Precomputed Frequency Modulated Halftoning Maps that Meets the Continuity Criterion

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

Abstract

Frequency modulated halftoning techniques using precomputed threshold matrices have a clear advantage in their speed of operation. Iterative techniques generally show more visually pleasing results but have very slow operation. This paper describes a technique that bridges the gap between the two techniques. It is based on a precomputed set of optimized halftoning maps with a semicontinuous transition in between. This results in a smooth halftoning of slowly varying tones without the graininess that threshold matrices usually exhibit. The technique is slightly more demanding in terms of required memory space but it compares favourably to simple threshold techniques in terms of speed.

Introduction

In recent years several techniques has been developed for frequency modulated halftoning. Unfortunately, the methods giving the best results are also very time demanding. Examples of such methods are the Direct Binary Search (DBS) algorithm[2] and the Hybrid Halftoning method.[2] The aim in both these methods is to minimize the perceived difference between the binary image and the original continuous-tone image. This is done in an iterative manner forcing the error metric towards a local minimum. However, in applications where the time spent on the halftoning process has to be minimized one is still forced to use techniques based on threshold matrices. Although some very competent methods have been presented in this area,[3,4] they fundamentally suffers from one significant drawback which is discussed below.

When halftoning homogenous tints, i.e. areas in the original image with constant tone value, we require that the relationship between the area of the plate and the number of halftoning dots corresponds to the tone value of the image. Further, we require that the halftoning dots are maximally spread over this area. Using a threshold matrix-based method, the second requirement can not be fulfilled for all possible tone values in the image. For instance, assume that a halftoning dot has been produced after the comparison between the image value and the threshold matrix at a certain position. For all image values greater than, or equal to this value, the algorithm will always produce a halftoning dot at this position. This means that we restrict tint plates of higher values to hold dots at certain positions. Thus, tints produced by a threshold matrix-technique will never exhibit patterns with the dots maximally dispersed.

In our method we abandon this constraint and produce a halftoning map for every possible value in the image. In each map the dots are maximally dispersed, thus producing tints with less graininess than threshold matrices-techniques. For instance, if we were to halftone images containing 256 tint values, we would precompute 256 halftoning maps consisting of ones and zeros only, and store these in a *halftoning volume*. When halftoning the image we use the image value as an index into the volume and simply copy the value of the map at the corresponding position into the resulting image. Thus, no comparison has to be done. However, if the precomputed maps are optimized one by one, not considering the other maps in the volume, we would get considerable discontinuity effects when halftoning slowly varying shades. Therefore the halftone maps has to be correlated in some way. The correlation can be done in several ways of which we will present two in this paper.

Several techniques to design tints with maximally dispersed dots have been proposed in the literature. Our method for optimizing the individual maps is strongly akin to the void-and-cluster method.[5] In this iterative method the tint is low-pass filtered and the maximum and minimum value of the result are localized. The positions where these are found corresponds to a halftoning dot and a 'hole'. To improve the quality of the tint, i.e. get more dispersed dots, the dot and the hole simply switch places. Another way to get optimized tints is to design the tints in the Fourier domain. The advantage with this technique is the instant control over the frequencies in the tint. It is for instance possible to omit the low frequencies which gives unpleasant grainy effects.

Computing the Halftoning Maps

The fundamental idea is to let every possible value in the original image be represented by an unique halftoning map and to store these maps together in a *halftoning volume*. Every map in the volume is accessible at any time through an index which is given by the tint value the map corresponds to. The size of the maps may be arbitrary, but should be the same for all maps in one volume. Each map will consist only of ones and zeros, where the ones corresponds to producing a halftoning dot in the resulting image. So for instance, assuming that a gray-scale image should be halftoned, the halftoning map for a totally white tint will consist only of zeros and a map representing a 33% black tint would consist of twice as many zeros as ones. Two criteria must be considered when constructing these maps:

- •C1 The dots should be maximally dispersed to avoid graininess.
- •C2 Maps close to each other in the volume must be correlated, i.e. approximately hold the same patterns to avoid unpleasant discontinuity effects.

In threshold matrix-techniques the criterion C2 is by definition always met, but again, C1 can not be met for all tone values. In our method we will try to meet both criteria even if they, in some sense, are in contradiction with each other.

To compute a map that meets C1, we start with a binary pattern generated from a thresholded white noise image. The ones and zeros in the result from this operation are rearranged until the ones are maximally spread, thus we produce a map with blue noise properties. The rearrangement can, as mentioned in the introduction, be done in several ways. We have chosen to make the rearrangement according to the following.

Optimization of the Halftoning Maps

The rearrangement process starts with the map, M_1, given by the threshold operation above. The following steps are iterated until no further change occur in the map which will give us a map with the desired properties.

- Convolve M_1 circularly with a Gaussian kernel, for instance:

$$g(x,y) = \begin{cases} e^{-((x-x_0)^2+(y-y_0)^2)/\alpha} & (|x|,|y|) \le \beta \\ 0 & otherwise \end{cases} \qquad (1)$$

where α and β is depending on the size of the map and the tint value to be computed. This operation produce a low frequency image of the map, LP_1.

- Overlay LP_1 with the map itself (i.e. pixel wise multiplication), and find the maximum value from the result. Its position corresponds to a badly placed 'one' (a halftoning dot) in the map. Remove this dot which gives us M_1^*.
- Make a locally deconvolution of LP_1 around this position, i.e. subtract the kernel from LP_1 at this position, giving LP_1^*.
- Overlay LP_1^* with $(1 - M_1^*)$ and find the minimum value. Its position corresponds to a badly placed zero in the map. Replace that zero with a one which gives us a new map M_2. Finally locally convolve LP_1^* with the kernel (i.e. add the kernel at this position) which gives us the low frequency image of M_2, LP_2.
- Iterate from the second step (now using the latest computed map) until no further change in the map occur, i.e. until a chosen dot is placed back at its original position.

In the operations above we move badly placed dots to a position as far away as possible from any other dot. The circular convolution gives us the desired wrap around effect which allows us to tile the maps to halftone the entire image.

Ensure Continuity

When halftoning slowly varying shades, we must make sure that the halftoned image gives an impression of continuity. To be able to do this we have to meet the second criterion, C2. We propose two methods for doing this.

In the first method (*Method 1*) we simply derive the maps from the same white noise image but with different threshold values. The threshold value is chosen to give a correct relationship between ones and zeros. This simple procedure actually will produce maps with similar patterns, at least for maps close to each other in the volume. The advantage with this method lies in the fact that all maps in the halftoning volume will be optimal, i.e. the dots in each map will be maximally dispersed. On the other hand, we do not have total control over the patterns produced which possibly could lead to an all too great difference in two adjacent maps.

In the second method (*Method 2*) we try to remedy this artifact by designing the maps at the dot level. Instead of optimizing all of the maps, as in the first method, certain maps, separated by a specific number of intermediate maps, are optimized. The optimization of these are done in the very same manner as above. The maps in between are given by interpolation which is done from the map with fewer dots, M_a, towards the next optimized map in the volume, M_b, according to the following.

Every dot in M_a is matched with the closest dot in M_b (Euclidean metric). The remaining, unmatched dots in M_b, are stored for later use. To compute the next map in the volume M_{a+1}, we let every dot in M_a move closer to its matched dot in M_b. The length of the step to make is determined by dividing the distance between the two matched dots with the number of maps between M_a and M_b. Finally, to get the required tint value, some of the previously stored dots are added to the map. The computation of the next maps is done in a similar manner but now by moving the dots in the latest computed map closer to the dots in M_b.

Since we move each dot individually we are ensured that maps close to each other in the volume holds a similar pattern, thus, the continuity criterion is met. Furthermore, the interpolation is always done between two optimized maps, if the distance between them is not too large, the maps in between should not be far from optimal. However, the interpolated maps are not perfectly optimized which could give the undesired graininess effect in some tints. To avoid this problem, thus producing better maps, we may be more careful when moving and adding the dots. For instance, we could let the size and the direction of each step vary within certain limits, allowing the dots to be positioned as far as possible from every other dot in each step. In the same manner, we may be more careful when choosing which dots to add and even to let these be moved a short distance a away from their given position, all this to produce maps where the dots are as dispersed as possible.

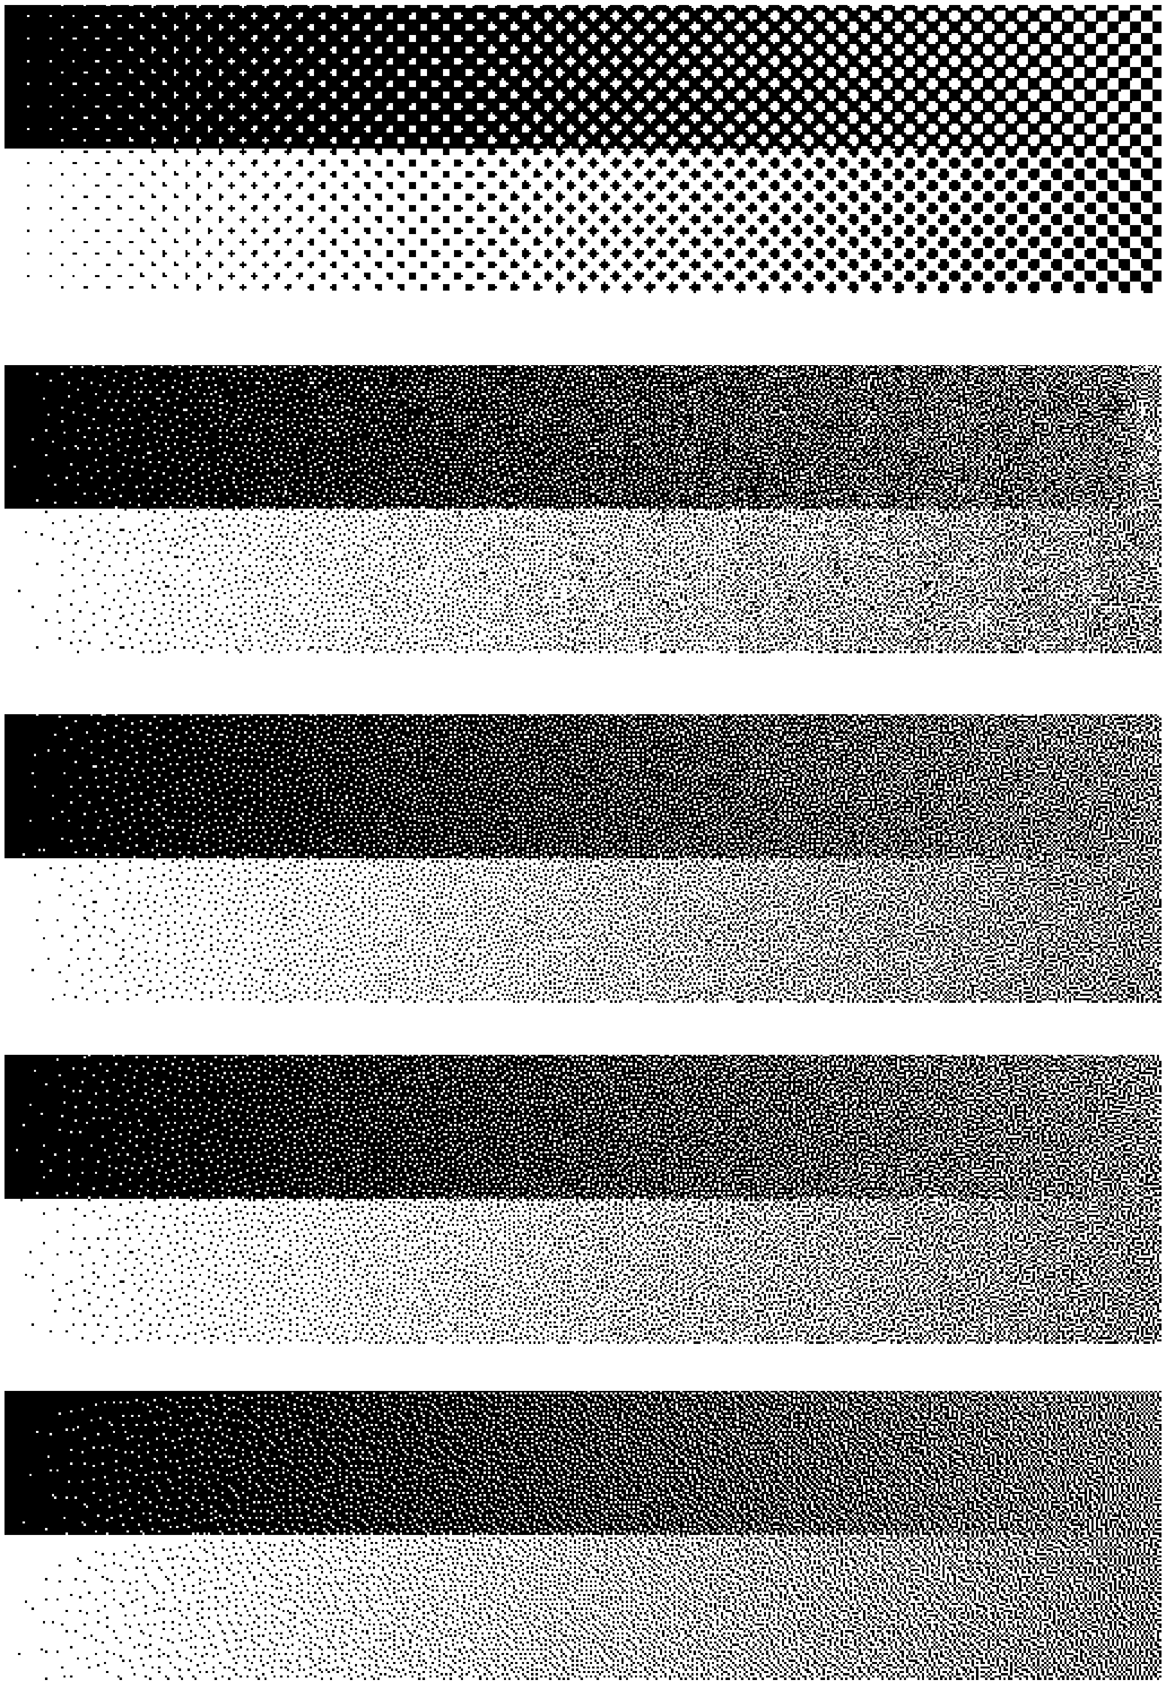

Figure 1. The original image is a 256-step gray-level wedge where every tint is 4 pixels wide. The wedge has been halftoning using: a) Conventional clustered dot halftoning, b) Optimized maps without correlation, c) Interpolation between optimized maps (Method 2, described in text) d) Optimized and correlated maps (Method 2, described in text) e) A commercial available Error Diffusion method.[6]

Results

In the following examples we have felt obliged to omit results from threshold matrices-techniques. The results we have produced with this technique are poor and it would not be relevant to compare these results with the following examples. Instead we have chosen to compare our method to a conventional clustered dot method and to a commercially available error diffusion method.[4,5] The halftoning volumes used in the following examples consists of 256 maps where each map has the size of 128×128 pixels. Please note that since the parameters of the reproduction process are unknown to us, none of the examples have been compensated for dot gain. Thus, when comparing the examples produced with clustered dot to the ones produced with dispersed dot, the latter will appear darker since the effect of the dot gain is greater in such techniques.

Figure 1 shows the results from halftoning a 256-step gray-level wedge with 5 different techniques. The first figure shows the result from a conventional clustered dot technique. In all of the next three figures b), c), and d) the wedge has been halftoned with the method discussed in this paper but with different policies when computing the maps. In the first of these figures, figure b), the computation is done without the necessary correlation between the maps. The result is a noisy image with unpleasant discontinuities. In the second figure, the maps are computed with the interpolation technique discussed above (*Method 2*) which results in a smoother shade without the undesired discon-tinuities. Finally, in the third figure, all of the maps are optimized and correlated according to the above (*Method 1*). This method provides a similar result as the second but gives an even better result at pixel level. For instance, it produces tints with even less graininess. From this point onwards, the examples referred to as produced with our method, have all been halftoned with *Method 1*. The last figure, figure e), shows the result from an error diffusion method.

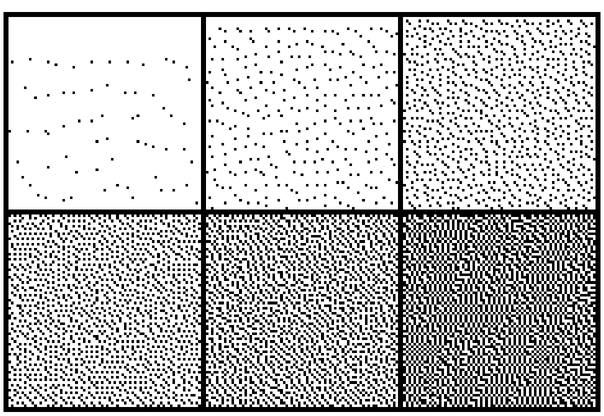

Figure 2. Tints halftoned with a commercial available Error Diffusion method.[6]

The error diffusion method provides a smooth and pleasant result without discontinuities. Though scrutinizing the image, even if not that disturbing, the typical patterns for error diffusion methods are visible in the result and these become even more obvious when halftoning larger tints as shown in Figure 2. The typical patterns are easily detected in low tint values and in the 50% tint the method gives a poor result with disturbing regularities.

Since our method (*Method 1*) is based on maps that all are optimized, the result when halftoning larger tints will still be optimal. This is shown in Figure 3 where the same tints as in Figure 2 has been halftoned but now, as expected, the result is more pleasant and without any obvious artifacts.

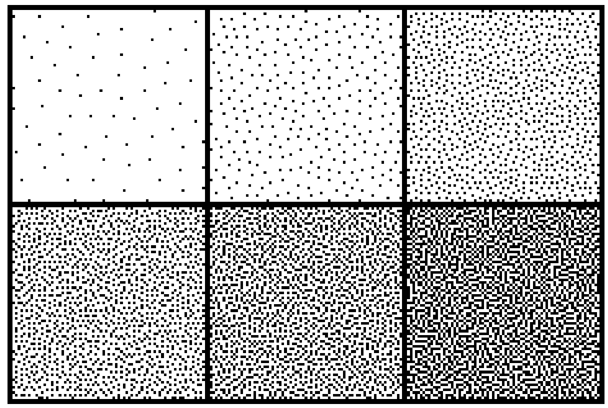

Figure 3. Tints halftoned with method described in text.

The last examples is two non-synthetical images. The first is a close up of a detail (Figure 4) and the second is an object from a natural scene (Figure 5). For comparison, we also provide the results from a conventional clustered dot method.

Discussion

When the halftoning volume is computed, the halftoning process is done in a table look-up manner which gives a great speed to the process. Due to this simplicity the method is very suitable for hardware implementation. This means that apart from the resampling process, which we always have to make to get the desired output resolution, we are able to produce the halftone dots at the same rate as we get the image information. Further, knowing the properties of the printer and the paper to use, we may compensate for the dot gain already when computing the maps, thus introducing the dot compensation directly into the halftoning process. The compensation may be done for each tint individually which implicate that each tint printed should correspond well to the original.

One drawback with this halftoning technique is that it requires more memory than a threshold matrix-based technique. For instance, if we were to halftone an image containing 256 tone values, our method would require 32 times as much memory as a conventional thresholding method using the same matrix size.

The results produced with the method discussed in this paper has encouraged us to consider halftoning colour images using this technique. There are, of course, additional problems when introducing colour, but we believe that these problems are far from invincible and that our method is well suited for colour as well.

(a)

(b)

Figure 4. A close up a halftoned detail. a) Method 1, described in text b) Conventional clustered dot method.

(a)

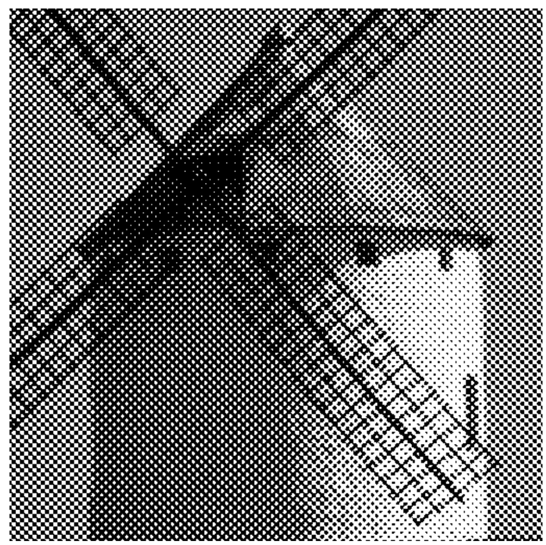

(b)

Figure 5. An object from a natural scene. a) Method 1, described in text b) Conventional clustered dot method.

Conclusion

With the frequency modulated halftoning technique presented in this paper, we have shown that it is possible to break the constraints of order dithering and let every tint be individually optimized without introducing unpleasant discontinuity effects. Further, the method produce tints without the graininess that threshold matrices usually exhibit and thus, produce halftoned images of improved quality. Since the halftoning procedure is done in a similar manner to ordered dithering, the technique can be used without any reduction in speed. The examples shows the increase in image quality compared to conventional techniques and also indicates the potential of this method. Finally, the advantages and drawbacks of this new technique were discussed.

References

1. D. A. Carrara, M. Analoui and J. P. Allebach, "Recent Progress in Digital Halftoning", *Proceedings of the IS&T 8th Congress on Advances in Non-Impact Printing Technologies,* Williamsburg, VA, October 25–30, 1992, pp. 265–269; ***see Recent Progress in Digital Halftoning, Vol. I, pg. 139.***
2. S. Gooran, M. Österberg and B. Kruse, "Hybrid Halftoning — A Novel Algorithm for Using Multiple Halftoning Techniques", *Proc. of the IS&T International Conference on Digital Printing Technologies (NIP12),* Sept. 1996; ***see pg. 187, this publication***.
3. T. Mitsa and K. J. Parker, "Digital Halftoning Technique Using a Blue-Noise Mask", *J. Opt. Soc. Am. A,* vol. **9**, pp. 1920–1929, 1992.
4. R. A. Ulichney, *Digital Halftoning*, The MIT Press, Cambridge, Mass., 1987.
5. R. Ulichney, "The Void-and-Cluster Method for Dither Array Generation", *Proc. SPIE, Human Vision, Visual Processing and Digital Display IV*, vol. **1913**, pp. 332-343, Feb. 1–4, 1993.
6. Adobe Photoshop 3.0

❋ Previously published in *NIP12*, pp. 72–77, 1996.