

Descreening printer-ready images

A. Ravishankar Rao, Frederick Mintzer, Gerhard Thompson
I.B.M., T. J. Watson Research Ctr.
Yorktown Heights, N.Y. 10598
e-mail: rao@watson.ibm.com
Tel. (914)-945-3553

Abstract

In this paper we address the problem of taking a bi-level image prepared for one printer, and processing it such that when it is printed on a second printer, the outputs of the two printers are as visually similar as possible. We present novel algorithms for this processing, which includes descreening and rehalftoning. The potential application for this algorithm is the printing of a digital image on a printer that is different from the one the image was originally formatted for.

Our algorithm segments the image into text and halftone areas based on gradient density. The halftone areas are analyzed through their Fourier power spectra to yield their repetitive frequencies. A smoothing filter is chosen based on the repetitive frequency to perform descreening. This results in an intermediate gray level image, which is then rehalftoned using a technique suitable for the second printer.

We implemented this algorithm, and found that it works satisfactorily on several test images.

Keywords: halftone, descreen, digital printing, binary images

1 Introduction

With the growing use of digital input means such as digital scanners and cameras, and digital printers, users will demand more flexibility in their choice of input and output devices. Thus it becomes important to be able to repurpose documents and images automatically and efficiently. If a better printer becomes available, a user will be interested in taking pictures scanned and formatted for an original printer and printing them on an improved printer. Since the original scanned document or object is invariably not available, it is crucial for this repurposing to be done automatically from the available halftone.

An instance of this problem is to take a bi-level image prepared for one printer, and process it such that when it is printed on a second printer, the outputs of the two printers are as visually similar as possible. An important component of this problem is known as "descreening" or "inverse halftoning" in the literature.

Currently, there is no satisfactory solution to this problem. Some of the obstacles in solving this problem are as follows. (1) Image and text areas need to be treated differently by a printer. Typically, image areas are halftoned with a repetitive screen, whereas text areas are not halftoned. Thus a composite document containing both text and images needs to be segmented accurately. (2) If a halftoned image prepared for one printer is printed on a second printer, there is a possibility of introducing Moire artifacts (interference fringes), which is a well known phonomenon. This shows up as an undesirable print quality problem, which can be avoided only with the proper processing techniques described in the current paper.

The problems of descreening and inverse halftoning has been studied in the literature. Wong [1] describes a method to reconstruct error-diffused images. Xiong *et al* [2] present a technique to perform inverse halftoning using wavelet transforms. Both these techniques focus on error-diffused halftones.

2 Methods

The input to our system is a bi-level image, containing one pixel per bit, where a "1" signifies that a dot is printed and "0" signifies no dot is printed. The input image is first segmented into halftone and text areas. The halftone areas are then analyzed to determine their periodicity. The periodic halftones are converted to an intermediate gray level representation which is then re-halftoned. The re-halftoned image areas are combined with the non-periodic image areas and text areas to generate the final bi-level image.

2.1 Segmentation into text and halftone areas

By separating the text and image areas, we can ensure that the text areas are not halftoned, thereby preserving their sharpness.

We develop a segmentation technique based on gradient information. Halftone areas contain more transitions between black and white values than text areas. Since gradients measure transitions, a simple measure to consider is the gradient activity per unit area. We expect the gradient activity per unit area to be high in halftone areas and low in text areas. This observation is borne out in practice, as shown in Figure 1, which shows the distribution of gradient activity for text and halftone regions. The two distributions are clearly different, and it is virtually impossible to find text regions which have a gradient activity of more than 50%. Whereas, there are a significant number of halftone pixels which have a gradient activity of more than 50%.

Thus, a simple rule to segment text from image halftone areas is:

```
If gradient-activity > gradient-threshold
    pixel is image halftone
else
    pixel is text
```

The size of the window used is $M = 25$, $N = 25$. The gradient threshold employed in the above rule is 50%. After the above segmentation, we compute bounding boxes of the halftoned image areas, which are then processed separately from the text areas as follows.

2.2 Detection of periodicity

We focus most of our attention on detecting the presence of screened halftone (as opposed to error-diffused halftones), as these are the most widely used halftones. The first step in processing is to idenitify the periodicity of the halftone. We developed a method which we call "shift and accumulate".

2.2.1 Shift and accumulate

Without loss of generality, let us assume we are trying to detect periodicity in the horizontal direction, across columns. The procedure is the same for the vertical direction, along rows. We start with an accumulator array, which is initialized to the first row. The second row is shifted by one pixel to the left and added to the accumulator. The third row is shifted two pixels to the left and added to the accumulator, and so on. By the time we have traversed the entire image, the accumulator array contains the accumulated sum of shifted projections. This accumulator array is a one dimensional representation of the entire image, which amplifies any periodicity present in the original image. Periodic components reinforce each other in the final accumulated sum.

Note that the shift-and-accumulate method as presented works best with a 45^0 halftone screen. However, it can be easily adapted to other screen angles by shifting every M^{th} row by N pixels such that the screen angle equals $tan^{-1}(N/M)$.

2.2.2 First order differencing

Let us denote the accumulated sum by a sequence s. We perform a first order differencing operation, as is customary in time series analysis to remove trends [3]. This results in a new sequence s_2 where the i^{th} element of s_2 is

$$s_2(i) = s(i+1) - s(i) \quad (1)$$

We then take the Fourier Transform of s_2 using an FFT implementation such as the one described in [4], and compute the corresponding power spectrum P. The power spectrum depicts the distribution of spectral energy content at different frequencies. The dominant peak in the power spectrum is identified easily as the maximum value in the sequence P. In order to test for periodicity, we calculate the height of the second peak, and find the ratio between the first and second peaks. If the ratio is high, say > 3, the original sequence is deemed to be periodic, and the period is determined by the location of the first peak. If the ratio is low, the spectral energy is distributed more or less equally, implying the original sequence does not contain dominant periodic function. Other tests of periodicity are possible such as [5].

2.3 Generation of an intermediate gray level image

Taking the horizontal and vertical periodicity detected in the previous step, we create an $M \times N$ smoothing box filter, where M is the half the vertical period and N the horizontal period. This smoothing filter is moved across the input image (each of the areas of the original image that is deemed to be a periodic halftone area) to generate an intermediate gray level image. The center of the $M \times N$ box is assigned the smoothed value.

All the coefficients within the $M \times N$ box filter are set to "1". However, other assignments are also possible, such as a Gaussian kernel. Any set of filter coefficients that is low-pass will perform comparably.

It is important to note that the intermediate gray level image does not contain any Moire patterns, as the size of the box filter is dependent on the periodicity of the original halftone.

2.4 Rehalftoning

This intermediate gray level image is then rehalftoned using any desired halftone technique. For instance, we can use a dither matrix. The important point to note is that since the intermediate gray level image has been properly generated, we do not get any Moire patterns during the rehalftoning process.

During this rehalftoning process, we can make use of the printer characterization parameters of the second printer in order to make the printer halftoned image match the output of the first printer. This is how we achieve similarity between the two printed outputs.

2.5 Combining

Finally, the rehalftoned areas are combined with the areas that were not processed (i.e. non-halftone areas such as text) in order to generate a complete image.

Figure 2 shows the image that is used to test our algorithm.

The halftone image areas are segmented from the text, descreened, and rehalftoned. The rehalftoned areas are combined with the areas that were not processed (i.e. non-halftone areas such as text) in order to generate a complete image. Figure 3 shows the result after using a calibrated 3x6 dither matrix to rehalftone the image areas.

This example shows the power of our descreening and rehalftoning approach.

3 Conclusion

We have demonstrated a technique that can be successfully used to descreen printer ready images. There are several potential uses for our application.

1. An original halftone prepared for one printer can be printed on another printer without loss of quality and the introduction of artifacts.
2. This technique can also be used to perform image enhancement tasks. For instance, if the original halftoned image had poor contrast, we re-create an intermediate gray level representation, enhance the contrast, and then re-halftone the enhanced image. We can also add some random noise to improve the appearance of texture in the textured areas of the image.

Acknowledgements

We are grateful to our colleagues Mikel Stanich and Jay Smith from the IBM Printing Systems Company, Boulder, for providing us the motivation and support to work on this problem, and for useful feedback.

References

[1] P. W. Wong, "Inverse halftoning and kernel estimation for error diffusion," *IEEE Trans. Image Processing*, no. 4, pp. 486–498, 1995.

[2] Z. Xiong, M. Orchard, and K. Ramachandran, "A wavelet-based approach to inverse halftoning," in *SPIE Vol. 3018*.

[3] W. Vandaele, *Applied time series and Box-Jenkins models*. New York: Academic Press, 1983.

[4] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1988.

[5] L. G. Tatum, "Control charts for the detection of a periodic component," *Technometrics*, no. 2, pp. 152–60, May 1996.

Author Biographies

A. Ravishankar Rao

A. Ravishankar Rao is a research staff member with the Image Applications group at the IBM T.J. Watson Research Center, Yorktown Heights, New York. His research interests include digital image processing, computer vision and data mining. He has published a book entitled "A Taxonomy for Texture Description and Identification," as well as several research papers. He chaired the SPIE Conference on Machine Vision and Applications from 1996-98.

He received his B.Tech degree in electrical engineering from the Indian Institute of Technology (Kanpur) in 1984, and his M.S. and Ph.D. degrees in computer engineering from the University of Michigan, Ann Arbor, in 1986 and 1989 respectively. He is a member of IEEE and SPIE.

F. C. Mintzer

Fred Mintzer received the PhD degree in Electrical Engineering from Princeton University in 1978. For the past ten years, he has led IBM projects devoted to developing high-quality image technologies for digital

IS&Ts NIP 14: 1998 International Conference on Digital Printing Technologies

library applications, including color-accurate scanning, image enhancement, image compression and image display. These technologies have been essential components to IBM projects with the America painter Andrew Wyeth, with the National Gallery of Art (USA) and with the Vatican Library. Recently, he and his colleagues have begun work on image processing methods that would discourage misappropriation of publicly-available images. Dr. Mintzer is a Senior Member of the IEEE and an IBM Master Inventor. His work has resulted in more than ten patents and thirty publications.

3.1 Gerhard R. Thompson

Mr. Thompson joined IBM in 1968 and worked on a variety of projects including performance evaluation, hardware and software modeling, and circuit and logic design of large-scale integrated chips.

Since 1982 he has been with Image Applications in the IBM Research Division in New York. Here he developed various image processing algorithms including binary and gray-scale rotation, scaling, and digital halftoning for both monochrome and color images. His present activities include developing algorithms and writing software for use in the scanning, display, and printing of digital images. He holds numerous patents in these areas.

Mr. Thompson received a B.S.E.E degree from Drexel University in 1967, an M.S.E.E from Princeton in 1968, and an M.S. in computer science from Syracuse University in 1983.

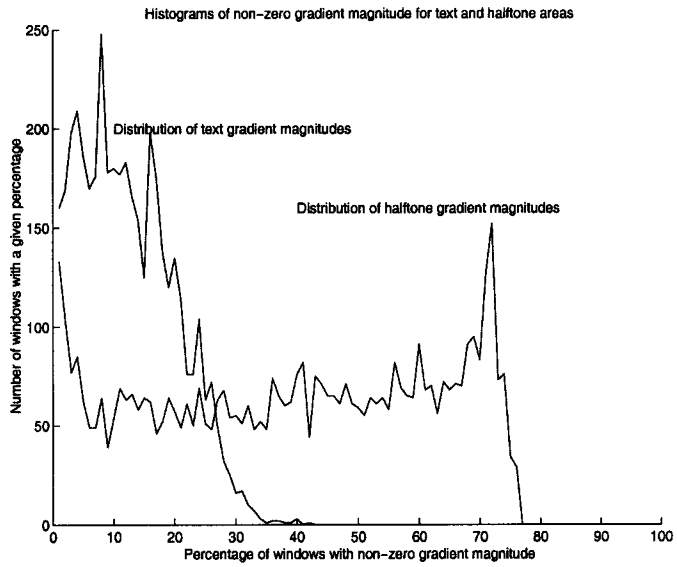

Figure 1: The distribution of gradient activity for text vs. halftone regions. A window size of 25x25 was used to calculate the gradient activity per unit area.

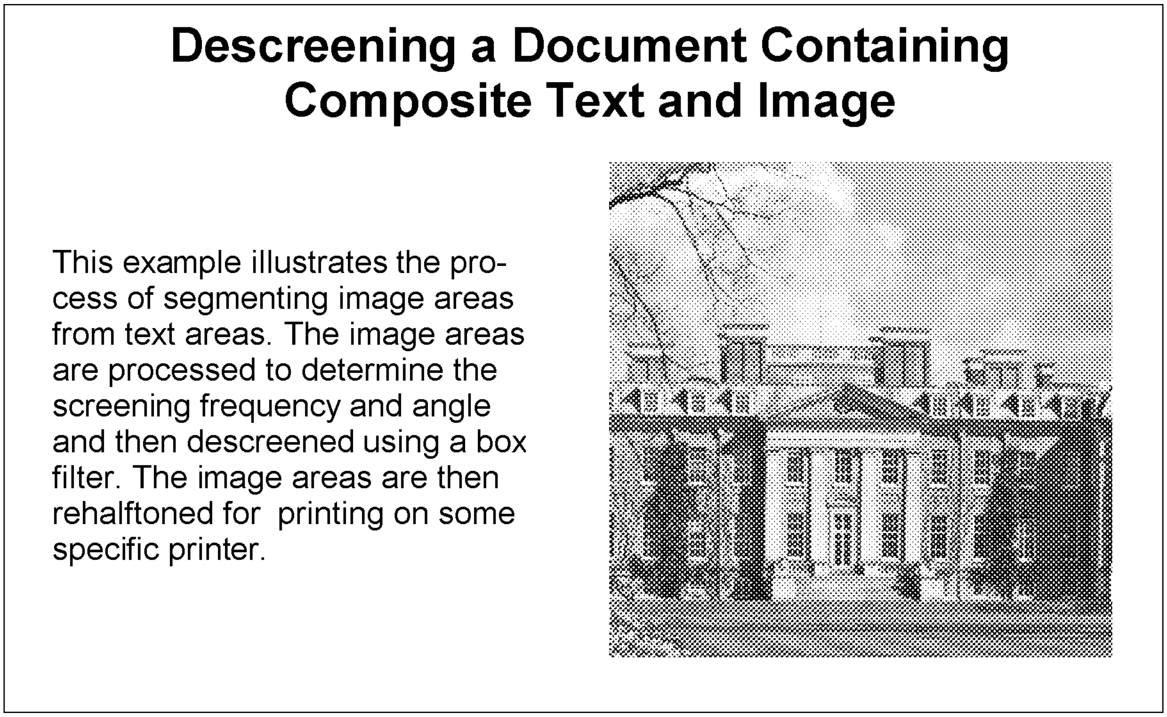

Figure 2: The test image that will be used to illustrate the various steps of our algorithm. This contains a combination of text and image areas. The image on the right is halftoned with a periodicity of 10 pixels. The original image size is 1812x1107 pixels.

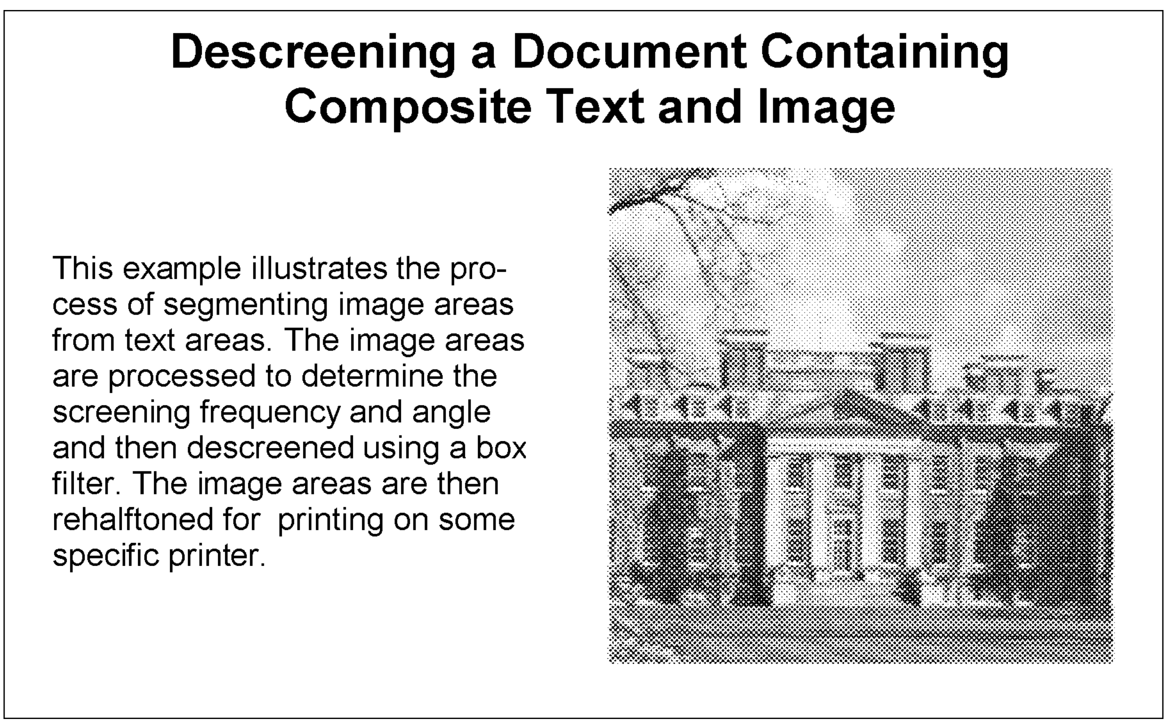

Figure 3: Result of using a calibrated 3x6 dither matrix.