RC22274 (W0112-018) December 4, 2001
Computer Science

IBM Research Report

Multilevel Color Halftoning

Joan L. Mitchell, Gehard Thompson,Chai Wah Wu
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Timothy J. Trenary, Yue Qiao
IBM Printing Systems Division
Boulder, CO 80501

Multilevel Color Halftoning

Joan L. Mitchell, Gehard Thompson,Chai Wah Wu
IBM T. J. Watson Research Center
Yorktown Heights, NY10598

Timothy J. Trenary, Yue Qiao
IBM Printing Systems Division
Boulder, CO 80501

Abstract

A new high-speed color printer uses new image screening methods that lead to superior-quality images. The screen rulings as well as the screening angles are varied. The cyan, magenta, and yellow screens are around 200 lines per inch. The black screen starts at 212 lines per inch at 45 degrees and progresses to 300 lines per inch at zero degrees when half the pels in the basic cell are on. This improves the details in the black color plane without introducing significant *moiré patterns*.

Since the printer is capable of 4 bits/color, threshold matrices are used to determine onset of printing. The actual output intensities are selected from look-up-tables indexed by the input value minus the threshold. Experiments showed that the intermediate intensity values reproduced best as leading pels on vertical marks. Results for mixing halftones at significantly different screen rulings are illustrated.

1. Introduction

The IBM InfoColor 130+ prints 130 pages per minute at a resolution of 600 dots per inch (dpi). Eight print stations deposit cyan (C), yellow (Y), magenta (M), and black (K) toners on the front and back of the continuous-roll paper. Two pages can be printed next to each other on both sides. The paper moves through the printer at approximately three inches per second.

Unlike most printers that print either fully saturated color(s) or nothing so that the paper shows through as the background color, the Xeikon marking engine prints fourteen intermediate intensities in addition to white (no toner) and a fully saturated color (i.e. 4-bits per color). However, large areas printed with only intermediate values showed objectionable mottling, paper texture, and sensitivity to operating conditions. Patches printed at a constant intermediate value looked more like leather than a smooth continuous-tone output. The next section details some of the experiments which led us to the conclusion that the intermediate values printed most consistently when placed next to fully saturated values. We call this printing of the intermediate values on the edges of a cluster of fully saturated pels as "gray-on-edges" since most of the experiments were performed on the black component where we could most easily see the effects.

Section 3 explains how binary threshold matrices can be combined with a simple lookup table (*LUT*) for multilevel gray-on-edges printing. Section 4 illustrates the mixing of different lines per inch screens to improve detail. Section 5 describes the design of the super cell clustered dither with stochastic interpolation binary threshold matrices. Section 6 shows the measured reflectance of the black component.

2. Experiments

An experiment with two inch square patches containing alternating white and fully saturated lines (i.e. 300 line pairs per inch printed vertically or horizontally) printed with I, 2, 3, or 4 colors yielded unexpected results. The colors printed as vertical lines appeared to be more saturated and uniform than when printed as horizontal lines. The patches containing vertical lines appeared continuous. Under magnification the white and saturated lines appeared to be about the same width.

Figure 1 shows some of the other experiments which helped determine how to obtain consistent intermediate intensity print quality. Figure la indicates the basic 6 inch x 5.5 inch grid of 1/2 inch square patches. The length of the 1 x M long mark changes along one axis from an isolated pel to a complete line through the N pel x 12 pel cells. The other axis increased the number (N-l) of white rows between marks. Since N started at 2, at least one white row separated the marks. A separate page was created for each output level L and color component. Figure lb shows the cell containing a mark with constant level L. Figure 1c modified the first pel in the mark to a fully saturated value (S) of 15. The rest of the pels (M-1) were left at the intermediate level L. Figure 1d shows the fully saturated pel moved to the last pel in the mark. The next two figures illustrate cells with the mark primarily made of saturated pels with a leading (Figure 1e) or trailing (Figure 1f) pel with level L. Not all experiments indicated by the indices were performed.

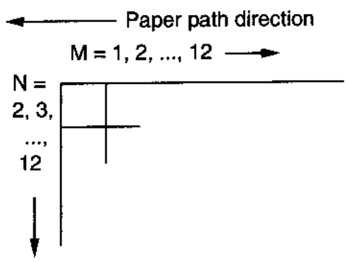

(a)

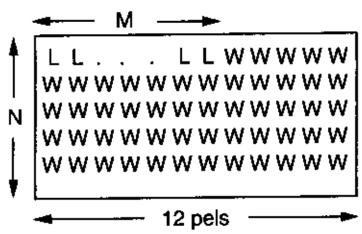

(b)

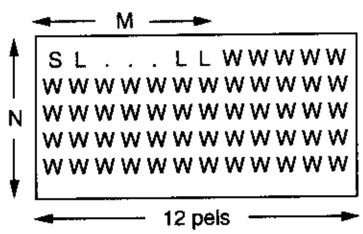

(c)

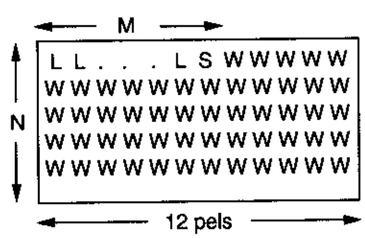

(d)

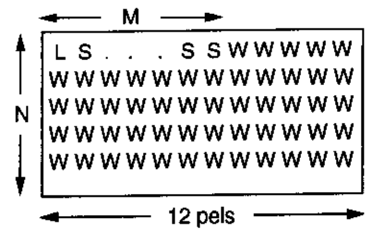

(e)

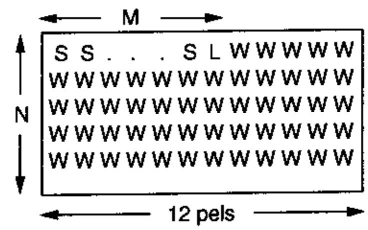

(f)

Figure 1. Experiments to determine consistent intermediate level print quality. a) For each level L (L=1, 2, ..., 15) and color (C,M,Y,K) print a 6 inch x 5.5 inch grid of 1/2 inch square patches filled with N x 12 cells containing a 1 x M mark in a white background. b) Cell with mark at level L. c) Cell with leading pel saturated (S=15). d) Cell with trailing pel saturated . e) Cell with non-saturated leading pel. f) Cell with non-saturated trailing pel.

The paper path direction is shown along the longer axis of the printed marks. A subset of these experiments was repeated with the paper path perpendicular to the mark direction and confirmed that marks of all types printed most consistently along the print direction. In addition, we concluded that given a choice, the intermediate value should be a leading pel of a vertical mark rather than the trailing pel where its smear into the white showed.

Some more experiments were designed to see if pairs of levels would print reliably once a certain difference between the levels was reached. Figure 2a shows a 7.5 inch x 7.5 inch grid of 1/2 inch square patches with the first level L1 changing from 1 to 15 across the page and the other level L2 changing from 1 to 15 down the page. A separate page was created for each color component. Each 1/2 inch square patch contained repeating 4 x 4 pel cells. In Figure 2b a pair of L1 and L2 pels is surrounded by the white background in the 4 x 4 cell. In Figure 2c a pair of L1 and L2 pels is followed by two saturated pels to create lines every fourth row. Until at least one of the levels was almost saturated, the print quality was not significantly improved for different levels as opposed to two pels of the same level in a white background. Two pels of intermediate levels touching saturated pels in a line printed well at almost every level.

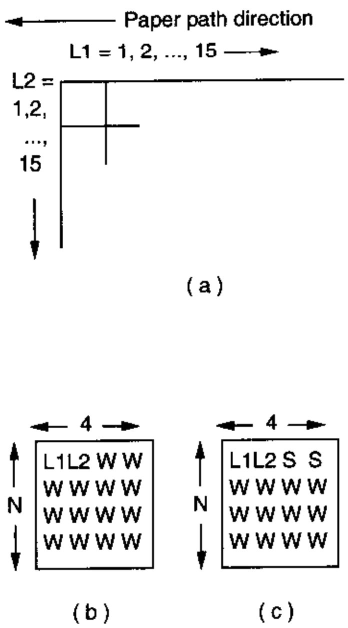

Figure 2. Additional experiments to determine consistent intermediate level print quality. a) For each color (C, M, Y, and K) print a 7.5 x 7.5 inch grid of 1/2 inch square patches filled with 4 x 4 cells. b) Cell with L1 and L2 pair of pels in a white background. c) Cell with L1 and L2 pair of pels followed by two saturated pels to make a line every fourth row in a white background.

3. Gray-on-Edges Halftoning

The experiments described in Section 2 showed the desirability of printing the intermediate levels next to saturated pels. Clustered dither threshold matrices add each new pel next to the cluster of previous pels. Figure 3 shows an example of the pattern growth for a 3x3 cluster. If, instead of instantly printing fully saturated pels, intermediate levels were used for some range of inputs greater than or equal to the threshold, a binary threshold matrix could still determine the onset of printing.

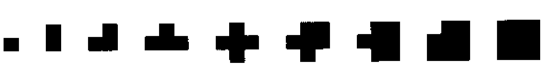

Figure 3. Dither algorithm basic (3x3) cell patterns.

Figure 4 gives a block diagram that shows how the threshold value (TH) for the given position (modulo the horizontal and vertical dimensions of the matrix) is subtracted from the input value (IN) to select the 4-bit output value (OUT). The LUT is shown as a 512 entry table. Since all negative values are forced to zero, a 256 entry table is sufficient and was used in custom high-speed hardware for halftoning [1].

The use of the LUT gives complete flexibility for the sequence of levels. We always used a ramp of increasing intermediate values. From experiments with isolated intermediate level pels (M=1 in Figure 1) we concluded that the first few levels were too unreliable (and did not measure as density differences) to use as isolated or leading pels on marks surrounded by white pels, so the LUT skipped at least the first two levels. Between saturated pels, the OUT values could start at 1 and have a measurable change in output for each increment.

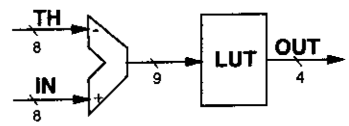

Figure 4. Block diagram of gray-on-edges halftoning where the look up table (LUT) output is 0 if the threshold (TH) is greater than the input (IN) and one of the output levels L (L=0,1,..,15) otherwise.

4. Mixing Different Lines Per Inch (LPI) Screens

Since the alternating vertical lines printed quite reliably, the black halftone screen basic cell was converted from a 2x4 horizontal block into a 4x2 vertical block. Instead of creating a 2x2 checkerboard pattern when the left half the pels were on, the pels within the 4x2 block were shuffled to grow upwards to fill in the left side as a 4 x 1 line. This created a 300 lpi screen (parallel to the paper path direction) when half the pels were on. Since the first pels were isolated pels, a 45 degree 212 lpi pattern was created for low input values which changed into a 300 lpi pattern at 0 degrees near the mid range. This improved preservation of black details.

Figure 5 shows an enlarged version of the moiré patterns generated by two overlapping halftone screens. One of the screens is rotated an additional 1.6 degrees between the top and bottom boxes. The moires are not noticeably visible in Figures 5a and 5a' when two screens of the same lpi have a 45 degree relationship. When both screens are at 0 degrees, a small shift in angle creates large, visible moiré patterns as shown in Figures 5b and 5b'. However, when the lpi of one of the screens is increased by 50% to 300 lpi, the moires are much less noticeable and too small for the unaided eye to detect (see enlarged versions in Figures 5c and 5c').

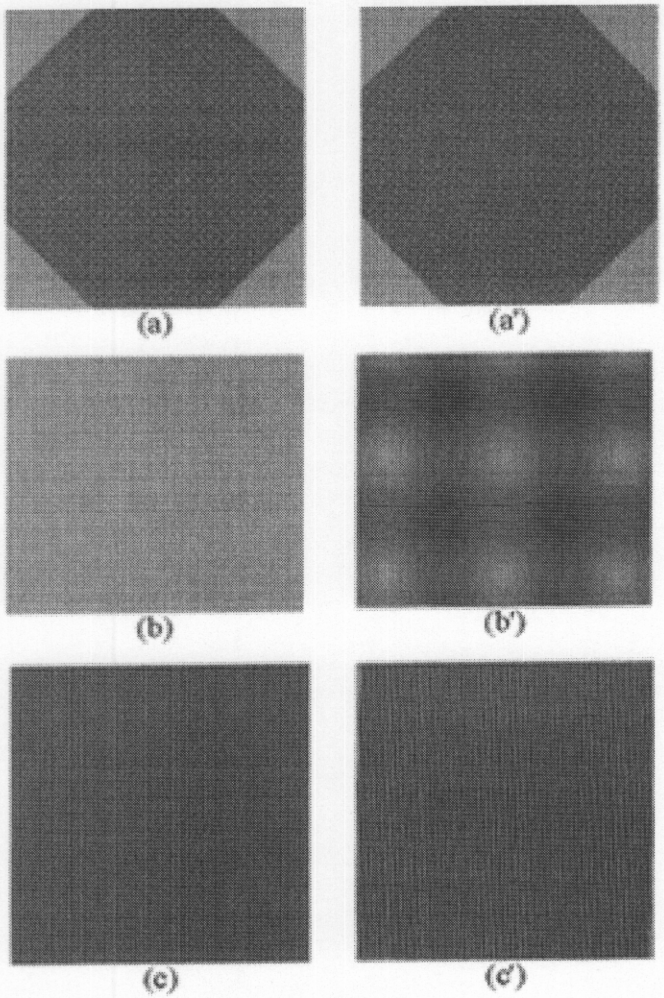

Figure 5. Enlarged moiré patterns generated by two overlapping halftone screens. a) 200 lpi at 0 and 45 degrees. a') 200 lpi at 0 and 46.6 degrees. b) 200 lpi at 0 and 0 degrees. b') 200 lpi at 0 and 1.6 degrees. c) 300 lpi at 0 degrees and 200 lpi at 0 degrees. c') 300 lpi at 0 degrees and 200 lpi at 1.6 degrees.

5. Clustered Dither Halftoning With "Stochastic" Interpolation

In traditional clustered dither halftoning methods, there is a tradeoff between cluster cell size and the number of gray levels in the dither; the smaller the cell size, the smaller the number of gray levels. A small cluster size generates halftoned images with sharper edges and finer details, whereas a large number of gray levels prevents contouring and posterization. To allow a large number of gray levels while maintaining a small cluster size, super cell techniques are used, in which several halftone cells are grown out of synch to obtain many gray levels. However, commonly used super cell clustered dither produce unpleasant patterns at certain gray levels. Constrained noise was used to replace the unpleasant patterns with more pleasing patterns. For more details on this algorithm, the reader is referred to Refs. [2, 3, 4, 5].

The goal is to create a dither mask such that for each gray level the corresponding image of uniform gray is rendered into a nice halftone pattern by the dither mask. Because of the way the dither algorithm is defined, these halftone patterns satisfy the stacking condition: the pattern for a lighter gray level g1 is a subset of the pattern for a darker gray level g2, (i.e. g2 > g1) (see Figure 3).

To increase the number of gray levels while preserving the cluster size, a super cell clustered mask can be constructed as follows. A collection of dither masks (each containing one halftone cell) are tiled to form a larger dither mask. The clusters in these cells are grown asynchronously, generating many gray levels while preserving the resolution of a single cell as shown in Figure 6. The unpleasant patterns are corrected as follows: Assuming gray levels g1 < g2 < g3, where the patterns for g1 and g3 are pleasing, while the pattern for g2 is not. The algorithm will replace g2 with a new pattern while preserving the patterns for g1 and g3. Let the pattern for gi be denoted by Pi. Because of the stacking conditions, the black pels in P1 are a subset of the black pels in P2 which in turn are a subset of the black pels in P3. Therefore the pattern for P2 is obtained by setting the black pels in P1 (which are called the *fixed* pels) and choosing some black pels which are in P3 but not in P1 (which are called the free pels). The choice will be determined by a stochastic algorithm, while satisfying the constraint that the fixed pels cannot move. A number of free pels are rearranged among all the free pels using a stochastic algorithm which also considers the fixed pels. In addition, a clustering constraint is satisfied to make sure the chosen free pels are adjacent to the fixed pels.

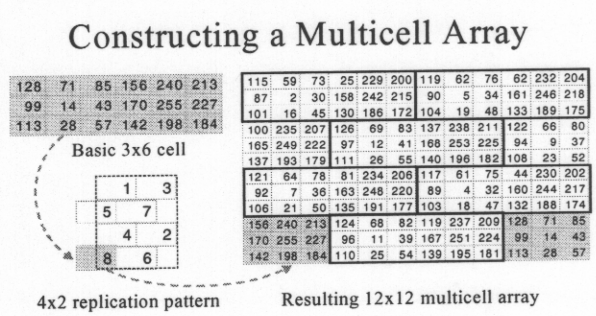

Figure 6. Constructing a multicell array.

alftone screens nominally at about 200 lpi had been developed for an earlier binary color printer using the clustered dither with stochastic interpolation technique described in this section. For this 4-bit per color printer, some of the visible noise in the output was suppressed by using a slow ramp in which the output was held constant for four inputs and then incremented to the next level for the next four inputs. Since the output reached full saturation at 44 levels beyond the threshold (TH+44), the threshold levels had to be shifted to smaller values. The shifted values were split equally into threshold values 1, 2, 3, and 4.

6. Lightness Measurements of the Screens

The GretagMacbeth SpectroMat was used to measure the gray scales of cyan, magenta, yellow, and black multilevel halftones. Figure 7 shows lightness (L*) as a function of the percentage of black toner. There is good agreement between the gray scales of the four halftone screens. The lightness changes smoothly in the light region and then linearly in the mid to dark regions. This preserves details in the light region and has sharp contrast in the mid to dark regions. The results for the yellow component are the noisiest because the lightness changed the least.

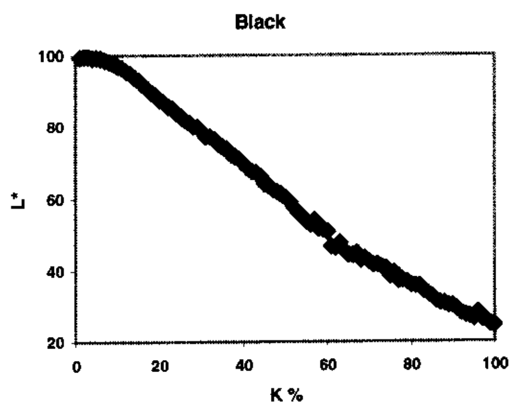

Figure 7. Plot of the Lightness (L) of the black color component.*

Conclusions

Cluster dithered threshold matrices with stochastic interpolation halftoning can be combined with intermediate intensity levels for high-speed color multilevel printing. Experiments proved the reliability and consistency of "gray-on-edges" for printing rather than uniform areas of intermediate values. Additional black quality was achieved by increasing the black line screen frequency for mid range values by about 50%.

Acknowledgments

We appreciate and acknowledge that many of our coworkers helped us achieve our best-of-breed halftone goal. In particular we would like to mention Danielle Dittrich for her insight into why the line screen mixing worked; Fritz Obermeyer for his help; Charles Morris III for the original hardware design; Michael Brady for hardware that did not re-halftone binary input data; Nenad Rijavec for the simulation environment, code and encouragement; Christine Basher for printing; and Charles Tresser for some of the stochastic halftones. We would also like to give a special thanks to our managers: Jim Christensen, Marco Martens, Fred Mintzer, Brenda Payne, Mike Simpson and Marilyn Smith.

References

1. Michael T. Brady, Charles H. Morris III, and Joan L. Mitchell, "High Speed Multilevel Color Halftoning Hardware," Color Hard Copy and Applications Conference, part of Electronic Imaging Symposium, San Jose, January 2001 (submitted).

2. Daniel L. Lau and Gonzalo R. Arce, Modern Digital Halftoning, Marcel Dekker, Inc: New York (c) 2001.

3. Gerhard Robert Thompson, Charles Philippe Tresser, and Chai Wah Wu, "Clustered Aperiodic Mask," US Patent 5,917,951 issued Jun. 29, 1999.

4. Gerhard Robert Thompson, Charles Philippe Tresser and Chai Wah Wu, "Multicell clustered mask with blue noise adjustments," US Patent 6,025,930 issued Feb. 15, 2000.

5. Chai Wah Wu, Charles P. Tresser, Gerhard R. Thompson, Mikel Stanich, "Supercell dither masks with constrained blue noise interpolation," NIP 17 Conference, Ft. Lauderdale, FL, 2001.

Biography

Joan L. Mitchell received BS, MS, and PhD. physics degrees from Stanford University and the University of Illinois in 1969, 1971, and 1974, respectively. She has worked in IBM Research, Marketing, and the Printing Systems Division in the fields of image processing and data compression. She has coauthored JPEG and MPEG books and 36 patents. She is a member of the IBM Academy of Technology, APS, IS&T, Sigma Xi, and an IEEE, and IBM Fellow.