

XEROX

PALO ALTO RESEARCH CENTER
Computer Sciences Laboratory
16 August 1976

To: CSL, SSL

From: Joe Maleson

Subject: Pictures, PRESS, and you

Filed on <MALESON>HalftoneMemo.Press; appendices on <MALESON>HalftonePlates.Press

This memorandum describes digitally sampled images, general techniques for dealing with them, and facilities now available in the 3100 Press printing program for printing them. An updated version of the PRESS File Format memo describes the new commands available for including sampled images in Press files.

A collection of pictures printed with the various parameter settings available is included as a set of two appendices.

What is a digitally sampled image?

A continuous tone image is a picture (like a photograph) with a range of intensities from black to white. A particular intensity is sometimes called a gray tone or simply tone. A continuous tone image may be stored on a digital computer by sampling it at some resolution. That is, the average intensity at some number of points is determined using a scanner (for example, a television camera). For a given point, the scanner returns a number ranging from 0 for black to some maximum value for white. That value is stored in a matrix. That matrix, when completed, is the sampled image. If the matrix has m columns and n rows, and the maximum intensity value (white) is w , then the resolution of the sampled image is given as " $m \times n \log_2 w$ bit samples."

Example

The image on the right was scanned in at the University of Utah. There are 512x512 8-bit samples. It is stored on maxc in AIS format as <AIS>Utah-Building.Ais. AIS stands for "Array of Intensity Samples." This format is described in [2].

Halftoning is the general term used to describe the printing of continuous tone images with only one color of ink (black). Thus, this kind of picture is often called a halftone.

The halftoning technique used here is error-distribution with a 65-line screen added. This technique is described in detail below.



Utah building scanned as 512x512 8-bit samples


How is a halftone printed?

The problem of halftoning is to simulate various gray tones using only one ink color (black). Conventional printing accomplishes this by adding a screen to a negative of the image. The screen is usually a regular pattern made up of square elements which are graded smoothly from white to nearly black. The continuous tone negative and screen are printed together on high contrast material, which develops black and white, but not gray.

The screen is usually too light to print by itself, so that image areas which are entirely white or black remain unchanged, but intermediate gray tones combine with the screen to form various size dots, with size proportional to the intensity of the image area. Note that any screen pattern can be used, and the square pattern



an example screen element



an example screen

shown here, while a common and useful screen, is not the only one available. Various screen angles may be used. The example screen is called a 0 (or 90) degree screen, because the square elements are collinear with the paper edges. Usually, the screen is turned 45 degrees. In color printing, different screen angles are used for each color. The number of square elements per inch is referred to as screen frequency.

There are many digital methods for producing halftone prints from a computer. The simplest technique is to print a black dot for each sample point below some pre-selected threshold. This kind of thresholding does not produce satisfactory results: while the print may hold some artistic interest, it does not do a good job of reproducing a continuous tone original.

A slightly better method is to define a pattern for each gray tone to be printed, and then printing the appropriate pattern for each sample point. (The popular "computer image printouts" choose patterns from a set of alphanumeric characters.) One major difficulty here is in the size of the patterns used. To reproduce 256 gray levels, the pattern must contain 256 dots. If we print 500 dots to the inch, this allows us only about 30 samples per inch, or 240 samples across an 8 inch page. The large pattern elements cause the print to look grainy. If a smaller pattern is used, the number of gray levels that can be represented decreases. A 5x5 pattern can only represent 25 distinct tones.

A close analog to the photographic screen technique can be achieved by sampling a continuous tone screen pattern (or generating a matrix equivalent to sampling the desired pattern), adding this digital screen to a sampled image, taking the sum as a new sampled image, and doing the simple threshold operation. This technique is often referred to as partial dotting because the shapes of the dot clusters printed will be asymmetrical whenever a digital screen element is added to an area of the original image containing an intensity contour. An intensity contour occurs when there is a rapid change from one gray level to another, as in the edge of a dark object against a light background (or vice versa).

The partial dot technique is not necessarily optimal. The print may be darker or lighter than the original, depending on chance combinations of threshold values with sample points. More importantly, it is difficult to perceive fluctuations in the original image which are smaller than the screen frequency. The so-called "Floyd-Steinberg" algorithm (after Robert Floyd and Leo Steinberg) makes use of the following observation: in the simple threshold technique, once the printing decision for a dot has been made, the error introduced by that decision is known. If a sample is 75% black, and a black dot is printed, then there is now 25% of a dot too much black on the print. Now, we can take this error value and distribute it among some neighboring sample points. When too much black is printed, the error is added to the neighborhood, effectively pushing each dot a little closer to white. When not enough black is printed, the error is subtracted. The choice of which points comprise the neighborhood, and what percentage of the error is distributed to different points, will affect the exact pattern of dots printed. So, the print reproduces exactly the average tone of the entire picture, and almost exactly the

average tone of the picture over local regions.

An interesting observation is that the choice of threshold for the error-distribution method is only important for the first few scan lines, since the error is calculated by measuring distance from white or black. For example, a threshold chosen too close to black would make the initial scan lines too light, but thereafter the negative error would bias successive sample points to translate a 50% black sample point to the chosen threshold value.

Printing problems in real-world situations

While the Floyd error distribution technique is nearly optimal in a formal sense, still, for a particular image and a particular printing device, other considerations than preservation of average gray value over local areas are important. For example, if there are fewer sample points than output points, (which is usually the case for high resolution printers like the 3100) then the array of sample points must be expanded. Simply replicating samples will eventually lead to a print covered with small squares of different intensities. Interpolation will remove sharp contours and give a "blurred" print. The printed dot may be so small that a Xerographic printer will not reproduce it. There may be mechanical effects which are biased against very different scan lines: on an interlaced CRT display, two adjacent scan lines with quite different content will appear to flicker; on the 3100 printer there is a periodic variation in scan line quality.


To produce pleasing prints in spite of these problems, it is often advantageous to combine the error distribution with the screen technique. By using a screen matrix with an equal number of positive and negative terms, the preservation of average gray level over small areas is maintained.

One of the benefits of combining error distribution with digital screening techniques is that the choice of what screen amplitude to use becomes less critical. Screen amplitude is the measure of the maximum intensity value of the screen which is added to the image. In the photographic screen process, higher amplitudes lead to darker prints, and lower amplitudes approach the simple threshold picture produced with no screen at all. The photographer can compensate for variations in images and screens by using several different exposures, and choosing the best result. However, by using error-distribution, satisfactory results are achieved over a wide range of values. The print cannot become "too dark", so that as long as the amplitude is above some minimum value, any increase only tends to smooth or blur the image. This is illustrated in Appendix 2.

The digital screen used by the Press printing program is illustrated on the preceding page. The digital analog of turning the screen 45 degrees is to form 2x2 array of square screen elements, where each of the two diagonally opposed matrices are the same, and the adjacent matrices alternate in sign. Using this 2x2 array as the "45 degree" screen element, the effect of using a rotated screen is achieved. Various screen frequencies at rotations of 0 and 45 degrees are illustrated in Appendix 1.



2x2 array of screen elements



45 degree digital screen

Present implementations

The Floyd error distribution algorithm, with the option of adding a screen of any amplitude and frequency, is implemented in the 3100 Press printing program. This program takes a file in Press format as input, converts it to a 384-dot per inch bit map on a Trident disk, and prints the bit map on the 3100, a solid-area development Xerographic machine. Press files combine text and graphics, including sampled images. A set of Press file commands has been specified so that parameters like the dynamic range of the image (values used for the darkest and lightest sample points) may be included in the file. The default value for black is 0. The default for white depends on the sample size. An n-bit sample size defaults white to $2^{*n}-1$ (8-bit sample defaults white to 255). Press file format is fully described in [1].

Information about the screen to be added may be specified in the file, or at the time the Press printing program is run. The parameters set in the file override parameters specified at run-time. Screen frequency is specified in lines per inch. Since the 3100 printer has 384 printing dots per inch, there are at most 384 matrices which can be used as screens of frequency greater than or equal to 1 line per inch. Specifying digital screens in lines per inch is deceptive, since the "ideal screen" is only roughly approximated. So, screen sizes of 193-384 lines per inch all use the same 1x1 matrix, whose only element is zero. Screen sizes of 129-192 use the same 2x2 matrix, etc. Conversion from lines/inch to matrix size is simply the integer part of $384/(\text{lines/inch})$. For screens at 45 degree angles, the size of each small matrix is the integer part of $384/(\text{ROOT2} \times \text{lines/inch})$, where ROOT2 is 1.4. Conversion tables for lines/inch to matrix size are:

Screens at 0 degrees	
lines/in : matrix size	
1	: 384
2	: 192
3	: 128
4	: 96
5	: 76
6	: 64
7	: 54
8	: 48
9	: 42
10	: 38
11	: 34
12	: 32
13	: 29
14	: 27
15	: 25
16	: 24
17	: 22
18	: 21
19	: 20
20	: 19
21	: 18
22	: 17
23-24	: 16
25	: 15
26-27	: 14
28-29	: 13
30-32	: 12
33-34	: 11
35-38	: 10
39-42	: 9
43-48	: 8
49-54	: 7
55-64	: 6
65-76	: 5
77-96	: 4
97-128	: 3
129-192	: 2
193-384	: 1

Screens at 45 degrees	
lines/in : matrix size	
1	: 274
2	: 137
3	: 91
4	: 68
5	: 54
6	: 45
7	: 39
8	: 34
9	: 30
10	: 27
11	: 24
12	: 22
13	: 21
14	: 19
15	: 18
16	: 17
17	: 16
18	: 15
19	: 14
20-21	: 13
22	: 12
23-24	: 11
25-27	: 10
28-30	: 9
31-34	: 8
35-39	: 7
40-45	: 6
46-54	: 5
55-68	: 4
69-91	: 3
92-137	: 2
138-274	: 1

Sample images not in Press format files can be stored in AIS format. AIS stands for Array of Intensity Samples. A subroutine library for dealing with AIS images is now being built. An AIS image is easily transformed into a Press format file. The AIS format is fully described in [2].

References

[1] Press file format, Newman&Sproull, on <GR-DOCS>Press.Ears

[2] AIS format, Baudelaire&Sproull, on <GR-DOCS>AIS.Ears