Presented at the *Seventh Workshop on Multi-dimensional Signal Processing,*
Lake Placid, NY, September, 1991.

Model-Based Halftoning

Thrasyvoulos N. Pappas, Nambi Seshadri
Signal Processing Research Dept.
AT&T Bell Labs
Murray Hill, NJ 07974
Telephone: (908) 582-2933
Email: pappas@research.att.com

David L. Neuhoff
EECS Department
University of Michigan
Ann Arbor, MI 48109
Telephone: (313) 764-6586
Email: neuhoff@dip.eecs.umich.edu

Digital halftoning is the process of generating a pattern of binary pixels that create the illusion of a continuous-tone image. It is necessary for display of gray-scale images in media in which the direct rendition of gray tones is impossible. The most common example is printing of gray-scale images on paper. We introduce new halftoning methods that exploit printer models as well as models of the human visual system to produce high quality images using standard laser printers (typically 300 dpi).

Halftoning relies on the model of the eye as a low-pass filter. More detailed models of visual perception have been developed, and we show that it is possible to implement halftoning methods that exploit them. The eye models we use are based on estimates of the spatial frequency sensitivity of the eye (e.g. by Mannos and Sakrison). In particular, we consider discrete-space models of the form

$$z_{i,j} = M(x_{k,l}, (k,l) \in N_{i,j}) \tag{1}$$

where the $x_{k,l}$'s are image samples, $N_{i,j}$ is a finite neighborhood of the site (i,j), the $z_{i,j}$'s are the model outputs (upon which cognition is based), and M is a sliding-window function.

Halftoning also relies on the assumption that the black area of a printed binary pattern is proportional to the fraction of ones in the pattern. This means that the area occupied by each black dot is roughly the same as the area occupied by each white dot. Thus, the "desired" shape for the black spots produced by a printer would be $T \times T$ squares, where T is the dot spacing. However, most printers produce circular spots. Their radius must be at least $T/\sqrt{2}$ so that they are capable of blackening a page entirely. Thus, black spots cover portions of adjacent spaces, causing the perceived gray level to be darker than the fraction of ones. We will refer to this fact as "dot overlap." Moreover, most printers produce black spots that are larger than the minimal size, which further distorts the perceived gray level. The most commonly used digital halftoning techniques protect against dot overlap by clustering black spots so the percentage effect on perceived gray level is reduced. Unfortunately, such clustering constrains the spatial resolution of the perceived images and increases the low-frequency artifacts. We propose methods that exploit printer distortions in order to increase both gray-scale and spatial resolution. A key element in such methods is an accurate printer model. We propose a general framework and develop some specific models for laser printers.

We used an HP Laserjet printer as the principal test vehicle. The printer is controlled by a binary array $[b_{i,j}]$, where $b_{i,j} = 1$ indicates that a black dot is to be placed at "site (i,j)" and $b_{i,j} = 0$ indicates that the site is to remain white. The gray level produced by the printer in the vicinity of site (i,j) depends in some complicated way on $b_{i,j}$ and neighboring bits. However, due to the close spacing of dots and the limited spatial resolution of the eye, the gray level can be modeled as having a constant value $p_{i,j}$ in this vicinity. Our printer model takes the form $p_{i,j} = P(W_{i,j})$, where $W_{i,j}$ denotes the window surrounding $b_{i,j}$. Thus, given the binary array $[b_{i,j}]$, our model generates a new array $[p_{i,j}]$ of gray levels which has the same dimensions as the binary array. This simplifies the processing task considerably.

For the methods considered here, it is essential that the window $W_{i,j}$ be finite. In this case, the possible values of P can be listed in a table. The elements of this table can be derived from a detailed physical understanding of the phenomena effecting gray level or from measurements of the gray level due to various dot patterns. A simple but very effective printer model considers the most elementary printer distortion: The fact that the dots are larger than the minimal covering size, as would occur if "ink spreading" occurred. This "dot-overlap" model can be described by the radius of the dots.

We propose two new halftoning methods that exploit the printer models. The first is a modification of error diffusion, which is considered to be the best halftoning technique for displays, such as CRT's, that do not suffer from substantial dot-overlap distortions. The new version exploits the printer model to extend the benefits of error diffusion to printers. Experiments show that it provides high quality reproductions with reasonable complexity. Error diffusion offers a substantial improvement over conventional clustered ordered dither in both spatial resolution and severity of low-frequency artifacts.

Like most halftoning schemes, error diffusion makes implicit use of the eye model. It shapes the noise so that it is not visible. It produces blue noise which has been found to be very pleasant to the eye. Error diffusion produces images with high spatial resolution, few low-frequency artifacts, and does not suffer from the false contouring problem. However, it is not entirely free of artifacts. For example, since it uses the eye model only implicitly, it allows the placement of single dots far apart so that they are very visible.

The second method, called least-squares model-based halftoning, exploits both a printer model and a model for visual perception. It produces an "optimal" halftoned reproduction, by finding the binary image that causes the cascade of printer and visual models to match (in the sense of minimizing squared error) the response of the visual model to the original gray-scale image. For one-dimensional halftoning, this can be implemented with the Viterbi algorithm. Experiments show that it successfully exploits the printer and visual models to produce more gray levels and better spatial resolution than conventional one-dimensional techniques. One-dimensional halftoning is seldom used in practice, however, because it does not exploit the two-dimensional characteristics of the eye and the printer.

Unfortunately, the Viterbi algorithm cannot be used in two dimensions. We therefore consider iterative methods to obtain the least squares solution. Experiments show that the two-dimensional least squares approach eliminates all the problems associated with error diffusion. More importantly, it establishes the limits of what can be achieved by a halftoning scheme given the characteristics of the eye and the printer.

Model-based halftoning can be especially useful in transmission of high quality documents using high fidelity gray-scale image encoders. In such cases halftoning is performed at the receiver, just before printing. Apart from coding efficiency, this approach permits the halftoner to be tuned to the individual printer, whose characteristics may vary considerably from those of other printers, for example, write-black vs. write-white laser printers.