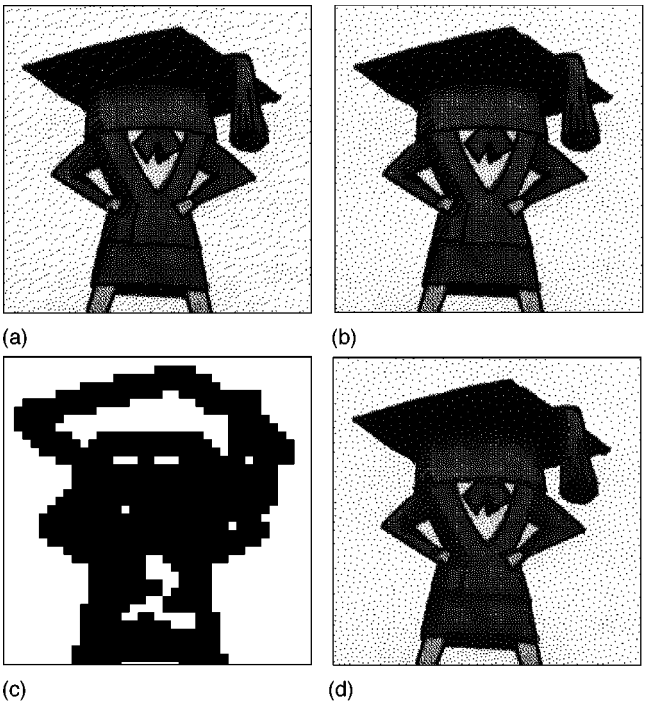


Fig. 8 Postprocessing results for a scanned image halftoned using error diffusion. (a) Input image scanned from a magazine advertisement, then halftoned using error diffusion. (b) Image (a) after postprocessing by the *Springs* algorithm. (c) Segmentation generated by *Springs* while processing image (a). (d) Image (a) after postprocessing, ignoring the segmentation.

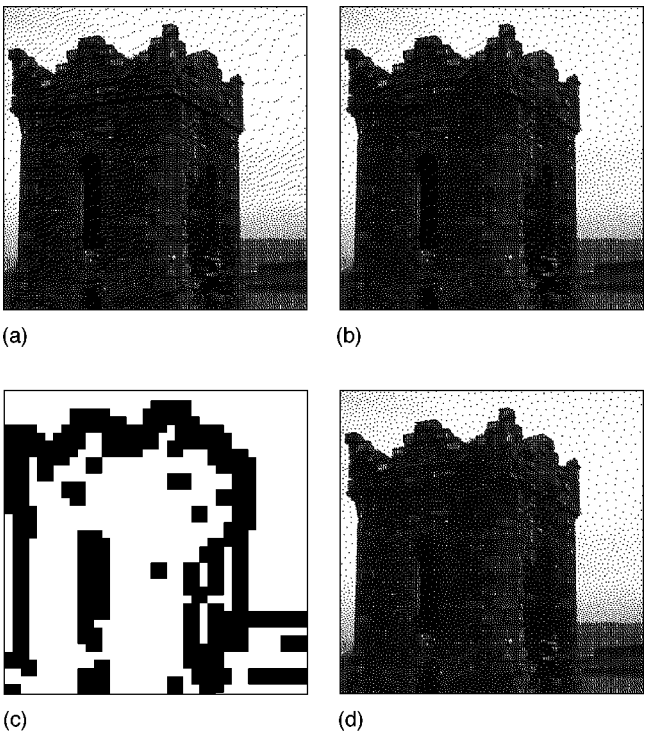


Fig. 9 Postprocessing results for an image halftoned using error diffusion. (a) Input image acquired from a photo CD library, then halftoned using error diffusion. (b) Image (a) after postprocessing by the *Springs* algorithm. (c) Segmentation generated by *Springs* while processing image (a). (d) Image (a) after postprocessing, ignoring the segmentation.

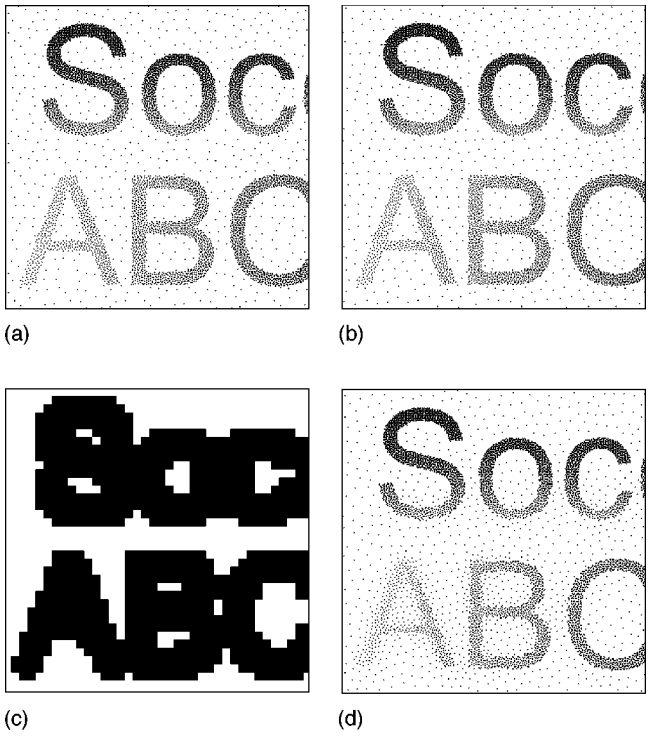


Fig. 10 Postprocessing results for a synthetic image halftoned by dispersed-dot screening. (a) Input image generated using Photoshop, then halftoned by dispersed-dot screening. (b) Image (a) after postprocessing by the *Springs* algorithm. (c) Segmentation generated by *Springs* while processing image (a). (d) Image (a) after postprocessing, ignoring the segmentation.

Since the basic action of *Springs* is to rearrange selected dots, two types of problems can occur: dots can be moved when they should not have been moved; or they can remain fixed when they should have been moved. These two scenarios are demonstrated in Fig. 11. To generate this figure, we first created a continuous-tone input image in Photoshop with the deliberate intent of ''breaking'' the *Springs* algorithm. The results of applying the *Springs* algorithm for this image appear for error diffusion on the left, and for dispersed-dot screening on the right. (Again, for error diffusion, we used Ref. 2; and for the dither array generation, we used Ref. 7.) For the edge segmentations, we used a value of $L=8$ for the block size; $K_1=0.0$, $K_2=6$ for the error diffusion result, and $K_1=0.2$, $K_2=2$ for the screening result. (These were the best sets of parameters that we could find for these images.) Note that for the halftone generated by error diffusion, some worming artifacts in the background were detected as edges and actually protected from the postprocessing. In both results, the edge between the light background and the darker region in the upper,

Table 1 Edge segmentation parameters used for Figs. 8–10.

Figure	L	K_1	K_2
8	8	0.0	8
9	8	0.2	6
10	8	0.0	8

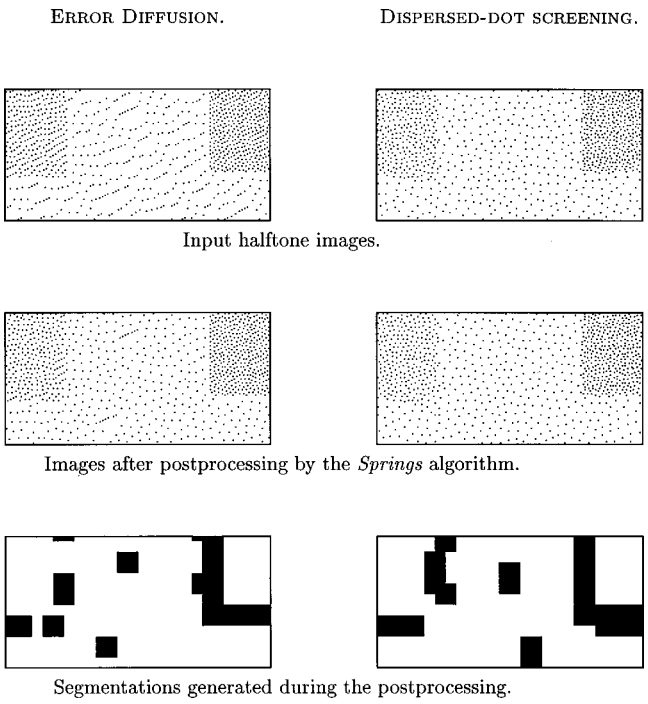


Fig. 11 Demonstration of some shortcomings of the *Springs* algorithm. The input images were generated by first using Photoshop to create the continuous-tone original, then halftoning by error diffusion (left column) and dispersed-dot screening (right column). (This image was generated with the specific objective of breaking the *Springs* algorithm.)

left-hand corner was not detected completely. Wherever that edge was not detected, the dots of the darker region appear to have ''leaked'' into the background as a result of the pixel rearrangement.

6 Conclusions

In this paper, we introduced *Springs*, an algorithm for improving halftone appearance by rearranging minority pixels in highlight and shadow regions. *Springs* moves a dot by first defining a cost functional based on the locations of neighboring dots, and then moving the dot to a local minimizer of the cost functional. This has the effect of relaxing the dot distribution. Since the dot-movement procedure has the effect of softening edges, it is only allowed in regions where edges are not detected. We demonstrated the effectiveness of the *Springs* algorithm on halftone images generated using both error diffusion and screening.

References

1. R. A. Ulichney, *Digital Halftoning* MIT Press, Cambridge, MA (1987).
2. R. Floyd and L. Steinberg, ''An adaptive algorithm for spatial greyscale,'' *Proc. SID* **17**(2), 75–77 (1976).
3. B. E. Bayer, ''An optimum method for two-level rendition of continuous-tone pictures,'' *Proc. IEEE Int'l Conf. on Comm.*, pp. 26.11–26.15 (1973).
4. R. A. Ulichney, ''Dithering with blue noise,'' *Proc. IEEE* **76**, 56–79 (1988).
5. R. Levien, ''Output dependent feedback in error diffusion halftoning,'' *Proc. of IS&T's 46th Annual Conference*, pp. 115–118, Cambridge, MA (May 1995).
6. R. Eschbach, ''Error diffusion algorithm with homogeneous response in highlight and shadow areas,'' *J. Electron. Imaging* **6**(3), 348–356 (1997).
7. R. A. Ulichney, ''The void-and-cluster method for dither array generation,'' *Proc. of SPIE/IS&T Symp. on Electronic Imaging Science and Tech.*, pp. 332–343, San Jose, CA (Feb. 1993).
8. Q. Lin, ''Improving halftone uniformity and tonal response,'' *Proc. of IS&T's Tenth Int'l Congress on Advances in Non-Impact Printing Technologies*, pp. 377–380, New Orleans, LA (Nov. 1994).
9. J. P. Allebach and Q. Lin, ''FM screen design using DBS algorithm,'' *Proc. of the 1996 IEEE Int'l Conf. on Image Proc.*, Vol. 1, pp. 549–552 (1996).
10. C. Gotsman and J. P. Allebach, ''Bounds and algorithms for dither screens,'' *Proc. of SPIE/IS&T Symp. on Electronic Imaging Science and Tech.*, Vol. 2657, San Jose, CA (Jan. 1996).
11. R. Eschbach, ''Pulse-density modulation on rastered media: Combining pulse-density modulation and error diffusion,'' *J. Opt. Soc. Am. A* **7**, 708–716 (1990).
12. P. W. Wong, ''Entropy-constrained halftoning using multipath tree coding,'' *IEEE Trans. Image Process.* **6**(11), 1567–1579 (1997).
13. R. Eschbach and R. Hauck, ''A 2-D pulse density modulation by iteration for halftoning,'' *Opt. Commun.* **62**, 300–304 (1987).
14. M. Analoui and J. P. Allebach, ''Model-based halftoning using direct binary search,'' *Proc. of SPIE/IS&T Symp. on Electronic Imaging Science and Tech.*, pp. 109–121, San Jose, CA (Feb. 1992).
15. J. B. Mulligan and J. A. J. Ahumada, in Ref. 14, pp. 96–108.
16. T. N. Pappas and D. L. Neuhoff, ''Least-squares model-based halftoning,'' in Ref. 14, pp. 165–176.
17. C. B. Atkins, J. P. Allebach, and C. A. Bouman, ''Halftone postprocessing for improved highlight rendition,'' *Proc. of the 1997 IEEE Int'l Conf. on Image Proc.*, Vol. 1, pp. 791–794 (1997).
18. J. C. Handley and E. R. Dougherty, ''Model-based optimal restoration of fax images in the context of mathematical morphology,'' *J. Electron. Imaging* **3**(2), 182–189 (1994).
19. E. R. Dougherty and R. P. Loce, ''Optimal binary differencing filters: Design, logic complexity, precision analysis, and application to digital document processing,'' *J. Electron. Imaging* **5**(1), 66–86 (1996).
20. K.-C. Fan and C.-C. Han, ''Parallel mechanism for detecting contours in binary images,'' *J. Electron. Imaging* **3**(1), 30–36 (1994).
21. B. Kolpatzik and C. A. Bouman, ''Optimized error diffusion for image display,'' *J. Electron. Imaging* **1**(3), 277–292 (1992).
22. P. W. Wong, ''Adaptive error diffusion and its application in multiresolution rendering,'' *IEEE Trans. Image Process.* **5**, 1184–1196 (1996).
23. T. Mitsa and K. J. Parker, ''Digital halftoning technique using a blue noise mask,'' *J. Opt. Soc. Am. A* **9**, 1920–1929 (1992).
24. J. R. Sullivan and L. A. Ray, ''Digital halftoning with correlated minimum visual modulation patterns,'' U.S. Patent No. 5,214,517 (1993).
25. R. A. Ulichney, ''Filter design for void-and-cluster dither arrays,'' *SID 94 digest of technical papers*, pp. 809–812, San Jose, CA (June 1994).

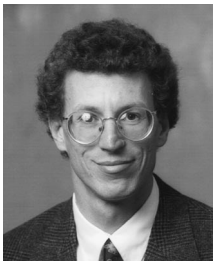


Clayton Brian Atkins received his BSEE, MSEE, and PhD degrees from Purdue University in 1992, 1994, and 1998, respectively. Since 1998, he has been with Hewlett-Packard Laboratories in Palo Alto, California. His research interests include image scaling and interpolation; image enhancement and restoration; and image halftoning and quantization.



Jan P. Allebach received his BS degree from the University of Delaware in 1972, and his MS and PhD degrees from Princeton University in 1975 and 1976, respectively, all in electrical engineering. From 1972 to 1975, he was a National Science Foundation Graduate Fellow. He was with the Department of Electrical Engineering at the University of Delaware from 1976 to 1983. Since 1983, he has been with the School of Electrical and Computer Engineering at Purdue University, where he holds the rank of Professor. Professor Allebach's research has primarily been in the areas of printing and display of images, scanning and sampling of multidimensional signals, and synthesis of digital diffractive elements. He has been a visiting summer faculty member at IBM Watson Research Center in Yorktown Heights, New York, Sandia Labs in Livermore, California, and Hewlett-Packard Labs in Palo Alto, California. In addition, he has consulted extensively for industry and

government laboratories. Professor Allebach is a Fellow of the IEEE, a Fellow of the Society for Imaging Science and Technology (IS&T), and a member of the Optical Society of America. In 1987, he received the Senior Award from the IEEE Signal Processing Society for his paper on time-sequential sampling published in the February 1984 issue of the IEEE Transactions on Signal Processing. During 1991–1992, he served as the IS&T Visiting Lecturer. During 1994–1995, he was a Distinguished Lecturer for the IEEE Signal Processing Society. Professor Allebach received an Outstanding Engineering Alumni Award from the University of Delaware in 1996. He has also received two teaching awards at Purdue University. Professor Allebach served on the Image and Multidimensional Signal Processing (IMDSP) Technical Committee of the IEEE Signal Processing Society from 1986 to 1995, and chaired it from 1991 to 1992. Professor Allebach served as Secretary of the IEEE Signal Processing Society from 1992 to 1995. Presently, he is serving as an elected member of the IEEE Signal Processing Society Board of Governors. He served as Chapters Vice-President of IS&T from 1992 to 1996. Professor Allebach is a past Associate Editor for the IEEE Transactions on Signal Processing. During 1992–1994, he served on the Editorial Board for the Journal of Electronic Imaging. Professor Allebach was Co-Technical Chair of ICASSP-93 held in Minneapolis, Minnesota, and Co-General Chair of the 9th IMDSP Workshop held in Belize City, Belize, in March 1996.



Charles A. Bouman received his BSEE degree from the University of Pennsylvania in 1981 and his MS degree from the University of California at Berkeley in 1982. From 1982 to 1985, he was a full staff member at MIT Lincoln Laboratory, and in 1989 he received his PhD in electrical engineering from Princeton University under the support of an IBM graduate fellowship. In 1989, he joined the faculty of Purdue University where he is an Associate Professor in the School of Electrical and Computer Engineering. Professor Bouman's research focuses on the use of statistical image models, multiscale techniques, and fast algorithms in applications such as multiscale image segmentation, fast image search and browsing, and tomographic image reconstruction. Professor Bouman is a senior member of the IEEE and a member of the SPIE and IS&T professional societies. He has been an associate editor for the IEEE Transactions on Image Processing, and is currently an associate editor for the IEEE Transactions on Pattern Analysis and Machine Intelligence. He was a member of the ICIP 1998 organizing committee, and is currently a chair for the SPIE conference on Storage and Retrieval for Image and Video Databases conference and a member of the IEEE Image and Multidimensional Signal Processing Technical Committee.