Digital Halftoning using Pre-Computed Maps

Fredrik Nilsson

Image Processing Laboratory, Department of Electrical Engineering, Linköping University, Linköping, Sweden

Digital halftoning techniques using threshold matrices have clear advantages in their speed of operation together with the possibility of controlling the produced dot patterns. However, due to an inherent restriction discussed in this article, the quality of dispersed dots halftones produced with this technique is not always satisfactory. In this article we describe a technique that overcomes this restriction, thereby allowing for an individual design of each tint value without any loss of speed. We also propose a design strategy for obtaining near optimal tints without losing the necessary correlation between tints. High quality reproduction of both tints and transitions between different intensity levels is thus possible. This technique also allows for new types of dispersed dot halftones with predefined micro structures. Halftones with micro structures are less sensitive to both mechanical and optical dot gain and can advantageously be used in processes with severe dot gain or low printing precision. In this article we present two types of micro structures together with several image examples.

Journal of Imaging Science and Technology 43: 153–159 (1999)

Introduction and Background

Due to the limitations of most printing devices, the reproduction of a continous-tone image requires that the image is converted to a bi-level image before it can be printed. This process, known as halftoning, can be divided into two main categories: clustered dot and dispersed dot halftoning. Extensive review of methods for digital halftoning including both categories can be found in the literature.[1] The use of dispersed dot halftones offers several advantages compared with clustered dot halftones where superior detail reproduction is one of the most important benefits. However, different strategies for the placement of the micro dots that compose the binary image may result in halftones with very different properties. The problem of placing the micro dots to obtain a binary image that gives an impression of a continous-tone image is a quite complex problem without an unambiguous optimal solution. Suggested strategies are often trade-offs between the image quality and the time required to produce the halftone. In recent years several methods for dispersed dot halftoning, all with their specific pros and cons, have been proposed.

The best results are probably obtained with iterative methods.[2,3] The general idea of these methods is to minimize an error function based on the difference between the binary image and the continous-tone image by iteratively placing or replacing the micro dots. Due to the massive amount of computation involved, these methods are not yet commercially interesting. Another strategy is the error diffusion technique that originally was proposed by Floyd and Steinberg in 1976.[4] Several variations and extensions based on this method have been suggested since then.[5–7] This strategy generally produces halftones of high quality. In some situations, however, disturbing worm-like dot patterns or sudden pattern changes might be introduced. These problems are particularly noticeable in very bright or dark tints and around the mid-tones. Error diffusion methods are in general faster than iterative methods but cannot compete speedwise with methods based on the use of a threshold matrix.[1,8]

Threshold-matrix-based methods are very attractive in terms of speed and can be implemented in a truly parallel manner, thus boosting their performance even further. Another advantage is that full control over the dot patterns produced in tints is possible. Certain less visually pleasant patterns can thus be avoided. For a good review on methods to generate threshold matrices see Ref. 9. With this approach, however, there is always a trade-off between the quality in the tints produced and the quality of smoothly varying shades. The principle of the threshold matrix approach does not allow for the highest possible quality in both. The reason for this is that the dot pattern produced for a certain tone value will be present in all tints of less intensity. A pattern at one intensity level will therefore restrict the appearance of all other tint values. This means that if there is a certain criterion for optimization of a tint for instance, requiring that the dots should be as dispersed as possible, this can not be fulfilled for all tint levels. This restriction is illustrated in Fig. 1.

We have developed a method for halftoning that overcomes the constraint of threshold matrix methods without any loss of speed. All tints produced have been individually designed and can thus have any desired property. The tints will therefore generally be less grainy than tints produced with a threshold matrix method and hence visually pleasant for both bright and dark tint values. Furthermore, the proposed design strategy allows for smooth transitions in halftoning of shades without any sudden structural changes. High quality in both tints and smoothly varying shades is therefore possible.

The fundamentals of the method have been described previously.[10] Improvements have been made since then, and new types of halftone characteristics have been de-

Original manuscript received May 29, 1998

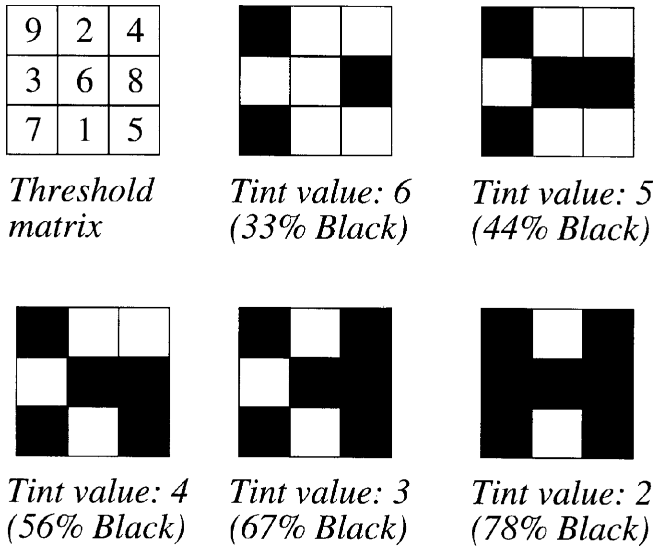

Figure 1. An illustration of the limitations of threshold-matrix-based halftoning techniques. A black micro dot is placed at every image position where the value of the threshold matrix is greater than the image value. Therefore, the dots produced for one tint value will be present in all darker tint values. This is a severe restriction when designing the tints possible to produce with the matrix.

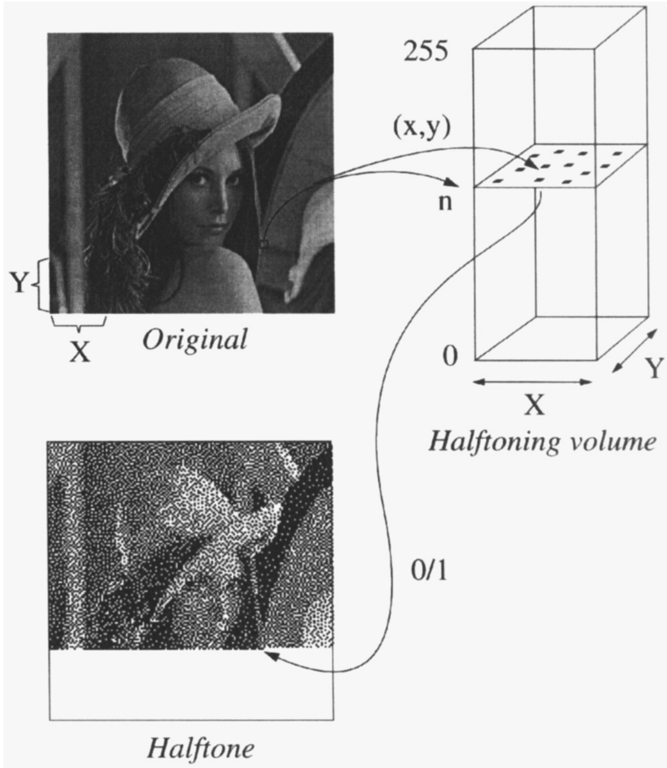

Figure 2. Halftoning with the Pre-Computed Maps technique described in the text. Maps describing the dot patterns for every possible image value (256 in the illustration) are stored in a halftoning volume. When halftoning, the image value is used as an index in the volume to point out the appropriate map and the position of it to compute the dot position within the map. The out pointed value of the map is then copied into the halftone at the given image position.

veloped. The ideas in the following sections are patent pending. Recently, other authors have also suggested the use of a collection of correlated maps instead of a threshold matrix for halftoning.[11]

The Pre-Computed Maps Technique

Instead of using a threshold matrix and a comparison technique to produce the halftones, the proposed method uses one pre-computed halftoning map for each possible tone value in the image. Each map is derived according to a specific optimization criterion and stored in a *halftoning volume*. The maps contain only zeros and ones, representing the presence and absence of a halftone dot respectively. When halftoning an image, the image value is used as an index into the volume. The position within the chosen map is derived from the image position in the same manner as for matrix-based-methods. The value of the map at this position is then simply copied into the corresponding position in the resulting halftone. Thus, no comparisons are necessary when halftoning. Halftoning with the Pre-Computed Maps (PreCoM) technique is illustrated in Fig. 2.

Like threshold-based-methods, PreCoM is easily implemented in a truly parallel manner because no information from neighboring pixels is required. However, the technique requires more memory than with threshold matrix methods. Given the task of halftoning an image with 256 gray levels (8 bits), the PreCoM method requires 256 one-bit maps whereas a threshold-matrix-based method requires 1 eight-bit matrix. If the same matrix map size is used respectively, the PreCoM method requires 32 times as much memory. As an example, assuming a map size of 64×64 pixels, the storage of the halftoning volume will require 128 kilobytes compared with the 4 kilobytes for the threshold matrix.

While PreCoM allows full control of the dot patterns for each tone value, severe discontinuity effects will be introduced when halftoning slowly varying shades if the maps in the volume are computed in isolation. Thus, adjacent maps in the volume must be correlated, i.e. exhibit similar dot patterns. The design of the maps must be done with this in mind. There is, however, no reason why the dot pattern of one intensity level should be restricted by the dot pattern at another level far from it. This implies that this limitation will not be as severe as for threshold-matrix-based methods. The computation of the maps and the process of correlating adjacent maps is described in the following sections.

Computing the Maps

The maps in the halftoning volume are derived individually in an iterative manner. The principles of the computation process are shared by all maps. However, by assigning a specific optimization criterion to each map, all maps can be given a specific characteristic. A variety of optimization criteria are possible, but for the production of visually pleasant halftones only some are of interest. One such criterion is to force the dots in a map to be as dispersed as possible so as to minimize disturbing patterns in the halftoned images. In fact, maps optimized with this criterion will possess the desirable blue noise characteristic.[1] The suggested optimization algorithm is closely related to the void-and-cluster method proposed by Ulichney.[12] The principle of the algorithm has also been used by Yao and Parker when deriving their threshold-matrix-based method the Blue Noise Mask.[13]

The computation of a map starts from any binary pattern with the appropriate ratio of white to black dots (ones and zeros), for instance, a thresholded white noise image. The idea of the algorithm is to rearrange the minority dots, i.e., black dots for bright tint values and white ones for dark values, until they are as dispersed as possible. This is done in a iterative manner by successively moving minority dots from the tightest clusters to the largest voids.

By representing the minority dots with ones and the majority dots with zeros, the tightest cluster of minority dots is located by finding the maximum of the low-pass filtered dot pattern. The minority dot at the maximum is removed and the low-pass image is updated by a local deconvolution, i.e., the low-pass kernel is subtracted from the low-pass image at that position. In a similar manner, the largest void is found by locating the minimum of the low-pass image. The previously removed dot is then placed at this location and the low-pass image is again updated, now by adding the kernel. The process of rearranging minority dots continues until no further change occurs, i.e. the selected minority dot is placed at the very same location it was taken from. The algorithm is straightforward but there are several factors that must be taken into consideration in order to ensure that the resulting maps behave well. The two most important of these factors are discussed next.

Because the image size is, in general, much larger than the map size, the maps have to be tiled for complete coverage when halftoning the image. If this is not taken into account when computing the maps, disturbing artifacts will be introduced at the borders of the maps. A straightforward solution is to use circular convolution when low-pass filtering the dot pattern. This means that the influence from one dot is allowed to spread over the map edge into the opposite side of the map. Thus, each generated map can be tiled without introducing any discontinuity effects at the borders of the map.

The characteristics of the low-pass kernel used in the algorithm has great influence on the final dot pattern. Its shape as well as its size is of importance and should be adapted to the current density level. For instance, in bright and dark maps, the distances between the minority dots will be large. Thus, a wide kernel should be used to avoid unevenly spread dots. On the other hand, for maps of around 50% coverage, dots have to be placed next to each other. A kernel designed to control the local behavior of the dot pattern is therefore needed.

The iterative process described above is fairly time consuming. The time taken depends on the number of minority dots, the initial binary pattern, and the size of the kernel used for the optimization. If, however, the initial binary pattern is close to optimal, the computational cost is dramatically decreased. This is utilized in the correlation process described next. Also, the computation is done off line and once and for all. The time required for generating the halftoning volume is therefore of minor interest.

Creating the Locally Correlated Volume

The strength of the PreCoM technique lies in the possibility of producing both near optimal tints and pleasant transitions between them. Without the correlation between adjacent maps, disturbing irregularities and discontinuities will be detectable at the transitions. The proposed method for the generation of the halftoning volume assigns an unique filter kernel to each map that is to be optimized. By changing the filter kernel gradually between adjacent maps, and by using the previously optimized map as the starting pattern for the next map, maps that produce images with both desired properties can be derived.

Before generating the halftoning volume, decisions have to be made about the map size and the number of tint levels that should be defined. Any size of the map could be chosen, but all maps in one volume must be of the same size. Furthermore, if the maps are too small, a disturbing periodical structure might be introduced into the produced halftones. Through experiments we have discovered that a map size of 64×64 positions seems to be sufficient for the elimination of such effects. Still, larger maps seems to give even better results. The image examples in this paper have been derived with a map size of 128×128. While it is possible that an optimal size could be established and proved, such experiments have not yet been done.

The image material we have been working with has a dynamic range of 8 bits, i.e., 256 gray levels are possible. It is doubtful whether more levels are meaningful for a human observer, at least for images reproduced in print. Possibly, it may be the case that even fewer levels could be acceptable. Until investigated further, however, the volumes will contain 256 pre-computed maps, one for each possible intensity level in an 8 bit image.

Having decided on these two properties, the number of minority dots that differs between two adjacent maps can be computed. A fixed difference in dots between maps results in a linear reflectance function for the digital volume. Unfortunately, due to the dot gain in prints, this linearity does not carry over to the reflectance function for the print. However, if the dot gain function is known, the structure of the PreCoM technique allows dot gain compensation to be performed when the volume is computed. Because each map is designed individually, each map can contain the number of dots that produces the desired linear reflectance function in print. Dot gain compensation can thus be done at the same time as the image is halftoned, making pre-adjustments of the original image's histogram unnecessary. With the number of dots for each map decided, the halftoning volume can be calculated. The proposed procedure for this is described below. Examples of results during the process are shown in Fig. 3.

The process is initiated by generating a matrix of the desired map size containing white noise in the range [0,1]. By thresholding the matrix at 0.2 for instance, a dot pattern with approximately 20% minority dots is obtained (Fig. 3a). If necessary, extra dots can be added or removed at random to get the exact number of dots that the map should have. By using the assigned filter kernel and applying the iterative optimization process described above, the first map of the volume is computed and stored (Fig. 3b). To calculate the next darker map, the voids of the first map are located and the desired number of extra dots are added iteratively in these positions by using the filter assigned to the first map. Note that no rearrangements of the dots are made at this stage (Fig. 3c). While the new dot pattern may be optimal with respect to the old filter, it will in general not be optimal with respect to the new filter that is assigned to this particular tint value. Thus, the dot pattern has to be optimized with the new filter before the new map can be stored (Fig. 3d). This procedure of iteratively deriving the next map from the previous map continues up to the level where the map consists of equal ones and zeros. At this stage the representation of minority and majority dots is switched, and instead of adding dots, dots at the tightest clusters are removed before optimization. The procedure continues until all minor-

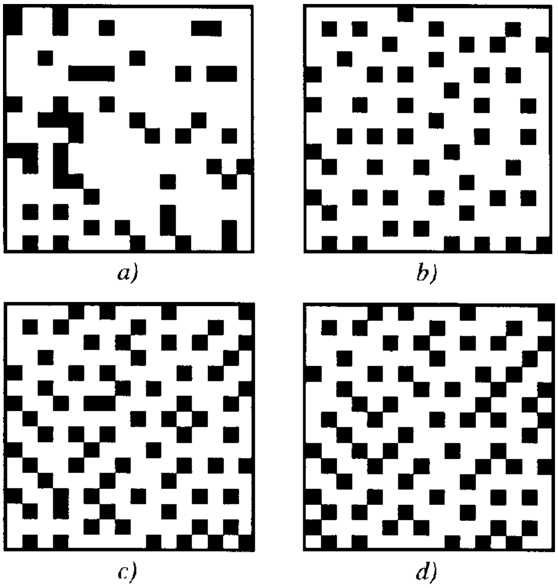

Figure 3. Examples from the generation of maps, (a) is a white noise image thresholded to get 20% minority pixels. The resulting pattern is used as the starting pattern. (b) The starting pattern is rearranged to get maximally dispersed dots. This map is stored in the halftoning volume. (c) Extra dots have been added to the dot pattern in (b) to get the initial map for the next level (25% level in the example). (d) The pattern is rearranged before being stored in the volume. Note that some of the dots in (b) have been moved to obtain the pattern in (d). This would not have been possible with a threshold-based-method.

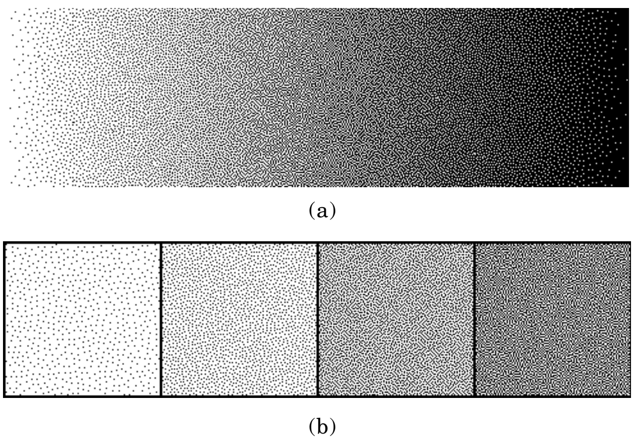

Figure 4. A synthetic test image consisting of a grayscale ramp and a collection of tints halftoned with the PreCoM method described in text. The halftoning volume used is optimized with isometric Gaussian low-pass filters. The maps are of the order 128 × 128.

ity pixels are gone. The maps between 0% and 20% are derived in a similar manner, starting from the first generated map at 20% and removing minority dots.

The filter kernel used to make the dots as dispersed as possible in each map is a Gaussian shaped kernel. The width of the kernel function is changed according to the tint value. The fewer the minority dots, the wider the kernel. The filter kernel $g(x,y)$ is defined by

$$g(x,y)=\begin{cases} e^{-\frac{(x^2+y^2)}{(c_1-c_2T)}} & (|x|,|y|)\le S/2 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Figure 5. A grayscale photograph halftoned with the PreCoM method. The halftoning volume used is optimized with isometric Gaussian low-pass filters. The maps are of the order 128 × 128.

where T is the number of minority pixel divided with the map size, c_1 and c_2 are positive constants chosen to make the expression c_1–c_2T positive, and where S is the size of the kernel in pixels. For an efficient implementation of the algorithm, the size of the kernel, S, could be made dependent on the function $g(x,y)$. While a wide kernel size is necessary for very bright or dark tints, a small kernel size is sufficient for tints around 50% density level. Therefore it is not necessary to use the large kernel size for all maps.

Experiments have shown that to get the smooth transition between tints around the 50% level, a change in filter characteristics is needed. Therefore, for maps around this level, a small amount of band pass characteristics is introduced. This is further discussed in the next section. In Fig. 4, a synthetic test image consisting of a grayscale ramp and a collection of tints has been halftoned with the PreCoM technique using a halftoning volume computed according to the above described procedure. The density levels of the four tints are 5%, 15%, 35%, and 50% respectively. In Fig. 5, a scanned photograph of a windmill has been halftoned using the same halftoning volume.

There are, however, some known problems with the described method. In rare cases, a dot is moved from its position in one map to a neighboring position in an adjacent map. This could result in a micro cluster in the halftoned image in transitions between these two maps. By examining the halftones in Fig. 4 and Fig. 5 thoroughly, both small white and black clusters can be found. It is doubtful if such micro cluster actually lowers the overall visible quality, but if this is the case, it should be possible to adjust such artifacts with a post-processing operation applied to the halftoned image. One could, for instance, compare the local average values of the original and the halftone and add or remove dots in those areas where the difference is too large. However, any post-processing operation will require additional time

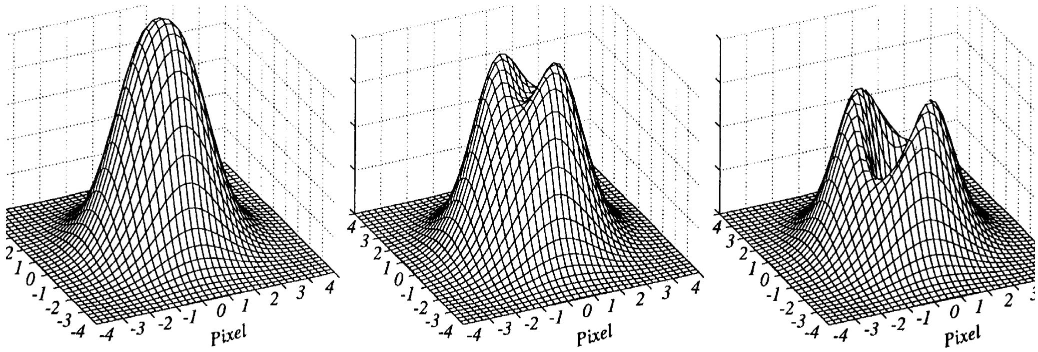

Figure 6. Some of the filter kernels used to compute maps with line-like micro structures. The kernels gradually change from a Gaussian low-pass filter for maps with fewer than 30% minority dots to a directional band-pass filter.

that in many situations is more of a problem than the possible decrease of image quality caused by the micro clusters.

Changing the Filter Characteristic

It is not obvious that the optimization criterion that results in maximally dispersed dots produces the most visually pleasant halftones. For instance, in maps with above 25% of minority dots, some of the dots must necessarily be placed in connection with each other, at least corner to corner. Moreover, for maps around the 50% density level, clusters of dots will always be generated. Due to the correlation criterion, these may not always be visually pleasant. One may therefore be motivated to try other optimization criteria for maps with more than 25% of minority dots. For instance, instead of spreading the dots as much as possible, they could be arranged into supervised micro structures.

The use of micro structures may both facilitate the correlation between maps, thus producing smooth transitions, as well as increase the visual quality of the resulting tints. The presence of structures appears to give a less grainy impression of the halftone. As an example, the visually pleasant impression of results from iterative and error diffusion methods could partly be explained by the micro structures introduced. The observer seems to exclude even clearly visible structures and focus on the image information it conveys. In addition, because the influence from the dot gain is greater for several dispersed dots than for one large dot, the dot gain will be smaller with micro structures than with the maximally dispersed dots. This could prove useful when using print setups with heavy dot gain. Moreover, micro structures could also be useful to facilitate halftone reproduction in setups with low printing precision.

In theory, the PreCoM technique allows for any micro structure to be used in the halftones. We have found through these experiments, two types of structures that result in less grainy halftones as well as an increase of the overall image quality.

Line-like Structures

Inspired by the conventional, but not commonly used line screens that result in fairly clear images, we have introduced line-like micro structures for the mid-tones in the halftoning volume. This is realized by using directional band-pass filters in the optimization of these maps. Still, in bright and dark tints, the dots can be, and presumably should be maximally dispersed. For these maps, a Gaussian low-pass filter is used. This means that to maintain the correlation between adjacent maps, the band-pass characteristics has to be introduced gradually. Up to a certain level, for instance 30% minority dots, the same Gaussian filter as discussed previously can be used. Passing this level, the Gaussian filter is modified by subtracting a function describing a directional low-pass filter. Choosing the bandwidths properly, the result from this operation will be a directional band-pass filter. By gradually increasing the influence from the subtractive term, the filter will become more and more band-pass-like and thus suitable for the higher levels. Examples of such kernels, with the gradual change, are shown in Fig. 6.

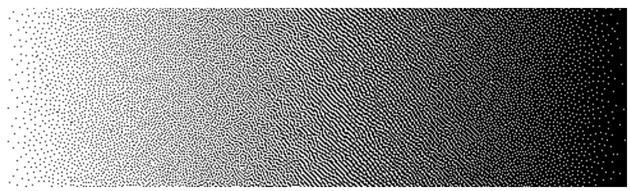

(a)

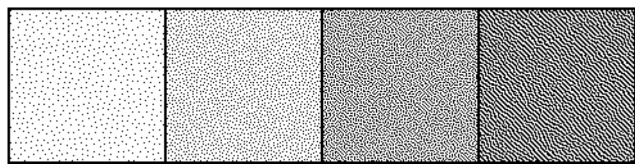

(b)

Figure 7. The synthetic test image halftoned with the PreCoM method. The halftoning volume used is optimized with directed Gaussian band-pass filters for mid tone maps. The maps are of the order 128×128.

Figure 7 shows the synthetic image halftoned using the halftoning volume derived with these filters. A slight asymmetry can be noticed in the ramp, where the line structure is visible further down on the dark side of the 50% level than on the bright side. The reason for this is that the initial binary pattern where the computation of the volume started was derived at 20% black dots. Because previously derived maps will affect the appearance of the next, the maps with more than 50% black dots will be affected by the heavy line structures in the 50% level. While it may be possible to derive a more

Figure 8. The grayscale photograph halftoned with the PreCoM method. The halftoning volume used is optimized with directed Gaussian band-pass filters for mid tone maps. The maps are of the order 128×128.

symmetrical volume, experiments have shown that if the initial map is derived at the 50% level, unpleasant discontinuity effects are introduced around this level. The reason for this may be that the optimization results in a local minimum from which it is hard to derive the next map without major changes.

As an example, an optimal map for the 50% level could be a perfect chessboard pattern, but the next map derived from this will either be visually unpleasant or have a very different characteristic. The starting map should therefore be derived at a level where a stable pattern is the optimum, i.e., a pattern where dots easily can be added to or removed from without really affecting the optimality or the characteristics.

However, the main reason why we have chosen to let the micro structures go further down in the dark levels than in the bright, and not put much effort in generating a symmetric halftoning volume, is due to the properties of the dot gain. The effect of the dot gain is usually more severe for dark tone values than for bright ones. Because micro structures are less sensitive to dot gain, the introduction of micro structures can resolve some of the dot gain related problems in those tone values.

As can be seen from the halftoned test image in Fig. 7, there are no obvious discontinuities in the shades due to the gradual changes in structure. The tints are slightly less grainy than without the micro structures. This is also true for several parts, as shown in Fig. 8, in particular, the background. A drawback with the use of micro structures is a slight impairment in the reproduction of details. This could, however, be compensated for by an increase in print resolution that is facilitated by the more robust and less dot-gain-sensitive halftones.

Curved Structures

For low print resolutions, the line structure may become too obvious and details along the orientation of the lines will be less sharp than others. A remedy to these problems is to use *curved* micro structures instead of straight lines. Such structures can be obtained by using *isotropic* band-pass-like filters in the optimization process. The computation of the halftoning volume is done in the same manner as discussed previously. The curved structure is gradually introduced by changing the properties of the filters. The closer to the 50% level, the more band-pass characteristic will the filters get. Examples of such filters are shown in Fig. 9.

In Fig. 10, the test image has been halftoned with the PreCoM technique using a volume derived with isotropic band-pass filters. The transitions in the shades are even smoother than for the line-like micro structure and the structures themselves are less obvious. Because there is no preferred orientation in the micro structure, the reproduction of details does not depend on their orientation. Moreover, even less graininess can be detected in both the tints as well as in the homogenous areas of Fig. 11. For more image examples and comparisons between the PreCoM methods and other dispersed dot halftoning methods, see Ref. 14.

Conclusions

Halftoning with PreCoM introduces the possibility of producing well-behaved tints as well as smooth transitions between them without any loss in speed compared to methods based on threshold matrices. The method can be implemented in a truly parallel manner, speeding up the halftoning procedure even further. Moreover, because every map is designed individually, under the constraint that the correlation is not lost, specific characteristics that improve the image quality can be introduced. This, together with the possibility of compensating for the dot

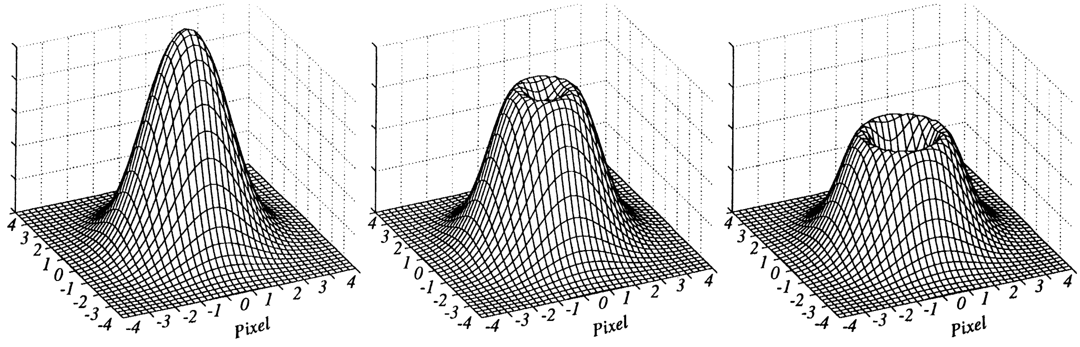

30% Minority dots *40% Minority dots* *50% Minority dots*

Figure 9. Some of the filter kernels used to compute maps with curved micro structures. The kernels gradually change from a Gaussian low-pass filter for maps with fewer than 30% minority dots to an isotropic band-pass filter.

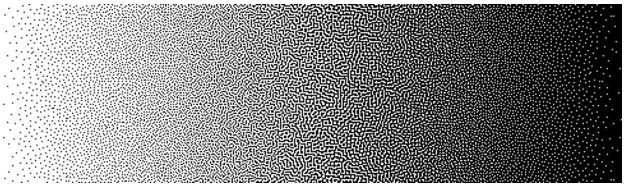

(a)

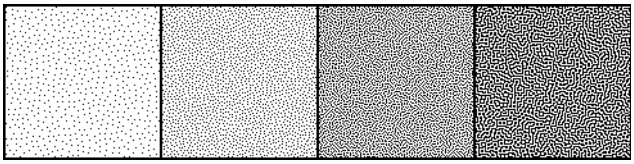

(b)

Figure 10. The synthetic test image halftoned with the PreCoM method. The halftoning volume used is optimized with isotropic Gaussian band-pass filters for mid tone maps. The maps are of the order 128×128.

gain in the halftoning procedure itself, makes this method attractive in applications where both speed and quality are important. In this article, we have proposed two kinds of micro structures, curved and line-like structures, that we believe increase the image quality. The use of micro structures facilitates the reproduction of tints and shades and gives a less grainy impression than maximally-dispersed-dot halftones. Moreover, halftones with micro structures are less sensitive to the effects of dot gain and to problems related to printing precision.

Figure 11. The grayscale photograph halftoned with the PreCoM method. The halftoning volume used is optimized with isotropic Gaussian band-pass filters for mid tone maps. The maps are of the order 128×128.

References

1. R. A. Ulichney, *Digital Halftoning*, The MIT Press, Cambridge, Massachusetts, 1987.
2. M. Analoui and J. P. Allebach, Model-Based Halftoning Using Direct Binary Search, *Proc. SPIE, Human Vision, Visual Processing, and Digital Display III*, **1666,** 96–108 (1992).
3. T. N. Pappas and D. L. Neuhoff, Least-Squares Model-Based Halftoning, *Proc. SPIE, Human Vision, Visual Processing, and Digital Display III*, **1666**, 165–176 (1992).
4. R. W. Floyd and L. Steinberg, An adaptive algorithm for spatial grey scale, *Proc. Soc. Inf. Displ.* **17,** 75–77 (1976).
5. J. F. Jarvis, C. N. Judice and W. H. Ninke, A survey of techniques for the display of continous-tone pictures on bilevel displays, *Comput. Graphics and Image Processing*, **5,** 13–40 (1976).
6. R. Eschbach, Error Diffusion Algorithm with Reduced Artifacts, *Proc. IS&T's 45th Annual Conf.*, IS&T, Springfield, VA,1992, pp. 133–136.
7. R. Levien, Output Dependent Feedback in Error Diffusion Halftoning, *Proc. IS&T's 46th Annual Conf.*, IS&T, Springfield, VA, 1993, pp. 115–118.
8. B. E. Bayer, An optimum method for two level rendition of continous-tone pictures, *Proc. IEEE Int. Conf. Commun.*, 26.11–26.15, (1973).
9. K. E. Spaulding, R. L. Miller and J. Schildkraut, Methods for generating blue-noise dither matrices for digital halftoning, *J. Electron. Imag.* 208–230, (1997).
10. F. Nilsson, Precomputed Frequency Modulated Halftoning Maps that Meets the Continuity Criterion, *Proc. IS&T Int. Conf. on Digital Printing Technologies (NIP12)*, IS&T, Springfield, VA, 1996, pp. 72–76.
11. P. Li and J. P. Allebach, Look-Up-Table Based Halftoning Algorithm, *Proc. IEEE Int. Conf. Image Processing*, 115–118 (1998).
12. R. Ulichney, The void-and-cluster method for dither array generation, *Proc. SPIE, Human Vision, Visual Processing, and Digital Display IV*, **1913**, 332–343 (1993).
13. M. Yao and K. J. Parker, A Comparison of the Blue Noise Mask and Error Diffusion, *Proc. IS&T's 47th Annual Conf.*, IS&T, Springfield, VA, 1994, pp. 491-493
14. F. Nilsson, Halftoning and Objective Quality Measures for Halftoned Images, Licentiate thesis, Linköping Studies in Science and Technology, Thesis no. 671, Linköping, 1998.