Computer Graphics, Volume 25, Number 4, July 1991

Digital Halftoning with Space Filling Curves

Luiz Velho*
Jonas de Miranda Gomes

IMPA – Instituto de Matemática Pura e Aplicada
Estrada Dona Castorina, 110
22460, Rio de Janeiro, Brazil

ABSTRACT: This paper introduces a new digital halftoning technique that uses space filling curves to generate aperiodic patterns of clustered dots. This method allows the parameterization of the size of pixel clusters, which can vary in one pixel steps. The algorithm unifies, in this way, the dispersed and clustered-dot dithering techniques.

Keywords: digital halftoning, quantization, dithering, space filling curves, bilevel display.

1. INTRODUCTION

The display of gray scale images on bilevel graphic devices requires a preprocessing step in order to adapt the data to the characteristics of the equipment. In particular, a process called *halftoning* creates the illusion of continuous-tone through the careful arrangement of the state of individual display cells. This process can be analog or digital, depending upon the underlying technology of the imaging system. The analog form of halftoning is well understood, and has been used in the printing industry for more than one century. *Digital halftoning*, also known as *spatial dithering*, is associated with the computer display of pictures, and has been object of intensive research.

1.1 MOTIVATION

The initial motivation for the development of dithering techniques was the popularity of graphic display devices, such as plasma panels, liquid crystal and CRT monitors. More recently, the availability of high resolution hardcopy devices such as laser printers and digital phototypesetters created a new motivation for the development of digital halftoning techniques.

The majority of existing dithering algorithms were designed for a class of graphic display devices that have a relatively low spatial resolution and allow precise control of individual pixels. These algorithms perform poorly on some hardcopy devices that do not have these properties and cannot properly reproduce isolated dots. An important class of devices of this type is the popular laser printer, based on electrophotographic technology.

The work presented in this paper addresses this problem. We propose an algorithm that is flexible enough to be used in a wide range of graphic devices.

1.2 OVERVIEW

The organization of the paper is as follows: In Section 2 we describe the architecture of an imaging system for bilevel displays; in Section 3 we give an introduction to digital halftoning; a brief review of space filling curves is provided in Section 4. The main aspects of the clustered-dot dithering method using space filling curves are described in Section 5. Implementation details of our method are provided in Section 6. Examples of images generated by the method and comparisons with other dithering methods are presented in Section 7. Concluding remarks and perspectives of future work are discussed in Section 8.

2. IMAGING SYSTEM FOR BILEVEL DISPLAYS

The imaging system must perform several preprocessing operations in order to generate the proper representation of a continuous-tone picture on a specific graphic display device. This process must also take into account the particular characteristics of the device to produce the best possible rendition of the picture.

The device's characteristics can be modeled as a mathematical function, defined on the space of images, called the *physical reconstruction function*. The *preprocessing operations* generally include: *tone scale adjustment*, *sharpening* and *halftoning*. This pipeline is illustrated in Figure 1 ([Ulichney 87]).

The *tone scale adjustment*, also known as *gamma correction*, is necessary because most devices have a non-linear intensity reconstruction function. This operation compensates, for example, the overlapping of contiguous dots, typical of some hardcopy devices. Detailed explanation on how to construct compensation tables for CRT monitors can be found in the literature (see for example [Catmull 79]). This procedure can be generalized for other types of graphic devices.

The *sharpening* is desirable because the dithering normally causes some reduction of the image spatial resolution. The quality of the final image can be greatly improved by an edge enhancement operation that emphasizes high frequencies bringing out the fine image details. Alternatively, the sharpening operation can be incor-

---

*Author's current address: University of Toronto.

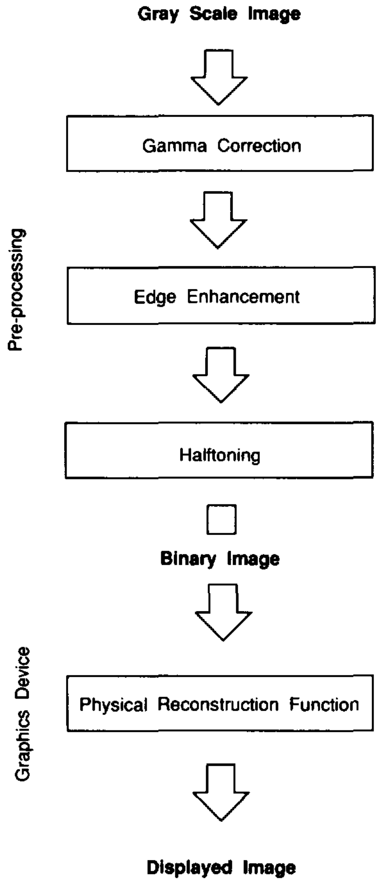

Figure 1 — Imaging pipeline for bilevel displays

porated into the halftoning process, as was observed by Jarvis [Jarvis et al 76].

3. HALFTONING

The existence of only two levels to be used in the display of continuous-tone images introduces visual artifacts, often manifested as false contours separating regions of different levels. Dithering alleviates this problem by properly controlling the distribution of bilevel intensities over the displayed image.

The dithering process is based on psychophysical characteristics of the human visual system. The eye integrates luminous stimuli over a solid angle of about 2 degrees [Wyszecki et al 82]. This means that we actually see the average intensities corresponding to small solid angles in our visual field. Dithering algorithms exploit this phenomenon, effectively redistributing the state of pixels in such a way that the average intensity in small areas of the dithered image is approximately the same of the original gray scale image.

Given a pixel P of the image with intensity $I(P)$, it will be mapped into a pixel P' of the dithered image whose intensity $I(P')$ is 0 or 1. The value of $I(P')$ is obtained by comparing the intensity $I(P)$ with a given intensity threshold I_0. The difference $I(P)-I(P')$ is the *quantization error* for the pixel P. In general, given a region of the image with N pixels, $P_1, P_2, \ldots, P_N$, $N+1$ intensity levels can be represented by turning these pixels "on" and "off". The quantization error for this region is the difference

$$\sum_{j=1}^{N} I(P_j) - \sum_{j=1}^{N} I(P_j'),$$

between the sum of intensities of the gray scale image in the region and the sum of the intensities of the corresponding region in the dithered image.

Dithering algorithms distribute the error over small neighborhoods of the image in such a way that the average quantization error is as close to zero as possible. There are two main strategies to define the states of the pixels on the dithered image in order to achieve this goal. One of them perturbs the intensity threshold I_0 in a predefined way, so that the error is statistically neglectable; the other strategy perturbs the threshold for a pixel P based on the quantization error in a neighborhood of P, obtaining an exact minimization of the error. In both cases, the perceived intensity of the dithered image at a given neighborhood will be close to that of the original image.

This technique implies in a trade-off between spatial and tonal resolution: as we spread the error over larger areas of the image, more tones can be represented at the cost of a poorer rendition of fine details. The gray levels are rendered as patterns of black and white pixels eliminating high frequency information. In this process contouring artifacts are transformed into patterning features.

3.1 DITHERING TECHNIQUES

Spatial dithering techniques can be classified according to the nature of patterns they generate and to the type of pixel configuration they produce. These two criteria capture the main features of the textures created to represent areas of uniform gray, one of the most important aspects of the halftoning process.

Textures can be rendered by *periodic* or *aperiodic* patterns. In general, periodic patterns are generated by deterministic processes based on regular sampling grids. Aperiodic patterns are generally associated with methods that can be modeled as stochastic processes.

The type of pixel configuration produced is determined by the spatial distribution of the "on" or "off" state of the image elements. *Dispersed-dot* methods depict a gray level by covering a small area with evenly distributed dots, while *clustered-dot* methods concentrate the dots in small groups.

3.2 PREVIOUS WORK

The most popular halftoning method is the ordered dither technique. It uses a deterministic perturbation to generate periodic patterns, and according to the distribution of perturbations it can produce dispersed or clustered dots. Other important methods are the error diffusion techniques. The well known algorithms in this category are the Floyd-Steinberg, and Knuth's dot-diffusion algorithm. They generate aperiodic patterns as the result of neighborhood operations. All published error diffusion algorithms fall into the dispersed-dot category.

The *ordered dither* algorithm determines a matrix of quantization thresholds that is replicated over the image. This is essentially a set of pseudo-random numbers uniformly distributed over the intensity range. The arrangement of thresholds is designed to avoid the introduction of low spatial frequency noise into the image. This algorithm is generally identified as a dispersed-dot technique [Limb 69], but if the intensity threshold levels are spatially concentrated it results in a clustered-dot dithering.

The *Floyd-Steinberg* algorithm [Floyd et al 75] computes the quantization error incurred in one image element and propagates it to the neighbors to the right and below. In this way, the local quantization error is distributed, minimizing globally the intensity difference between the original and quantized images.

The *dot diffusion* algorithm [Knuth 87] combines some characteristics of ordered dither and error diffusion techniques. Similarly to ordered dither it uses a matrix that is replicated over the entire image. This matrix gives the order by which the quantization error in one display cell will be distributed among its neighbors in the cell.

A comparison between dithering algorithms can be found in the survey [Jarvis et al 76]. A comprehensive study of dithering techniques with an analysis of the statistical properties can be found in [Ulichney 87].

We propose a digital halftoning method based on space filling curves, which uses the path of the curve to distribute the quantization error over the image. Witten and Neal [Witten et al 82] also described a dispersed-dot dithering algorithm that propagates the quantization error along a Peano curve.

Our technique parameterizes the dot aggregation factor allowing a precise control of the cluster size, which can vary in one pixel steps. This is the first algorithm that effectively unifies the dispersed and clustered-dot techniques. When the cluster size is one pixel it reduces to a dispersed-dot dithering using error diffusion. Therefore, Witten and Neal's algorithm is a particular case of our method.

As mentioned before, a large class of hardcopy devices cannot reproduce well configurations of sparse "on" and "off" pixels. For this reason, most page description languages employ clustered-dot ordered dithering, as the standard halftoning method [Adobe 85]. The method presented in this paper offers an alternative solution to the halftoning problem. It works very effectively in graphic displays as well as in hardcopy devices, and has potential applications in higher resolution printing.

4. SPACE FILLING CURVES

A continuous *plane curve* is a continuous map $c{:}I \rightarrow \mathbf{R}^2$ from the unit interval $I = [0,1]$ of the real line to the two-dimensional euclidean plane $\mathbf{R}^2 = \{(x, y)\,;\, x, y \in \mathbf{R}\}$. The image $c(I)$ is called the *trace* of the curve c. A *space filling curve* is a continuous curve such that its trace covers the unit square $I^2 = [0,1] \times [0,1]$ of the plane. Therefore, for each point P in the square I^2 there exists a real number t in the interval I such that $c(t) = P$. Intuitively, this means that the curve provides an ordered way to visit all points of the square as the parameter t moves from 0 to 1.

Space filling curves were first discovered by the Italian mathematician Giuseppe Peano in 1890, and they constitute the first examples of the mathematical objects that Benoit Mandelbrot called fractal sets [Mandelbrot 77].

The mathematical construction of a space filling curve c is done as a limiting process. We consider a sequence $c_n{:}I \rightarrow I^2$ of curves in the unit square, and we define c as the limit

$$c = \lim_{n \to \infty} c_n$$

when this limit exists. The curves c_n constitute approximations of c, and as we increase n it visits a greater number of points in the unit square. It is possible to construct space filling curves for which each curve c_n is simple, i.e. the map is 1–1. This means that it does not visit a point in the square more than once. In general it is possible to construct the sequence $c_1, c_2, \ldots, c_n, \ldots$ of approximating curves in a recursive way. In a certain sense a space filling curve defines a relationship between the area of subregions of the unit square I^2 and the length of subintervals of the unit interval I.

4.1 COMPUTATIONAL METHODS

Space filling curves can be properly specified by a formal geometric language. Sentences in this language are defined by a parallel graph grammar, and they are constructed by recursively applying a set of rewriting rules. Each sentence corresponds to a curve c_n from the approximating sequence of the space filling curve. We will refer sometimes to this approximation itself as a space filling curve. A discussion about computational methods to generate space filling curves can be found in [Prusinkiewicz 90].

4.2 CLASSIC CURVES

The classic space filling curves are the *Peano curve*, the *Hilbert curve*, and the *Sierpinsky curve*. Figure 2(a)(b)(c) shows an approximation of these curves. All curves in the approximating sequence of these curves are simple.

4.3 IMAGE SCAN

When each curve c_n in the approximating sequence of a space filling curve is simple, we obtain a method to visit, in a unique and ordered way, a subset of points of the square. The number of points visited increases as we increase the value of n. If we consider the square grid defined by the pixels of a raster image it is possible to address uniquely all pixels using a simple approximating curve c_n of a space filling curve. Therefore, these curves constitute an effective method to scan a raster image. This idea has been exploited in the field of Digital Image Processing [Koo-Yan-Too 88], [Stevens et al 83].

The scan method described above has several advantages over the traditional scanline method for some class of image operations. The recursive nature of the construction of space filling curves allow a subdivision of the image into regions where each region is mapped to some subinterval of the unit interval I. This implies in a certain sense a reduction of the dimensionality of the problem, and simplifies immensely algorithms that deal with small regions of the image, as well as the computations involved.

The path followed by the space filling curve results in an image scan free of directional features presented by the traditional scanline raster pattern.

5. APERIODIC CLUSTERED-DOT DITHERING

The digital halftoning method using space filling curves exploits the properties of these mathematical objects to perform neighborhood operations essential to the spatial dithering process. This section presents the overall structure of the method and describes in detail its main aspects.

5.1 THE METHOD

The method consists of the following steps:

- Subdivision of the source image into small regions based on the trace of the space filling curve;
- Computation of the average intensities of each region;
- Determination of the dot patterns of the dithered image corresponding to each intensity;

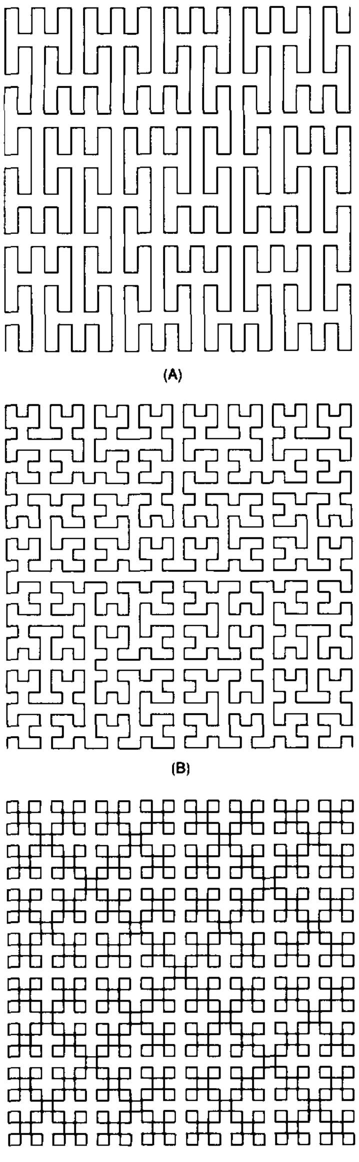

Figure 2 — Approximations of: (a) Peano, (b) Hilbert, (c) Sierpinski space filling curves.

5.2 IMAGE SUBDIVISION

The method takes advantage of some properties of space filling curves that allow a subdivision of a raster image into regions with desirable characteristics. Let $c_n{:}I \to I^2$ be an approximation of a space filling curve c that visits uniquely all pixels of the image. Let $I_1, I_2, ..., I_n$ be a subdivision of the unit interval I into n subintervals. By restricting the curve c_n to each subinterval I_j we obtain n subregions $R_1, R_2, ..., R_n$ of the image. The size of each region R_j varies proportionally with the length of the corresponding subinterval I_j. This gives an ordered way to visit all regions R_j, and also to visit all points in each of these regions. Besides this, the restriction $c_j{:}I_j \to R_j$ is by itself a space filling curve, that is a scaled version of the original curve c, because of the self-similarity properties of the space filling curves. This characteristic minimizes the grid effect often manifested in dithering methods that use standard methods of image scan.

5.3 DOT GENERATION

The dot generation strategy is a direct consequence of scanning the image with a space filling curve. The objective is to produce, for a given region, a configuration of clustered dots that will result in a perception equivalent to the intensity of the original image. This depends on the area of the region, the average intensity over the region, and the graphic device's physical reconstruction function.

As described above, the trace of the space filling curve determines a relationship between the area of the region and the length of the curve. Suppose that the average intensity of a region R is I. Ideally, the desirable perceptual results would be obtained by partitioning $R = R_1 \cup R_2$ into two subregions R_1 of white pixels, and R_2 of black pixels, such that R_1 corresponds to a subinterval of length proportional to I and R_2 corresponds to a subinterval of length proportional to $1-I$. In practice, this subdivision cannot be done exactly because there is a discretization process involved that is influenced by the physical characteristics of the output device.

The graphics output device is able to display only a discrete number of fixed size dots at a determined resolution. In general, the shape of the dot is not completely regular, and there is some overlapping between contiguous dots. This fact implies in a degree of non-linearity in the reconstruction function. As mentioned in Section 2, it is possible to account for the device's non-linear response by means of an independent preprocessing step.

The dot configuration produced by the space filling curve method results in an aggregate of pixels connected not only sequentially by the curve, but also in other directions because of the intertwined way the space filling curve traces the region. Consequently, the cluster of dots obtained is confined within the limits of a ball that has an area close to the area of the region. As a whole, the patterns generated by this type of dots are evenly distributed but not periodic.

In order to account for the fine details of the image, it is desirable that the dot configuration grows outwards from the point of highest intensity of the region. This can be accomplished by centering the white subregion with a proper translation of the corresponding subinterval.

Figure 3 illustrates clusters of dots corresponding to intensities 15/16 to 0 for the Hilbert curve, in a region of 4×4 pixels. In Figure 4 we used the method to render a black to white gradation using different sizes for the dot aggregation.

5.4 ERROR DIFFUSION

The discrete nature of the reproduction process, as we have seen, may result in quantization errors. This error can be propagated along the path of the space filling curve in order to minimize the total quantization error. This is similar to the dispersed-dot error diffusion dithering techniques, but works on display cells of more than one pixel.

6. IMPLEMENTATION

The halftoning method presented in this paper was developed under the VISGRAF project, as part of an image processing system in the Computer Graphics laboratory at IMPA.

The computing environment is integrated by a network of Sun workstations and the primary graphics hardcopy devices are 300 dpi Postscript laser printers.

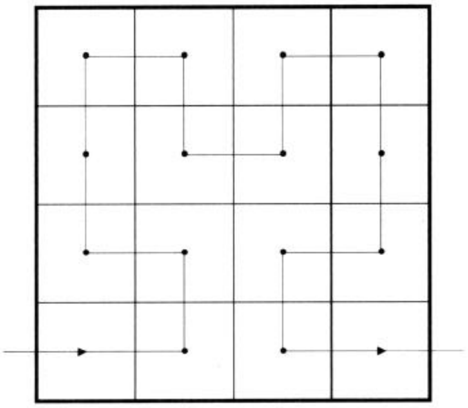

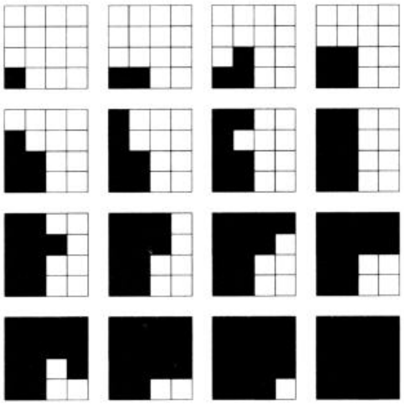

Figure 3 — Configuration of dots corresponding to intensity levels 15/16 to 0, for a cluster of 16 pixels using the Hilbert space filling curve.

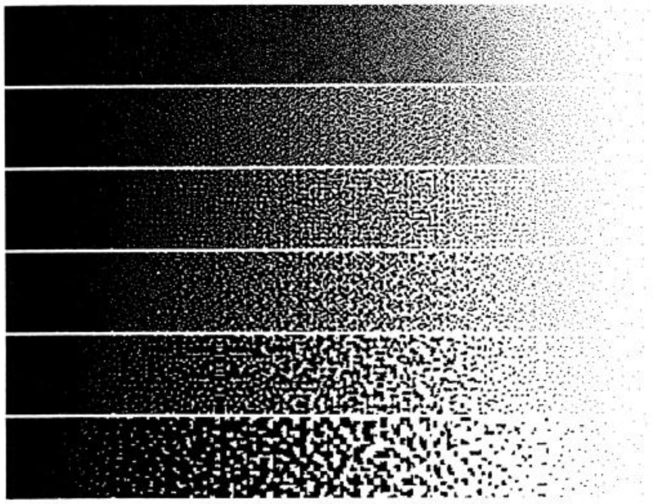

Figure 4 — Stripes with gradation dithered with the space filling curve algorithm (Hilbert curve) using different cluster sizes. From top to bottom, clusters of 2, 6, 12, 20, 32, 60 and 120 pixels.

The algorithm was implemented using the C language in the Unix operating system.

6.1 SCAN LIBRARY

The image scan pattern generation is implemented by a library of functions with a common interface. This simplifies the addition of new types of space filling curves to the dithering operation, and encourages experimentation.

The library's front-end consists of two functions. The first one selects the curve to be used for the image scan and, if necessary, executes initialization and setup procedures. The second function moves forward and backwards along the path incrementally returning the coordinates of image points to be visited. It should be called once for each element processed.

6.2 ALGORITHM

The pseudo-code below gives a description of the basic algorithm. R is the maximum pixel intensity (255 for images with 8 bits of resolution), and N is the cluster size in pixels.

```
Select image scan curve ;
Initialize intensity accumulator ;

While (image elements to be processed ) {
        Advance image pointer along the scan
        path to the end of interval ;

        Move backward N pixels, accumulating
        the intensity of the input image ;

        Move forward N pixels along the path,
        setting the output pixels :
        if ( accumulator = R ) then {
                decrement R from accumulator ;
                set output pixel "on" ;
        } else {
                set output pixel "off" ;
        }
}
```

Note that the algorithm implicitly accounts for the quantization error, propagating it along the path.

The processing structure of the algorithm allows the same buffer to be used for both input and output image.

7. RESULTS

Although the method works well in low resolution devices, the clustered-dot dithering using space filling curves is primarily intended for medium to high resolution bilevel devices that cannot accommodate isolated black or white pixels. For this reason, the tests of the method were performed using a 300 dpi laser printer as the graphics output device.

7.1 EXAMPLES

Two different images were chosen as representatives of the common types of pictures in graphics applications. The first image, Figure 5, was captured from a black and white photographic reproduction of a study for the mural painting, "Escola dos Jesuitas", by the Brazilian artist Candido Portinari. This drawing of an indian boy head was done using charcoal, red ocher and sepia on paper, and dates from 1938. The image was digitized using a 300 dpi, 8 bits gray scale scanner. The second image is a

Figure 5 — Digitized test image: A drawing by the Brazilian artist Candido Portinari (1938).

computer generated image designed to include a wide range of features. It consists of a circular gradation inside a disc over a background with horizontal bands. Both images contain areas of smooth intensity variation as well as areas of high contrast and fine detail.

In the preprocessing step, only tone scale adjustment was performed prior to the halftoning operation. We decided not to do any edge enhancement in order to have a better feeling on how the algorithm handles fine details.

Figures 6 and 7 illustrate the clustered-dot dithering algorithm using Hilbert's space filling curve. The clustering size was of 11 pixels. Before dithering the two images were scaled down to 150 dpi. By increasing the viewing distance we can simulate the behaviour of the algorithm in higher resolution.

Figure 8 (A), (B) and (C) shows halftoned versions of the two images processed respectively by the space filling curve, the Floyd-Steinberg and the clustered-dot ordered dither algorithms. They were included to compare the results of the new method with both a standard error-diffusion technique and with the clustered-dot method used in most hardcopy devices. For the last comparison we used a 8×8 matrix in the clustered-dot ordered dither and a cluster size of 32 pixels in the space filling curve dither. These choices produce clusters of approximately the same size. Before dithering the two images were scaled down to 75 dpi.

7.2 ANALYSIS

The space filling curve dithering algorithm generates aperiodic patterns of evenly distributed dots without directional artifacts. It renders well the gray levels, and captures the fine details. These

Figure 6 — Indian boy head at 150 dpi dithered with the space filling curve algorithm (Hilbert curve), using clusters of 11 pixels.

features are evident in both images, in particular in the face, eye and hair of the indian boy.

The Floyd-Steinberg algorithm, as was expected, did not produce satisfactory results on the laser printer. This is because the dispersed-dot method is not appropriate for this device. Groups of nearby individual small dots tend to be merged into a large blob. For this reason, the gray levels are not reproduced well, increasing the image contrast.

The clustered-dot ordered dither algorithm simulates the traditional analog halftoning screen. It reproduces very well the gray levels, but it blurs slightly the image. Depending on the cluster size contouring is more or less noticeable.

8. CONCLUSIONS

In this paper we introduced a new digital halftoning technique. The dithering method described is based on the trace of space filling curves to generate aperiodic patterns of clustered dots. To our knowledge this is the first algorithm that exploits random looking patterns of pixel agglomerates as a solution to the problem. Furthermore, the algorithm encompasses the dispersed dot error diffusion technique [Witten et al 82] as a particular case.

8.1 DISCUSSION

The space filling curve method has several advantages over previous ones. It generates patterns without the regular structure of the clustered-dot ordered dither. The patterns created are perceptually pleasant with similar characteristics to the photographic grain structure. The clustering factor can be easily parameterized,

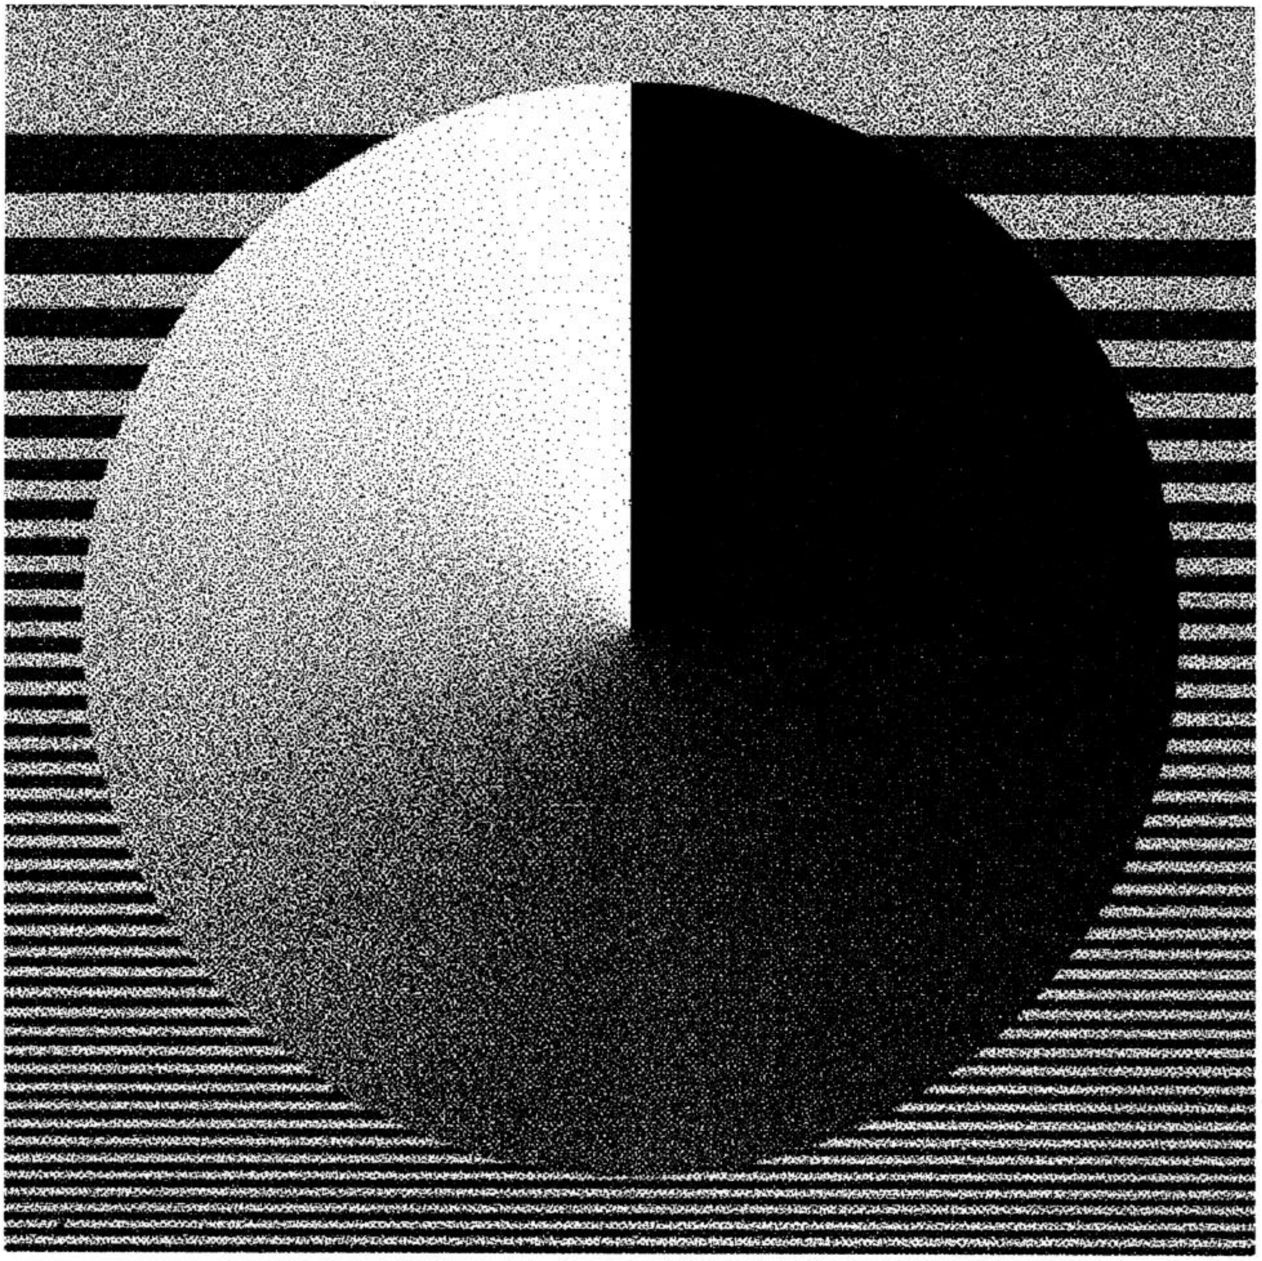

Figure 7 — Computer generated picture at 150 dpi dithered with the space filling curve algorithm (Hilbert curve), using clusters of 11 pixels.

allowing the image rendition to match precisely the limits of the physical reconstruction function of the display device. The algorithm is computationally efficient requiring only 1 addition, 1 subtraction and 1 comparison per image element processed.

The main drawback of the algorithm is its high memory requirement, since it buffers the entire image because of its non-standard access pattern. This is probably not a serious restriction, except for very high resolution images. In this case, the problem can be addressed in two ways: the image can be subdivided in small blocks, and the algorithm is performed more or less independently in each one. This requires buffering of small strips of the image. Another solution is to store the image in a non-standard way such that its structure favors the access pattern. This is discussed by Blinn in the context of texture mapping [Blinn 90].

One inherent limitation of the method is that it is not truly bidimensional. For this reason, the error propagation is not totally uniform. This weakness is shared to some extent with all the published dithering techniques. The error diffusion can be cast as an equilibrium problem, which can be solved by relaxation techniques, such as simulated anealling [Kirkpatrick et al, 1982], [Fiume 89]. The computational effort required for an accurate solution is very expensive and has not yet been tried for this type of application.

8.2 FUTURE RESEARCH

Future work includes the extension of the method to process full color images, experiments with higher resolution graphics devices and the investigation of adaptive clustering techniques.

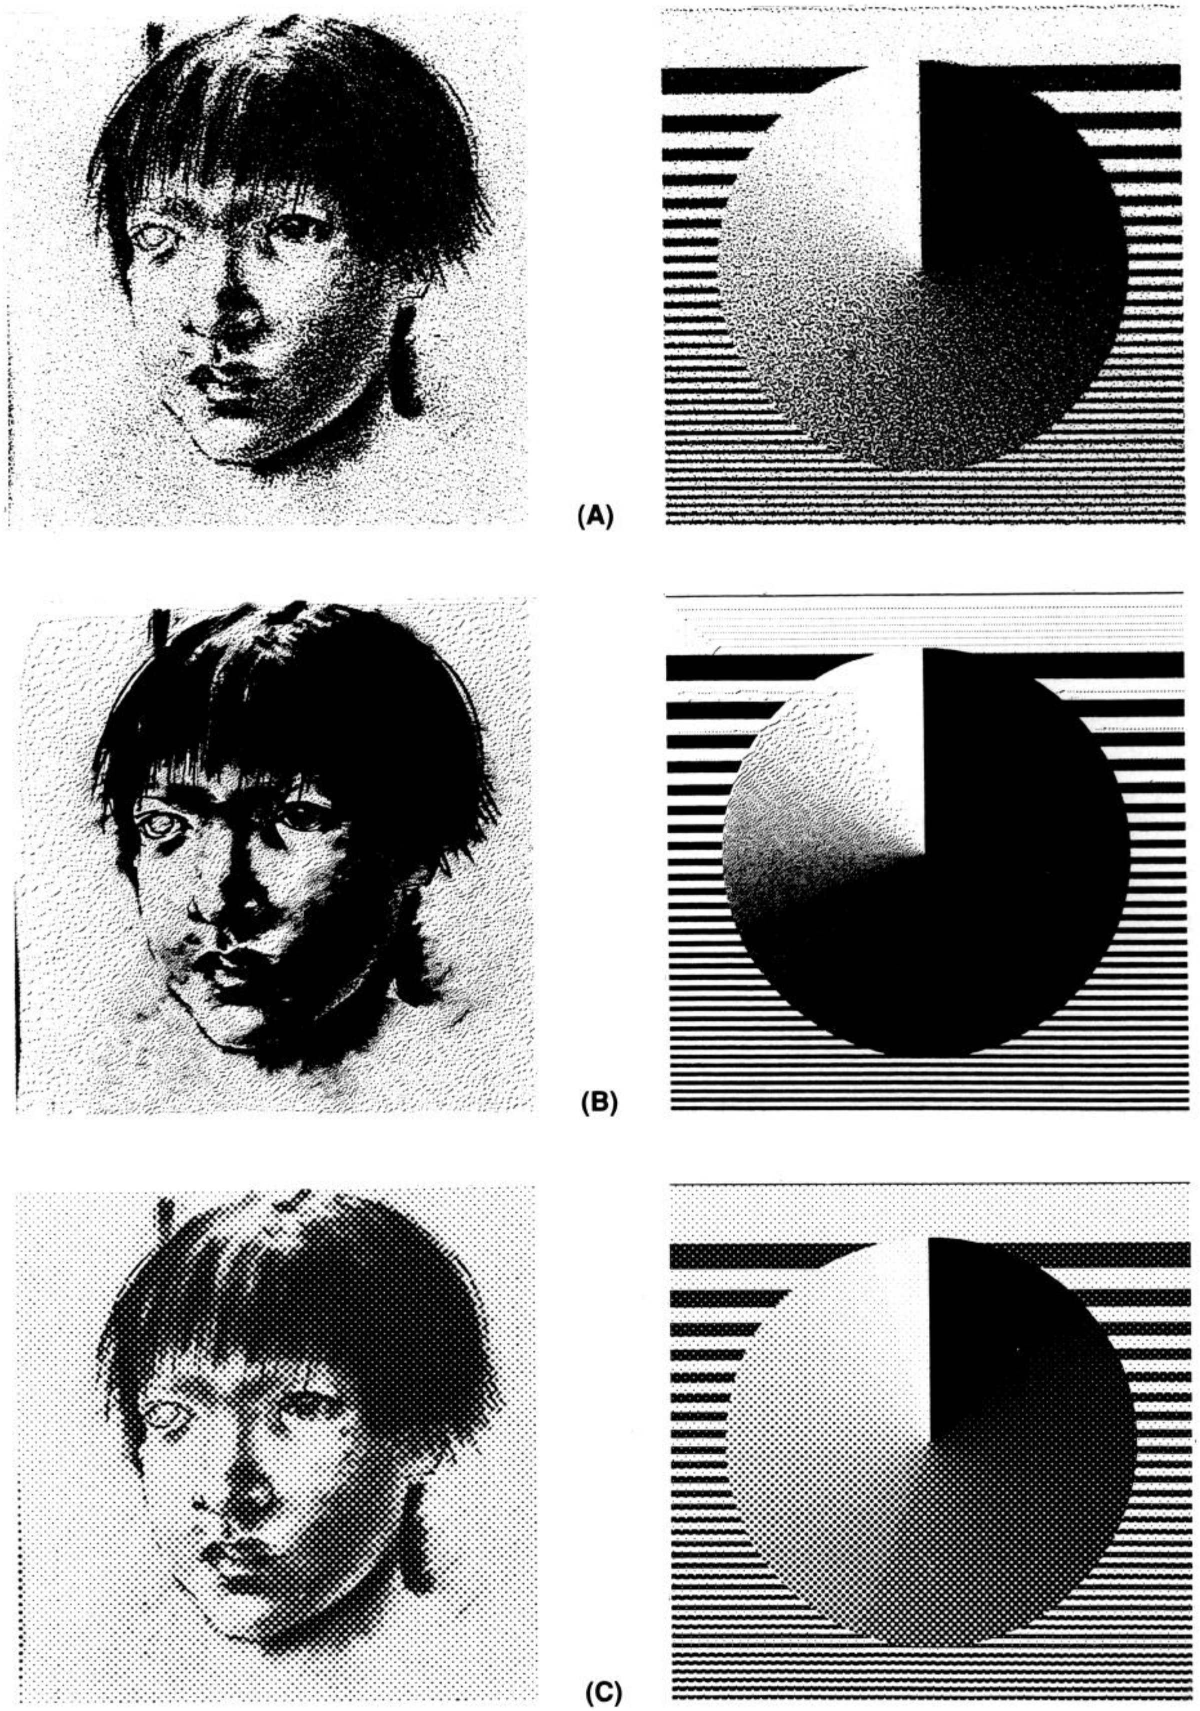

Figure 8 — The two test images at 75 dpi dithered with three different algorithms: (A) Space filling curve algorithm (Hilbert curve), using clusters of 32 pixels; (B) Floyd-Steinberg algorithm; (C) Clustered-dot ordered dither algorithm, using a matrix of order 8.

The method has also a potential to be used for illustration purposes. Other kinds of rendering effects can be obtained by a combination with image processing techniques. An example of this process, simulating pen-and-ink drawing, is shown in Figure 9.

9. ACKNOWLEDGEMENTS

The authors would like to thank João Candido Portinari, and Portinari Project's team for kindly providing the photographic reproduction of the study for the indian boy head. We also appreciate the helpful and encouraging comments provided by the reviewers.

10. REFERENCES

Adobe Systems, (1985): *Postscript Language Reference Manual.* Addison-Wesley, Reading Massachusetts.

Blinn, J. (1990): *The Truth About Texture Mapping.* IEEE Computer Graphics and Applications, March 1990, 78-83.

Catmull, E. (1979): *A Tutorial on Compensation Tables.* Proceedings SIGGRAPH '79, in Computer Graphics, Vol. 13, 1-7.

Fiume, E. and Ouellette, M. (1989): *On Distributed Probabilistic Algorithms for Computer Graphics.* Proceedings of Graphics Interface 89, 211-218.

Floyd, R. and Steinberg, L., (1975): *An Adaptive Algorithm for Spatial Gray Scale.* SID Symposium, 1975, 36-37.

Geist, R. and Reynolds, R. (1990): *Colored Noise Inversion in Digital Halftoning.* Proceedings of Graphics Interface 90, 31-38.

Jarvis, J., Judice, C. and Ninke, W., (1976): *A Survey of Techniques for The Display of Continuous Tone Pictures on Bilevel Displays.* Computer Graphics and Image Processing, n. 5, 13-40.

Kirkpatrick, S., C. D. Gelatt Jr., and M. P. Vecchi, (1982): *Optimization by Simulated Annealing*, IBM Research Report RC 9355.

Knuth, D., (1987): *Digital Halftones by Dot Diffusion.* ACM Transactions on Graphics, V. 6 N. 4, Oct 1987, 245-273

Koo-Yan-Too, H. C., (1988): *A Peano Scan Approach to Multivariate Data Clustering, with an Application.* Master Thesis, Dept. C. S., Univ. of Regina.

Limb, J. O., (1969): *Design of Dither Waveforms for Quantized Visual Signals.* Bell Systems Technical Journal, v. 48, n. 7, 2555-2582.

Mandelbrot, B., (1977): *The Fractal Geometry of Nature.* W. H. Freeman, New York

Prusinkiewicz P., and A. Lindenmayer (1990): *The Algorithmic Beauty of Plants.* Springer-Verlag, New York.

Sonnenberg, H., (1983): *Designing Scanners for Laser Printers.* Lasers & Applications, April 1983, 67-70.

Stevens, R. J., Lehar, F. A. and Perston, F. H. (1983): *Manipulation and Preservation of Multidimensional Image Data using the Peano Scan.* IEEE Trans. on Pattern Analysis and Machine Intelligence, 5, 520-526.

Ulichney, R., (1987): *Digital Halftoning.* MIT Press, Cambridge Massachusetts.

Witten, I. H., and Neal, M., (1982): *Using Peano Curves for Bilevel Display of Continuous Tone Images.* IEEE Computer Graphics and Applications, May 1982, 47-52.

Wyszecki, G. and Stiles, W. (1982): *Color Science: Concepts and Methods, Quantitative Data and Formulae.* Second edition. John Wiley & Sons.

Figure 9 — A pen-and-ink drawing effect obtained using image processing and the space filling curve dithering.