

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

**Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600**

UMI[®]

University of Alberta

TEXTURE CONTROL IN DIGITAL HALFTONING

by

Oleg Veryovka ©

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Department of Computing Science

**Edmonton, Alberta
Fall 1999**



**National Library
of Canada**

**Acquisitions and
Bibliographic Services**

395 Wellington Street
Ottawa ON K1A 0N4
Canada

**Bibliothèque nationale
du Canada**

**Acquisitions et
services bibliographiques**

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-46940-9

Canada

University of Alberta

Library Release Form

Name of Author: Oleg Veryovka

Title of Thesis: Texture Control in Digital Halftoning

Degree: Doctor of Philosophy

Year this Degree Granted: 1999

Permission is hereby granted to the University of Alberta Library to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only.

The author reserves all other publication and other rights in association with the copyright in the thesis, and except as hereinbefore provided, neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatever without the author's prior written permission.

Oleg Veryovka

Oleg Veryovka
309 Michener Park
Edmonton, Alberta
Canada T6H 4M5

Date: *May 12, 1999*

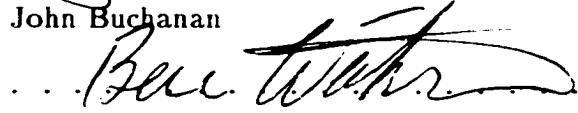
University of Alberta

Faculty of Graduate Studies and Research

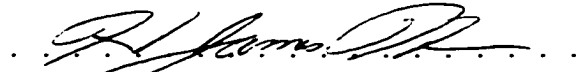
The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled **Texture Control in Digital Halftoning** submitted by Oleg Veryovka in partial fulfillment of the requirements for the degree of **Doctor of Philosophy**



John Buchanan



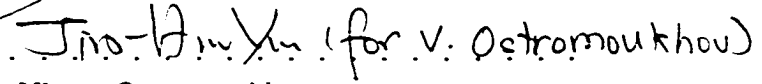
Ben Watson



Jim Hoover



Desmond Rochfort



Victor Ostromoukhov

Date: *May 12, 1999*

Abstract

There are two ways to represent visual information: photographic and artistic. Photographic approaches attempt to approximate an image despite the limitations of the output medium. Traditional halftoning takes a photorealistic approach. In an artistic rendering visual information is interpreted by the artist and displayed accordingly using the chosen medium. The non-photorealistic rendering area of computer graphics develops tools and techniques to enable interpretive rendering in digital media.

Texture is an inevitable artifact of halftoning. The challenge of photorealistic halftoning is to preserve image features — tone, edges, and textures — and to hide the halftoning texture. To the contrary, in artistic rendering texture is often used as a visual cue and an expressive mean. In this thesis I explore the use of halftoning texture to enhance the representation of visual information in both photorealistic and interpretive rendering applications. However, this use of texture requires methods and techniques to control texture in halftoned images. Thus, the objective of this work is to control the appearance of texture in the resulting images.

My technique of texture control is based on previous halftoning algorithms: ordered dither and error diffusion. I use the ability of the ordered dither algorithm to define halftoning texture through the arrangement of threshold values in its dither matrix. The thesis describes two methods of generating dither matrices: image processing and procedural texturing. The use of texture based dither matrices guarantees the appearance of desired textures in the halftoned image. The strength of the resulting texture is controlled by combining ordered dithering with the error diffusion process.

The ability to define and control texture in the halftoned image leads to the use of this texture as an expressive mean in image rendering. A user may introduce a variety

of artistic effects into the image. Examples include embossing an image with a texture or text; approximation of traditional art styles and rendering techniques — pencil drawing, carving, oil brush painting. The thesis also includes techniques that allow us to map texture features to enhance representation of image gradient, 3-D scene information and subjective user defined information.

This study is a contribution to both photorealistic and artistic halftoning of images. It is a new approach to non-photorealistic rendering. Unlike previous techniques, interpretive halftoning is not limited to any particular style of rendering. Moreover, artistic effects generated by previous techniques may be approximated.

Investigation of texture control in halftoning extends photorealistic dithering techniques. It turns out that the use of image based dither matrix improves rendering of the original image textures and edges. Also, the thesis includes investigation of an image quality measure that allows us to analyze halftoned images. This measure is based on multi-scale analysis of image edges and thus enables us to quantify edge distortions introduced by the halftoning algorithm.

Acknowledgements

I thank everybody who inspired me to do this research,
who made it possible and enjoyable:

my parents,

my family,

my friends,

my instructors and support staff at University of Alberta.

Table of Contents

1	Introduction	1
1.1	Motivation and context	1
1.2	Research goals	2
1.3	Outline of the approach.	3
1.4	Overview of the thesis	5
2	Fundamentals of tone reproduction	7
2.1	Properties of the human visual system	7
2.2	Conventional halftoning approaches	9
2.2.1	Threshold matrix dithering	10
2.2.2	Error diffusion	12
2.2.3	Optimization approaches to halftoning	14
2.3	Numerical measures of image quality	15
2.4	Summary of photorealistic halftoning	16
3	Expressive and interpretive rendering	18
3.1	Two approaches to expression of reality	18
3.2	The main styles of non-photorealistic rendering	19
3.2.1	Modeling of traditional art materials	20
3.2.2	Computer graphics as a medium in itself.	20
3.2.3	Comprehensive rendering	20
3.3	Expressive tone reproduction	21
3.3.1	Artistic styles of tone reproduction	22
3.3.2	Extensions of halftoning to expressive rendering	22

3.3.3	Comprehensive tone reproduction	23
3.4	Summary of expressive rendering	24
4	Fundamentals of Texture Control in Halftoning	25
4.1	Generalized ordered dithering	26
4.2	Conversion of an arbitrary image into a dither matrix	26
4.3	Procedural dither textures.	31
4.3.1	Generalization of essential properties of a dither matrix	31
4.3.2	Procedural dither kernels	32
4.3.3	Hybrid dither textures	34
4.4	Control of texture contrast	37
4.5	Discussion	39
5	Texture Based Interpretive Halftoning	42
5.1	Application of texture to interpretive halftoning.	42
5.2	Image based control of texture features	44
5.2.1	Texture tone and contrast	45
5.2.2	Texture shape	46
5.2.3	Texture direction and scale	47
5.2.4	Summary of the image based texture control	50
5.3	Interpretive halftoning of two-dimensional images	50
5.4	Interpretive Halftoning of 3-D Scenes	55
5.4.1	Scene based approach to the halftoning of 3-D scenes	55
5.4.2	Image based approach to the halftoning of 3-D scenes	59
5.4.3	Discussion of the 3-D scene based halftoning	64
5.5	Discussion of the interpretive halftoning	65
6	Texture Control in Photorealistic Halftoning	68
6.1	Preservation of image details in halftoned images	68
6.1.1	Error diffusion with various thresholds	69
6.1.2	Applications to the non-clustered and clustered halftoning	71

6.2	Multiscale analysis of halftoned images	74
6.2.1	Multiscale analysis of edges.	75
6.2.2	Evaluation of edge artifacts.	78
6.2.3	Analysis of edge distortions in halftoning.	81
6.2.4	Summary	83
6.3	Quantative evaluation of texture control techniques	83
6.3.1	Evaluation of the AHE parameters	84
6.3.2	Evaluation of the error diffusion parameters	86
6.3.3	Evaluation of photorealistic halftoning results	86
6.3.4	Summary of the quantative results	89
6.4	Discussion of texture control in photorealistic halftoning	90
7	Conclusion	92
7.1	Contributions to photorealistic halftoning	93
7.2	Contributions to non-photorealistic rendering	93
7.3	Limitations of the technique	94
	Bibliography	96

List of Tables

6.1	Strength of edge artifacts across multiple scales	80
6.2	Edge distortions of the photograph “Lena”	88

List of Figures

2.1	Principles of ordered dithering	10
2.2	Examples of clustered and dispersed ordered dithering	11
2.3	Examples of error diffusion with different approximation regions	14
4.1	Halftoning images with textures as a dither matrix.	28
4.2	Generating dithering texture by clipped adaptive histogram	30
4.3	Technique textures with statistical properties of dither screens	33
4.4	Procedural dither kernels with texture displacements	34
4.5	Hybrid textures: functions and applications	36
4.6	Influence of error diffusion on halftoning textures	38
4.7	Embossing images with textures	40
5.1	Local control of texture appearance	45
5.2	Image based control of texture shape	46
5.3	Ambiguity of texture scale control	47
5.4	Local control of texture scale	48
5.5	Local control of texture direction	49
5.6	Interpretive halftoning of the “Lena” image. Part 1	51
5.7	Interpretive halftoning of the “Lena” image. Part 2	53
5.8	Interpretive halftoning of the “Lena” image. Part 3	54
5.9	2D texture mapping of dither matrices	56
5.10	Halftoning with 3-D textures	58
5.11	Image based 3-D halftoning using id-buffer and n-buffer	60
5.12	Image based 3-D halftoning using l-buffer and z-buffer	62

5.13	Image based 3-D halftoning with derivative of the z-buffer	63
5.14	3-D halftoning with the importance map and <i>id</i> and <i>n</i> buffers	64
5.15	Interpretive halftoning of the Beethoven bust	66
5.16	Comparison between the 2-D and the 3-D halftoning	67
6.1	The effect of error diffusion on halftoning with various dither screens. . . .	70
6.2	Reproduction of edges in non-clustered halftoning	72
6.3	Clustered halftoning with textured enhancement	73
6.4	Detection of multiresolution edges	77
6.5	Detection of edge artifacts produced by popular halftoning algorithms . . .	79
6.6	Edge distortions produced by clustered dot algorithms	82
6.7	Evaluation of the AHE parameters	85
6.8	The influence of error diffusion on the halftoning texture	87

Chapter 1

Introduction

The beauty of the world around us, our thoughts and designs to be shared and documented must be put into images. These images are representations of visual information rendered despite the limited means of the chosen medium. This thesis deals with rendering of images on binary output devices such as black and white monitors or printers. The objective of this research is to provide tools to control texture in binary images, thus preserving image details and enhancing display with artistic or illustrative elements.

1.1 Motivation and context

There are two approaches to expression of visual information: photographic and artistic. While both approaches share the same goal they treat displayed information in a different manner.

The goal of the photographic approach is to copy visual information and to approximate it in the target medium. One can take a photograph of a scene to record visual information automatically. Current computer graphics techniques allow us to create photorealistic images through modeling of objects and their materials. The resulting images must be halftoned into appropriate representation for printing on conventional output hardware. These halftoning algorithms create the illusion of multiple shades by a combination of black and white dots. The arrangement of dots forms an artificial texture in the resulting halftoned image and is called the *halftoning layer*. Thus, texturing of images is the inherent artifact of halftoning dictated by the current limitations of the output hardware. The halftoning layer may interfere with display of image tones, edges, and textures. Hence, the challenge of photorealistic tone reproduction is to make the halftoning layer less visible.

There is a fundamental difference between photorealistic and artistic rendering:

a photograph is a trace or a print of reality, whereas painting - even one meeting

photographic standards of execution - can only be an interpretation. [LS95]

Thus, artistic approach to information display is based on the artist's interpretation. The photographic accuracy can be sacrificed for the clarity of this display. As a result, not only all the information is accounted for in the final image, but also this information is interpreted and filtered by the artist.

The challenge of artistic rendering is to select the most essential information and to find the best possible representation of it in the chosen medium. The tone reproduction layer — e.g. pencil strokes — is not only used for approximation of continuous tone image, it is also an additional channel of visual information. Thus, tone reproduction layer is an expressive medium that is carefully manipulated by an artist to enhance image display.

To summarize, the artistic approach differs from the conventional halftoning approaches in the treatment of graphics primitives used to create the illusion of shades and tones. In artistic rendering, pencil and brush strokes are used to enhance information display, while the goal of conventional tone reproduction is to hide the halftoning layer.

In both photorealistic and artistic image display graphics primitives form a texture layer. The goal of my investigation is to control this texture and to use it for the enhancement of image display. In the case of artistic halftoning the texture is adjusted to users' interpretation of the scene and thus allows the viewer to "read between the lines". In the photorealistic display I use textures of the original image to shape the halftoning layer and to enhance reproduction of the original image details.

1.2 Research goals

The objective of my research is to enable an artist to use the halftoning texture as an expressive medium in artistic rendering and technical illustration.

The expressive use of halftoning is a departure from conventional methods of tone reproduction. My approach is to accept the fact that artificial texture is an inherent feature of halftoning and to use it as an extra channel of information. Thus, the "*straight*" mechanical reproduction of image features is substituted by *interpretive non-photorealistic* rendering that enhances display of image features and provides additional image information. The halftoning texture layer can be used:

- to depict objective information, present or implied by the image. Examples of objective information are image gradient, scene geometry, and illumination. The objective of my research is to use texture as a visual cue to enhance the expression of this objec-

tive information. Visual parameters of texture such as direction, scale, and variance are used for this purpose;

- to supply subjective information: artistic, cultural, and emotional. I convey subjective information by varying the style of texture and use the semantic meaning of texture shape, its resemblance to a natural material, or a particular illustration style.

In both cases the goal of interpretive halftoning is to enhance appreciation of image information by the viewer.

Clearly, the quality of such rendering is hard to estimate: it is not likely we are ever to be able to measure human emotions! On the other hand, one of the goals of my research is to outline a framework for an image quality measure suitable for expressive rendering. The intent in developing such a measure is to estimate the influence of different parameters of a technique on resulting image features. The measure quantifies the appearance of edges and textures in the halftoned image.

1.3 Outline of the approach.

My approach to the expressive rendering in binary media is based on conventional tone reproduction techniques: ordered dithering and error diffusion [Uli87]. Therefore, unlike previous NPR research [WS94, SPR⁺94, PB94, Elb95] I do not develop a new tone approximation system and use well known halftoning results to control the style and accuracy of the display. In particular, I control the shape of the halftoning texture by designing threshold matrices with desired properties. I use the resulting matrices to modulate thresholds of the error diffusion algorithm, thus controlling the accuracy of tone approximation and the contrast of the halftoning texture. I utilize current NPR results and adapt texture shape, scale and direction to the displayed information. I evaluate my results by quantifying the appearance of edges and textures with the edge based image quality measure.

There are three main parts to the thesis:

Introduction of arbitrary textures into halftoned images by the use of dither matrixes.

The following approaches to the construction of dither matrixes are investigated:

1. An image based texture is converted into a dither screen using an image processing algorithm. This technique allows an artist to introduce textures of natural materials: stone, brick, paper, canvas, pencil stroke.

2. Dither screen is generated procedurally. This approach allows for parameterized control of texture shape and generation of un-repeated textures of a desired size and scale.
3. Image based and procedural textures are combined into one dither matrix. For example, procedural pencil stroke is overlaid over image based paper texture.

The design of dither matrixes from a wide variety of textures is the corner stone of this research. The intended use of multiple texture styles is to convey artistic and esthetic image information through the semantic meaning of introduced textures.

Adaptation to image information. The display of images may be enhanced by adapting dithering texture and other halftoning parameters to the interpretation of image information. The following types of information are considered:

1. **Two-dimensional image information:** image gradient, edges, textures.
2. **Three-dimensional information:** object geometry and position, scene illumination and camera position.
3. **Subjective information** supplied by the user: the center of attention, importance of features, completeness (in the case of technical designs).

I develop image based and scene based techniques that enable an artist to modify texture shape, scale, direction, and contrast to the desired information. These techniques are applicable to rendering of photographs and computer generated scenes.

Image quality measure for evaluation of the halftoning algorithms. I develop an image quality measure that is based on multiscale edge analysis. This measure is applicable to both photorealistic and interpretive tone reproduction. The objective of the image metric study:

- to evaluate halftoning artifacts, such as quantization bands, abrupt changes of texture, warm-like dot alignment;
- to estimate distortion of image edges.

These measures are used to study the influence of various halftoning parameters and textures on the reproduction of image details.

1.4 Overview of the thesis

My research is an extension of conventional tone reproduction algorithms and is a new approach to non-photorealistic rendering. Thus, the thesis starts from the analysis of current results in both fields of computer graphics.

The discussion of the background material (Chapter 2) presents an overview of conventional approaches to halftoning. I identify the main properties of the human visual system that enable us to create an illusion of continuous tones and influence our perception of halftoned images. I present the most widely used dithering techniques: threshold matrix and error diffusion. The advantages and shortcomings of these approaches are discussed. The chapter is concluded by an analysis of image quality measures used for evaluation of halftoned images.

The current non-photorealistic rendering techniques are presented in Chapter 3. I outline objectives of expressive image synthesis and identify three main areas of the NPR research:

- approximation of art materials and techniques;
- development of rendering styles specific to computer generated images;
- comprehensive rendering.

I focus the discussion of current NPR results on the rendering of non-photorealistic images for binary displays. I conclude presentation of the background material by pointing out the shortcomings of the existing non-photorealistic rendering approaches.

The general concept and the fundamental texture control techniques common to all applications are found in Chapter 4. In particular, this chapter addresses the control of texture shape and contrast in halftoned images. My approach is based on the definition of texture shape through the use of threshold matrices in the ordered dithering algorithm. I identify essential threshold distribution properties of a dither matrix and approximate these properties by an image processing and procedural texturing techniques. I combine ordered dithering with the error diffusion algorithm to control the contrast of the halftoning texture.

The halftoning texture is used for artistic purposes in Chapter 5. I adapt texture shape, direction, scale and contrast according to the image information. Two approaches to artistic rendering are developed: image based and scene based. Image

based techniques use image buffers and are applicable to expressive rendering of photographs and 3-D scenes. The scene based approach is specific to 3-D rendering and is based on conventional texture mapping.

Chapter 6 discusses the issues of texture control in photorealistic rendering. I apply the texture control technique to enhance textures and edges of the halftoned image and present a numerical measure of image quality. This measure is based on multiscale edge analysis and allows us to quantify image distortions.

The thesis is concluded by the discussion of contributions and limitations of my texture control approach.

Chapter 2

Fundamentals of tone reproduction

In order to print continuous tone images we must approximate their intensities with a distribution of black and white pixels. The *halftoning* or *dithering* approach exploits spatial integration properties of the human visual system. That is, an area with interspersed black and white pixels is perceived as an intermediate gray tone at a distance. The main objective of research in halftoning is then to choose an effective arrangement of black and white dots for an accurate representation of the original image.

This chapter presents fundamental principles of halftoning. I start by identifying the most relevant properties of the human visual system. I overview the halftoning literature and discuss image quality measures used for the evaluation of dithered images.

2.1 Properties of the human visual system

Halftoning is a process of approximating continuous tone images by a combination of graphic primitives (e.g. black and white dots of various sizes). The objective is to create an illusion of the original image without the introduction of significant visual artifacts. While interpretation of images by a viewer is a complex process that depends on many factors including cultural background and viewing context, the successful creation of intended illusions is impossible without an understanding of the basic properties of the human visual system. In this section we review characteristics of the human vision that are relative in the context of halftoning. The full treatment

of this topic can be found [Sny85, Mar79]

Human vision is often modeled as a collection of photo-receptors interconnected into receptive fields. Each receptive field is tuned to a certain size, frequency, and orientation of visual stimulus. Any region of the retina contains receptive fields with varying properties. Thus, the perception of visual information is based on the combined response of overlapping receptive fields. The following properties of human perception are explained by this model:

- **Spatial integration.**

Receptive fields act as low pass filters of various sizes. When the input frequency becomes higher than the frequency response of these filters, we cannot distinguish between individual visual elements. Instead, we perceive the averaged response of receptive fields. A halftoned representation of a gray scale image is possible due to this property. When a dithered image is viewed at a distance, our visual system averages the black and white dots of the halftoning layer and the continuous tone is perceived.

- **Contrast sensitivity and edges.**

Our perception of brightness and darkness is based on contrast. While we are inaccurate in the judgment of the absolute lightness, we are able to perceive small variations of light intensities between neighboring regions. We think of sharp lightness discontinuities as edges. Due to the various sizes of receptive fields our visual system can be modeled as a collection of multiscale edge detectors. Psychophysical research [Mar79] suggests that the analysis of edges at multiple scales is a foundation of image understanding. Thus, reproduction of relative contrast and an accurate depiction of edges is the important goal in the design of halftoning algorithms.

- **Contrast sensitivity and frequency.**

Perception of contrast depends on the frequency of tone variations. Contrast sensitivity is not uniform and decreases for low and high frequencies. This property is important to classification of halftoning artifacts. The high frequency “blue” noise artifacts are less disturbing to the viewer [Uli87].

- **Orientation tuning.**

Sensitivity of the visual system is dependent on orientation of visual stimulus. It is the highest in horizontal and vertical direction. In the context of halftoning,

this property suggests that alignment of horizontal and vertical artifacts should be avoided by orienting graphics primitives in other directions.

These properties of the human visual system help to explain the effects of halftoning. I review the main techniques of halftoning in the following section and present numerical measures used for evaluation of the resulting images.

2.2 Conventional halftoning approaches

The objective of traditional halftoning algorithms is to approximate images with photographic accuracy without the introduction of visible artifacts. That is, a halftoned binary image should be perceived in the same way as a photograph or a continuous tone rendering on a color monitor. The body of research in *photorealistic* halftoning can be divided into the following classes:

Threshold matrix halftoning is a point process. A dither screen is used to threshold an image. Gray scale values of the original image are compared to the dither screen elements to determine a binary value of the output.

Error diffusion is a region based process. Approximation of the pixel in the original image depends on the errors accumulated in the approximation of the pixels in the neighborhood. Error diffusion algorithms differ in the order of image traversal and the size and shape of the neighborhood that influences approximation of the output pixel.

Optimization methods are based on minimization of approximation errors computed using perceptual image quality metrics.

Hybrid methods use some combination of point and neighborhood operations.

Due to the variety of binary display hardware it is important to adapt halftoning algorithms to the specific properties and limitations of the target output device. Spatial resolution and accuracy in setting individual pixels should be considered in the design of a halftoning process. Monitors have very limited resolution (typically around 75 dpi.), but are good at addressing isolated pixels. In the contrast, most printers have significantly higher resolution that ranges from 300 dpi. to 2400 dpi. Unfortunately, many printers are unable to set individual pixels without overlap. This property results in “bleeding” of ink from black pixels into white areas and is called “*dot gain*”.




Figure 2.1: Dither matrix (a) is tiled to produce a dither screen (b). The arrangement of thresholds in the matrix controls the appearance of the texture introduced into the halftoned image. The artificial texture is visible in the halftoned gray ramp in image (c).

Thus, in the following discussion I identify halftoning algorithms appropriate to either monitors or printers.

2.2.1 Threshold matrix dithering

Ordered dither halftoning is based on threshold matrices and is the simplest and the most widely used halftoning approach. These algorithms use an array of thresholds arranged into a dither matrix (Figure 2.1a). Dither matrices, tiled over the surface of an image, constitute a dither screen (Figure 2.1b). Gray scale values of the original image are compared against the corresponding thresholds in this screen. A gray scale value that is smaller than the threshold is approximated by a black pixel, and white otherwise.

The arrangement of thresholds in the dither screen is independent from the halftoned image. Thus, it is impossible to control the accuracy of tone reproduction within an ordered dither algorithm. The design of dither matrixes is optimized for images with slow varying gradient [Bay73, Uli87], therefore such images are approximated well by ordered dither. The following properties of a threshold array should be satisfied [Bay73, MP91, Uli93]:

1. *Uniform distribution of threshold values.*

The dither screen should contain the same number of pixels with the same values. In the 8x8 dither matrix (Figure 2.1a) all the values between 0 and 63 appear exactly once. This property enables the uniform reproduction of the maximum range of gray tones.

2. *Homogeneous spatial distribution of threshold values.*

Pixels with the same threshold values should be uniformly spread through out the dither screen. Thus, gray tones are approximated in the same fashion in different



Figure 2.2: Distribution of thresholds in the dither matrix defines the appearance of texture that is introduced into the halftoned layer. Clustered dither produces strong texturing artifacts. Texturing is less visible in the image with dispersed dither. Due to the fact that printers cannot set isolated pixels, clustered dither tone approximation is better than the one produced with dispersed dither matrix. Isolated pixels in the image on the right “bled” into the white space resulting in a strong shift of image contrast.

regions of the image. The homogeneity property is satisfied automatically when a small dither matrix is tiled to generate a dither screen (Figure 2.1b).

Limitations of the target display hardware must also be taken into account. It is advantageous to cluster thresholds together to minimize the effects of dot gain in most printing hardware. However, clustered dither leads to introduction of strong texturing artifacts that interfere with the display of fine image details (left image, Figure 2.2). Dither matrixes with dispersed distribution of thresholds produce less noticeable texture artifacts. Unfortunately, the use of dispersed dither creates many isolated pixels. Therefore, ordered dithering with dispersed threshold matrixes is applicable only for monitors and suffers from pixel “bleeding” in printers (image on the right, Figure 2.2).

The arrangement of threshold values within the dither matrix defines the artificial texture introduced into the halftoned image. Research has thus focused on developing dither matrixes that produce less noticeable texturing effects. The matrixes proposed by Bayer in 1973 [Bay73] are still been widely used. Bayer’s dither matrixes are typically small threshold arrays (up to 8x8 pixels) with diagonal, horizontal and vertical arrangement of thresholds. Thus, halftoning with Bayer’s dither arrays results in introduction of periodic regular structured texture into the halftoned image (Fig-

ure 2.1c). Ostromoukhov and Hersch [OHA94] reduced these regular texture artifacts by rotating Bayer's threshold matrixes. Thus, alignment of horizontal and vertical components of these dither matrixes is avoided.

It is possible to avoid the periodic texture artifact of small dither arrays if larger stochastic matrixes are used. There are a number of algorithms that construct large threshold matrixes: void and cluster [Uli93], blue-noise [MP91, MP92, SP94]. The design of large dither screens is partially based on our understanding of the texture perception by the human visual system.

The main advantage of threshold matrix halftoning is its simplicity. Ordered dither can be adapted for rendering images on either monitors or printers by changing the distribution of thresholds in the dither matrix. Also, the texture introduced into the halftoned image is easily controlled through the modification of the threshold matrix. However, ordered dither uses the same threshold matrix for all the images. Thus, important image structures such as edges are not accounted for in the halftoning process.

2.2.2 Error diffusion

Error diffusion is a region based halftoning algorithm. The value of the binary pixels in the region depends on the input gray values, on the threshold and on the weighted errors diffused from the neighboring regions. The error associated with this approximation is propagated to un-processed neighbor regions, thus compensating for the approximating error over a region. When compared with threshold dithered images, the results of error diffusion exhibit better rendition of image features and less visible halftoning artifacts.

The shape of approximating regions and the order in which they are processed determines the artifacts introduced by the error diffusion methods. In most cases single pixel regions and scan line processing paths are used. The most popular error diffusion technique was presented by Floyd and Steinberg [FS76]. While being simple and often producing good quality output, this algorithm creates strong correlated "worm-like" patterns in dark and light areas. Numerous techniques were proposed to alleviate the problem. Stucki [Stu79] reduced the worm-like artifact by diffusing errors from a larger neighborhood (more than 4 pixels). Ulichney in [Uli87] described a serpentine scanning approach, which reverses the direction of processing every scan line. Witten

and Neal [WN82] avoided worm-like artifacts by using error diffusion along a Peano space filling curve. Knuth [Knu87] specified the order of processing and error distribution within a matrix. Error diffusion in a wavefront fashion is developed by Naiman and Lam [NL96].


The above error diffusion algorithms process an image one pixel at time. Such an approach leads to creation of many isolated pixels in the halftoned image. Due to the dot gain in most printers, these isolated pixels may not be reproduced when printed (image on the left, Figure 2.3). Dot clustering may be achieved by changing the size of processing regions. Velho and Gomez [VG91] divided an image into segments along a space filling curve. The appropriate number of pixels in the resulting regions is set to black, thus the average intensity over the region is approximated. The resulting approximation error is then diffused into the next segment along the space filling curve.

One of the main advantages of error diffusion techniques over threshold matrix methods is the ability to adapt the halftoning process to the local image features. In particular, the perceived information detail in the binary image is increased if the edges of the original image are preserved. Thurnhofer and Mitra [TM94] modified the standard error diffusion process to use a nonlinear filter for feature extraction. Thus, approximation errors are not diffused across sharp edges. Buchanan and Verevka [BV95] and Velho and Gomez in [VG95] modified clustered space filling curve halftoning to handle edges. One-dimensional edge detection filter is used to ensure that processing regions do not cross significant edges. The result of halftoning an image with this clustered space filling curve algorithm is presented (image on the right, Figure 2.3). Clustering of pixels allows us to compensate for dot gain in printers (compare with Floyd-Steinberg algorithm, image on the left, Figure 2.3). Also, error diffusion with edge detection improves rendition of edges in the image.

Hybrid methods that combine error diffusion and ordered dither are used to improve quality of halftoned images. In standard error diffusion the threshold value is constant for all the pixels in the image. Billottet-Hoffman and Bryngdahl [BHB83] used threshold matrixes to modulate thresholds in the error diffusion process. This threshold modulation approach allowed them to create clustered halftoned images of high quality. Also, the worm-like artifacts of error diffusion and regular texturing of dither arrays were reduced. The use of image dependent threshold by Eschbach and Knox



Floyd-Steinberg algorithm



Space filling curve with edge detection

Figure 2.3: Error diffusion algorithms differ in the size of approximating regions and in the order of image traversal. Floyd-Steinberg method uses scan-line image processing and approximates one pixel at a time. Space-filling curve error diffusion processes images in clusters along a Peano curve. Due to pixel clustering, the later algorithm is more appropriate for printing. Also, when compared with clustered dithering Figure 2.2, space-filling algorithm produces less structured texture artifacts.

[EK91] induced edge enhancement in error diffused images.

2.2.3 Optimization approaches to halftoning

Threshold matrix and error diffusion algorithms are the most widely used and are well investigated by computer graphics researchers. While these approaches are able to produce good quality output, their performance is hard to improve due to the inherent limitations of sequential order of pixel processing. That is, only information about the previously processed pixels can be used in the approximation of the current gray scale value. Moreover, once the value of a binary pixel is set, it is impossible to change it afterwards to improve the overall appearance of the image. Therefore, multiple passes over the halftoned image are needed. An example of how results of error diffusion can be improved by multi-pass processing is given in [Fan94]. The worm-like artifacts produced by the Floyd-Steinberg algorithm are reduced by the second pass over the halftoned image.

A new class of non-sequential halftoning algorithms based on optimization approaches has been developed recently [ZLE92]. The objective of this approach is to improve rendition of a gray scale image iteratively. In general, these algorithms proceed as follows:

1. An initial approximation of the continuous tone image is chosen. Results of the ordered dither or the error diffusion can be used as such initial approximation.
2. The accuracy of the current approximation is evaluated using an image quality measure. The measure may also account for the influence of the halftoning artifacts (e.g. texturing) on the perception of the input image features.
3. An optimization technique changes the halftoned output to improve approximation of the original image. The following optimization techniques have been used in the past: direct binary search [ZLE92], neural networks [GRS93], simulated annealing [GRS93, MA92].
4. Go back to the step 2, or stop the process if the maximum number of iterations is reached.

Even though generating halftoned images using such an optimization approach takes considerable computation resources, the resulting output is perceptually superior to the traditional dithering methods.

2.3 Numerical measures of image quality

The discussion of halftoning techniques in the previous section was mainly based on visual observations of image quality. These visual observations are subjective and difficult to reproduce. On the other hand, quantitative metrics provide an objective and reproducible measure of image quality. Moreover, optimization algorithms use such quantitative metrics to improve the quality of halftoned images. For these reasons visual modeling and the derivation of reliable image quality metrics has recently received increased attention in the computer graphics research community.

Most image quality metrics used in halftoning are based on frequency-domain models of human perception. Ulichney [Uli87, Uli88] examined the spectral characteristics of textures produced by halftoning using the modulus of the discrete Fourier transform. The influence of different frequencies on perception of errors is not uniform and can be described by the Contrast Sensitivity Function (CSF). The CSF can be adjusted for different viewing conditions and properties of the printing device. Also, the CSF can account for direction dependent sensitivity of the visual system. CSF-weighted Fourier transforms were used to measure errors in halftoned images by several researchers [AA92, PN92, ZLE93, MV93].

In the previous discussion (Section 2.1) I pointed out that the human visual system can be modeled by radially symmetric narrow-band filters in the spatial frequency domain. A spatial filtering model for texture analysis was used as a halftoning quality measure by Mitsa et. al. in [MVA93]. Their model uses radially symmetric filters represented by Gaussian functions of different width. A frequency domain error measure is applied to each narrow-band filtered image. The overall error is computed as a CSF weighted sum of errors in different bands.

In recent years wavelet based algorithms for image coding and analysis have been developed. Wavelet decomposition is a multiscale representation of image information produced by a bank of filters of variable width. The important property of wavelets is their ability to combine spatial and frequency information. Wavelet filters have been used in the design of dither matrices in [MB95]. Wavelet based error measures were proposed for the evaluation of various rendering algorithms in [GMY97]

The work presented so far has focused on the measurement of general image distortion. These algorithms are well suited for overall assessment of image quality. Unfortunately, the existing general purpose measures are difficult to adapt for the evaluation of specific image structures. Since our perception of images is dependent on the analysis of edge information (Section 2.1), edge specific image measures are important to investigate. *Edge correlation* [ML80] is the only edge based metric for analysis of halftoned images developed thus far. Mitsa [Mit93] found poor correspondence of results produced by this measure and psychovisual data. The reason is that this criterion only considers preservation of edges in the halftoned image and completely ignores gray scale rendering.

I have developed a set of measures based on multiscale analysis of edge information. These measures were applied to evaluation of halftoned images in the previous research [VFB98]. The further improvement of these measures and their application to the analysis of texture control techniques are investigated in this thesis (Section 6.3).

2.4 Summary of photorealistic halftoning

The main goal of photorealistic halftoning is to generate a binary image that is the best possible approximation of the continuous tone image and to avoid the introduction of significant artifacts. A number of techniques have been developed over the

years. The solutions range from fast and simple threshold matrix to complex computationally expensive optimization approaches. The improvement of perceived image quality have been often achieved by tailoring the halftoning algorithm to local image features such as edges. The increasing knowledge of the human visual system guides the development of new dithering techniques free from disturbing artifacts. Unfortunately, due to limitations of the current display hardware, the halftoning artifacts cannot be avoided completely.

In this thesis I develop tools to control the halftoning texture by extending conventional halftoning approaches. I define the shape of the texture using the arrangement of thresholds in the dither screen. I combine ordered dithering with the error diffusion algorithm to improve the approximation of image tones and to control the contrast of the halftoning texture. I apply the texture control approach to both photorealistic and artistic rendering of images. I also develop two edge based measures suitable for analysis of halftoning artifacts and distortions. I use these measures to quantify the parameters of my technique and to assess the accuracy of image reproduction.

Chapter 3

Expressive and interpretive rendering

In the previous chapter I presented the main approaches to photorealistic tone reproduction on binary devices. The notion of image quality in that context is closely related to the accuracy of *copying* visual information from one medium to another. The following question should be asked: *“Is the correct reproduction of image tones and colors the only way to ensure our understanding of this image?”* This chapter is a discussion of alternative means for representation of visual information. In particular I review the existing computer graphics approaches to expressive image display: *non-photorealistic rendering*.

3.1 Two approaches to expression of reality

What is the best way to express information in an image? It appears that the answer to this question depends on the purpose of the image on the first place. Clearly, a photograph is a must in the cases when an accurate recording of the visual information is required. Although, there are many cases when visual information is best represented through artistic drawing or painting.

What is the relationship between an artistic rendering and a photograph or a photorealistic computer rendering? While, both are *“expression of form upon a plane surface”* [Spe17], a photograph is a mechanical recording of reality, and the art work is the result of artistic interpretation of this reality:

... it (*a photograph*) can only at best give mechanical accuracy, whereas art gives the impression of live, individual consciousness. [Spe17]

In the case of a photograph, we rely on the quality and accuracy of a camera and photographic materials to create an exact copy of what we see. In the case of computer rendering, we model the objects in the scene, their optical properties and the light propagation through them, and only then compute an approximation of our world. Clearly, the images generated by these photorealistic methods are representations of only what the camera sees and what the rendering algorithm can compute. All the information about the world that is beyond the sensitivity of photographic materials and the accuracy of our models is not accounted for.

On the other hand, the artistic interpretation of reality allows an artist to create effective representations of visual information using a particular medium. Thus, all visual information needed for the purpose of communication is accounted for and is carried through in the image. Also, the artistic interpretation of the world leads “*to selection of significant and suppression of non-essential*” [Spe17] in the final image, thus the intended message is not obscured by unwanted details.

The need for expressive and selective display of information is reflected by the current practice in technical illustration and graphic design. Even though, modern cars and airplanes are modeled and tested entirely on computers, the repair manuals are still drawn by hand. Graphic designers often use photorealistic 3-D renderers and then manually re-touch the resulting images using 2-D paint programs.

The field of computer graphics that addresses the issues of expressive image synthesis is called *non-photorealistic rendering* (NPR). I review the main results of this area in the following section.

3.2 The main styles of non-photorealistic rendering

The objective of research in non-photorealistic rendering is best expressed by Strothotte et. al. [SPR⁺94]:

... to study alternative methods of rendering images to provide designers with the ability to express more than just the geometry or the photometric qualities of their designs

In other words, this research is directed towards development of tools to assist designers and illustrators in creation of expressive images. While photorealistic rendering techniques are governed by our understanding of laws of physics, the NPR approaches

are directed by artistic interpretation of the world. The laws of physics in non-photorealistic rendering can be sacrificed to a more expressive and understandable display of visual information. In most cases in NPR the image synthesis is controlled by the user input. Thus, the author of an image can choose the best means for expression of the desired information. The choice of expressive means depend on the chosen style of rendering. For the purpose of this discussion I separate different NPR approaches depending on the style of the resulting rendering.

3.2.1 Modeling of traditional art materials

Most of the effort in non-photorealistic image synthesis has focused on modeling of artistic media. The objective of this research is to imitate traditional art materials. Strassmann [Str86] presented a system that simulates a brush stroke for interactive rendering of images in the style of *sumi-e* paintings. Curtis et. al. [CAS⁺97] described how various artistic effects of watercolor can be simulated. This watercolor model can be used as a part of an interactive paint system or as a mechanism for automatic rendering of non-photorealistic three-dimensional scenes. An extensive modeling of interaction between pencil and paper has been proposed by Sousa [Sou98]. The objective of this system is to encode the rules of pencil sketching for automatic illustration.

3.2.2 Computer graphics as a medium in itself.

While the modeling of art materials allows designers to generate images that resemble traditional illustration styles, Landsdown and Schoffield [LS95] stated the importance of creating rendering style specific to computers. That is, instead of approximating what a brush or a pencil stroke looks like, it is important to investigate what graphics primitives should be used in computer generated expressive renderings. For this purpose they developed a system of expressive marks. Similar stylized “brush” strokes that do not model any particular medium are presented by Haeberli in [Hae90].

3.2.3 Comprehensive rendering

The non-photorealistic rendering techniques reviewed thus far concentrated on developing a pleasing alternative style of images suitable for artistic work. The other goal of expressive image synthesis is to improve viewer’s comprehension of displayed

geometric objects. This style of rendering is more applicable to the field of technical illustration.

The early works in this area concentrated on comprehensive wire-frame rendering of objects. Appel et. al. [ARS79] achieved the illusion of depth by introducing the haloed line effect. Kamada and Kawai [KK87] improved the wire-frame display by changing the style of lines. Dolley and Cohen [DC90] further enhanced the treatment of line drawings by controlling the line style according to importance tags and illustration rules. Strothotte et. al. [SPR⁺94] designed a sketch-renderer. They studied how the variation of line styles and limited shading can direct viewer's attention to the desired parts of the image.

Instead of creating an NPR renderer, Saito and Takahashi [ST90] presented a post-processing approach to enhancing the appearance of geometric forms. They augmented a photorealistic image by simple lines and textures, thus edges of the displayed objects and surface directions are highlighted. Similarly, Interrante et. al. [IFP96, IFP97, Int97] studied the use of texture to improve the display of surfaces in visualization of medical data. Understanding of transparent objects was improved by modulating a directed texture on their surfaces. Also, the effective (or visually optimal) direction of textures to convey surface orientation were identified.

3.3 Expressive tone reproduction

Most work in the field of NPR does not take into consideration the reproduction of the generated images on printers. Unfortunately, the use of traditional halftoning methods may significantly alter the appearance of textures and colors in non-photorealistic images. Therefore, much of the effort in enhancing display of visual information may be wasted when these images are halftoned by conventional techniques.

In this section I discuss the expressive rendering styles that produce binary images. While some of the previously mentioned techniques [ARS79, KK87, DC90, SPR⁺94] render black and white images, their development was concentrated on line rendering, and did not take into account shading and continuous tone. Hence, the following discussion is focused on expressive tone reproduction.

3.3.1 Artistic styles of tone reproduction

Much of the work in non-photorealistic rendering was inspired by traditional line art such as engravings, wood cuts, pen and ink illustrations.

A digital engraving system DIG^i Dürer was developed by Pnueli and Bruckstein [PB94]. Continuous tones are reproduced by controlling the density of graphics elements. The graphics elements lie on level contours of the potential field induced by the image via Eikonal equation. The system allows a user to specify initial conditions to this equation to alter the appearance of the emerging halftoning layer.

Leister [Lei94] presented a system for rendering images in copper plate style. Leister's approach uses a ray tracer to render surface curves on the depicted objects. The illusion of shading is created by varying the width of these curves in the post-processing step. The resulting images resemble engravings. However, the proposed system fails to reproduce subtle tone variations.

Elber [Elb95, Elb98] created the illusion of shading by covering free form surfaces with isoparametric curves. Tone variations are achieved by changing the density of these curves. The lack of control over curve width does not allow this technique to reproduce smooth continuous shading, and is thus results in strong quantization artifacts in generated images.

Many of the shortcomings of the previous line based rendering systems are addressed in pen and ink work at University of Washington. Winkenbach and Salesin [WS94] created automatic illustrations in the pen and ink style by texture mapping procedural textures on polygonal objects. These procedural textures are able to depict not only subtle variations of tone but also material properties of surfaces. This work was extended to rendering parametric surfaces [WS96]. Similar texture based approach is used in an interactive pen and ink system by Salisbury et. al. [SABS94, SWHS97].

3.3.2 Extensions of halftoning to expressive rendering

The modeling of line rendering styles allows us to create images that resemble traditional illustrations. However, the existing systems either fail to reproduce subtle tone variations [Lei94, Elb95] or involve computationally expensive operations (e. g. rendering of equipotential contours [PB94]). The other approach is to utilize the results from photorealistic halftoning and to apply these conventional methods to artistic

dithering of images.

Buchanan [Buc96] showed that by altering various parameters of the clustered error diffusion method a wide variety of textures can be introduced into the halftoned image. These textures are generated by adjusting the shape of the curve and the cluster size. Unfortunately, this approach lacks intuitive control over the appearance textures and only a limited style of textures can be introduced.

Ostromoukhov and Hersch [OH95] developed a halftoning algorithm that enables artists to design single screen elements. These micro-screens are outlined by curves and can take a desired geometric shape. A great variety of screen elements have been presented: calligraphic letters, oriental polygonal patterns, images of fish and birds. However, since a screen shape must be outlined by a curve, this approach does not lend itself to introduction of natural textures such as paper, wood, or stone.

When compared to the line art algorithms (Section 3.3.1), these extensions of conventional halftoning methods result in accurate and computationally efficient tone reproduction in artistic style.

3.3.3 Comprehensive tone reproduction

The expressive tone reproduction in the previous section was achieved by changing the style of halftoning textures. Thus, cultural, emotional, and artistic information can be conveyed. The other approach is to enhance viewer's appreciation of displayed information by highlighting important image features.

A technique similar to comprehensive rendering of Saito and Takahashi [ST90] was used to improve rendition of edges by Buchanan et. al. [BSV98]. Edges in any image dithered by a conventional halftoning algorithm are highlighted in the post-processing step. The authors concentrated on clustered dithering, thus images printed by standard hardware are enhanced.

Sloan [SC93] improved comprehension of image gradient by constructing a dither screen from directional dither matrixes. A dither matrix was inserted into the screen if its direction coincided with the local gradient of the image. Thus, texture features are aligned with image gradient in the halftoned output. Unfortunately, the use of small dither matrixes limits a number of possible directions and creates texture segmentation artifacts.

Streit and Buchanan [SB98] developed an importance driven approach to tone reproduction. In their system the placement of graphics elements is controlled by an arbitrary image importance function. By varying these importance functions different types of information can be highlighted at different scales of the halftoned image. As an example, the authors presented how perception of image contrast, variance and gradient can be enhanced.

3.4 Summary of expressive rendering

This chapter was devoted to the discussion of non-photorealistic rendering techniques. Inspired by artistic methods of visual expression, these techniques allow a user to enhance representation of visual information.

The current research in NPR ranges from modeling of traditional illustration styles to development of expressive means unique to computer generated images. Much of the work is devoted to improve comprehension of depicted information for visualization and illustration purposes.

The expressive methods of tone reproduction are most relevant to the current research. At the present time techniques that generate images in the style of line art illustrations lack simple means to control the accuracy of shading. A number of proposed extensions of conventional halftoning are able to generate high quality tone reproduction, but are limited in style of the introduced texture features. Moreover, the reviewed artistic approaches are either hard to parameterize or it takes considerable time and skill to master [OH95, SABS94, SWHS97]. The goal of my research is thus to develop a halftoning system that allows a user to tailor its parameters to an interpreted information and to generate images in a variety of artistic styles. My system extends conventional halftoning approaches to NPR and is based on the control of texture in binary images. The fundamentals of texture control are developed in the following chapter.

Chapter 4

Fundamentals of Texture Control in Halftoning

Previous chapters presented two approaches to rendering images in bi-level media: photorealistic halftoning and non-photorealistic, or interpretive tone reproduction. The difference in objectives in those approaches is reflected in the role of the halftoning texture. Photorealistic halftoning research attempts to hide the halftoning texture and to preserve all the features of the original image. Interpretive rendering techniques use texture to enhance presentation of visual information and to hide unnecessary details.

Even though photorealistic and artistic approaches have different objectives and different relationship to texture, they use texture for fundamentally the same purpose: approximation of tones. Thus, it is possible to extend photorealistic methods of texture control to interpretive halftoning. This is the approach taken in this work.

This chapter generalizes results of the traditional halftoning and lays a foundation of texture control. My approach is based on the ability of the ordered dither algorithm to define halftoning texture through the arrangement of threshold values in the dither matrix. This chapter states and generalizes the important properties of the dither matrix: uniform and homogeneous distribution of threshold values. I present two methods for generating dither matrixes based on image processing and on procedural textures. The use of a desired texture in the dither matrix guarantees the appearance of this texture in the resulted image. The strength of this texture is controlled by combining ordered dithering with an error diffusion algorithm.

4.1 Generalized ordered dithering

The ordered dither algorithm (Section 2.2.1) is based on thresholding gray scale values of the original image with thresholds in the dither screen. Thus, the value $b_{i,j}$ of the output pixel (i, j) is determined by the following formula:

$$b_{i,j} = \begin{cases} 0 & \text{if } g_{i,j} \leq t_{i,j} \\ 1 & \text{otherwise} \end{cases} \quad (4.1)$$

where, $g_{i,j}$ is the gray scale value, and $t_{i,j}$ is the corresponding threshold from the dither screen.

In most previous research a dither screen is constructed by tiling a dither matrix \mathcal{T} . Therefore, in the case of $m \times n$ dither matrix \mathcal{T} , a threshold value $t_{i,j}$ is found as:

$$t_{i,j} = \mathcal{T}_{i \bmod m, j \bmod n} \quad (4.2)$$

where $a \bmod b$ denotes the remainder of a division. Equation 4.2 represents a mapping between a dither matrix and a dither screen.

In this thesis I generalize construction of a dither screen to an output of a function $T(u, v)$, where u and v are normalized texture coordinates between 0 and 1. In most cases $T(u, v)$ function represents a combination of a *kernel* function $\mathcal{T}(s, t)$ and a mapping function \mathcal{M} . Thus, a kernel function is a generalized dither matrix and *tiling* is a special case of a mapping function \mathcal{M} . Other examples of a mapping are:

Scale: $T(u, v) = \mathcal{M}_s \circ \mathcal{T} = \mathcal{T}(s_u u, s_v v)$, where s_u and s_v are scaling factors.

Rotation: $T(u, v) = \mathcal{M}_\alpha \circ \mathcal{T} = \mathcal{T}(\cos(\alpha)u - \sin(\alpha)v, \sin(\alpha)u + \cos(\alpha)v)$, where α is the rotation angle.

The arrangement of threshold values in the dither screen defines resulting halftoning texture. Thus, the challenge is to build a kernel function with a desired texture. Two possible approaches exist in computer graphics: to use images of existing textures, or to generate a desired texture procedurally. The following two sections explore both approaches.

4.2 Conversion of an arbitrary image into a dither matrix

In this section I investigate the use of arbitrary images as kernel functions of a dither screen.

In general, a kernel function is defined on a continuous interval, however, an image is a discrete two dimensional array of values. In order to use an image as a dither kernel, threshold values in between image pixels must be defined. This problem has been addressed in previous texture mapping research by filtering an image with a smoothing function [FvDFH90]. Thus, in the rest of the thesis I use images as functions defined on a continuous interval. Moreover, in this section I assume that there is a generalization of a dither matrix to a dither kernel function and present an algorithm that converts an arbitrary image into a dither matrix.

Arbitrary gray scale textures can be used as threshold matrices in the ordered dither algorithm. However, resulting halftoning is often a poor approximation of the continuous tone image. Consider the bark texture presented in the top left corner of Figure 4.1. When this gray scale image is used as a dither matrix it failed to reproduce a simple ramp (black on the left white on the right). Thus, the texture must be transformed to satisfy essential properties of *good* threshold matrices developed by the previous research [Bay73, MP92, Uli93] and discussed in Section 2.2.1:

1. Uniform distribution of threshold values.
2. Homogeneous spatial distribution of threshold values.

The task of uniform distribution of pixel values (property 1) has been addressed in image processing by the *histogram equalization* (*HE*) algorithm [Jai89]. The *HE* transform is based on global cumulative distribution of image values $H(g)$. The gray-scale value g is transformed as follows: $\bar{g} = H(g)$.

In order to achieve the spatial homogeneity of threshold values (property 2), one can split an image into small blocks and apply the *HE* algorithm independently to each block. That is, pixels of the block are transformed using cumulative distribution of values $h_b(g)$ within this block. Thus, each block becomes equivalent to a small dither matrix within a large dither screen. Since this dither screen has both desired properties outlined above the resulting halftoned images are also good approximations of the input gray-scale values. Unfortunately, this localized use of histogram equalization leads to strong blocking artifacts and destroys some characteristic features of the input texture.

These limitations were resolved by *adaptive histogram equalization* (*AHE*) approaches developed for local contrast enhancement. While many versions of the *AHE* exist



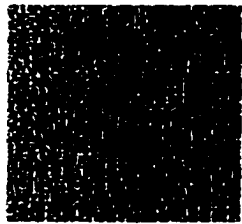
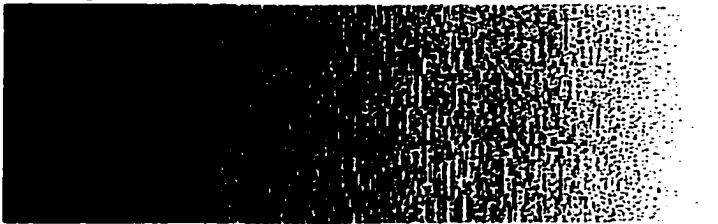
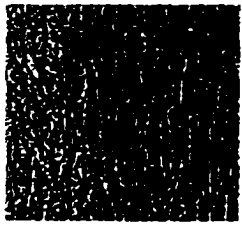
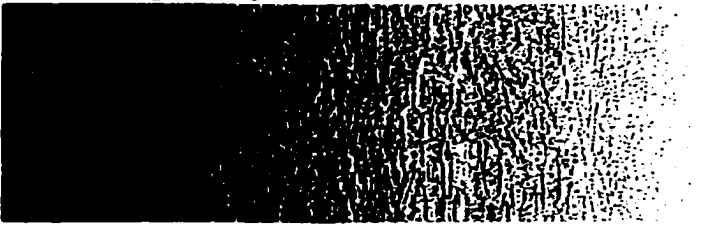
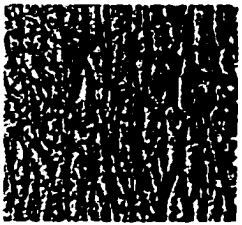
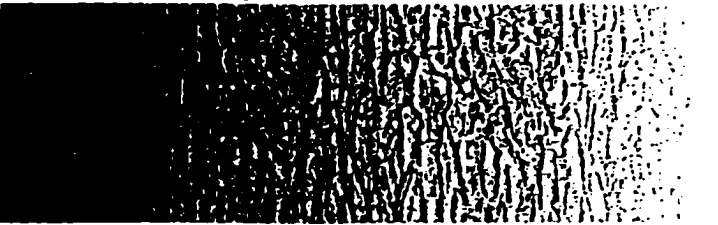
Halftoning images with textures as a dither matrix	
Texture	Halftoning result
Unprocessed bark texture	
	
Histogram equalization of 6x6 texture blocks	
	
Adaptive histogram equalization $w = 6$	
	
Adaptive histogram equalization $w = 12$	
	

Figure 4.1: Bark texture is used as a dither matrix. The halftoning with unmodified texture fails to reproduce the gray ramp. When the histogram equalization (*HE*) algorithm is applied to small blocks of the image, the resulting dither screen has the desired uniform and spatially homogeneous distribution of threshold values. Unfortunately, characteristic features of the texture are destroyed. The *AHE* algorithm preserves local characteristics of the input texture, while approximating the desired distribution properties of the dither screen. The size of blocks in the *AHE* algorithm controls the size of texture features introduced into the halftoned image. The images are printed at 300 dpi.

[Gau92], I have chosen the following algorithm:

1. An input texture is split into small blocks.
2. Cumulative distributions $h_b(g)$ are computed for each block.
3. Local distribution $h(g)$ for a pixel within a block is computed as a linear approximation between $h_b(g)$ of neighboring blocks.
4. The texture is transformed using local distributions: $\bar{g} = h(g)$

The resulting texture based dither kernel is free from blocking artifacts. Also, the desired properties of threshold value distributions are closely approximated. The size of blocks in the splitting process controls the size of texture features introduced into the halftoned images (Figure 4.1). Large blocks correspond to strong texturing effects. I found that block sizes between 3x3 and 8x8 produce the best results.

The AHE algorithm is suitable for a large number of textures. However, this texture filtering algorithm preserves texture features of a limited size. Texture features that are significantly larger or smaller than the window size of the AHE algorithm will likely to disappear from the halftoning layer. Thus, the appearance of textures with a combination of large scale and small scale features is distorted.

Typically, the presence of texture features that are larger than the AHE window can be detected by a large number of pixels with the similar intensities. In other words, the pixel histogram becomes strongly un-balanced. Zuiderveld [Zui94] suggested clipping histogram values to a threshold and then redistributing excessive values to other intensity ranges. Such clipping and redistribution leads to more balanced histograms. Adaptive histogram equalization with the resulting histograms produces limited contrast enhancement in homogeneous image regions. Thus, the algorithm is called contrast limited AHE (CLAHE).

Figure 4.2 demonstrates the effect of histogram clipping in the CLAHE algorithm. The 100% clipping threshold corresponds to no clipping and is equivalent to the original AHE algorithm. On the other hand, low values of clipping threshold reduce the amount of histogram equalization and are equivalent to halftoning with un-processed texture. Halftoning with limited histogram clipping enables as to preserve texture features at a wide range of scales. However, clipping histograms in the CLAHE algorithm compromises uniformity of pixel distributions in small regions. Thus, CLAHE is







Generating dithering texture by clipped adaptive histogram equalization	
Variations of AHE	Clipped AHE over 4x4 regions
Original "Wrinkled paper" texture 	Clipped by 25% 
Region 16 x 16 	Clipped by 50% 
Region 4 x 4 	Clipped by 75% 

Figure 4.2: Contrast limited adaptive histogram equalization (CLAHE) enables us to preserve texture features at multiple scales. The scale of texture elements is controlled by changing the region size of the AHE algorithm (left column). However, this method preserves either large scale features (for 16x16 regions) or small scale features (for 4x4 regions). Therefore, important texture details may be lost. CLAHE limits the amount of histogram equalization in regions with similar texture values. As a result texture features of a larger scale are introduced. However, limited histogram equalization results in poor tone reproduction. The test images are printed at 300 dpi.

a compromise between the accuracy of tone reproduction and preservation of texture features at multiple scales.

Overall, the adaptive histogram equalization algorithm modifies distribution of pixel values in an image to satisfy homogeneity and uniformity properties of a dither matrix. Therefore, a texture image can be used as a dither matrix and be introduced into the halftoning layer. The scale of texture features is controlled by two parameters of the CLAHE algorithm: processing window size and histogram clipping. Unfortunately, the goal to preserve large scale texture features often leads to poor local tone approximations.

4.3 Procedural dither textures.

The previous section described image based dither kernels. It is a powerful approach enabling the use of texture images in halftoning. However, image based dither kernels are limited by the image resolution and rely on filtering to generate sub-pixel values. Moreover, the appearance of texture features and the accuracy of tone reproduction are influenced by parameters of the CLAHE algorithm. Most of these limitations are eliminated by procedural dither kernels.

Procedural texture modeling is an active research area of computer graphics [EMP⁺94]. The following advantages of procedural textures are essential to this work:

- Procedural dither textures are generated directly without an image processing step.
- Procedural textures do not limit the kernel size. Thus, non-repeated textures of any size can be produced.
- The scale of texture features and the amount of local detail is controlled directly and can be parameterized. That is, a large variety of similar textures can be produced at different scales.

4.3.1 Generalization of essential properties of a dither matrix

Previous research in texture modeling enables generation of a great variety of textures. However, only textures with homogeneous and uniform distribution of values can be used for halftoning. Here I express the uniform distribution of values for a function $\tau(s, t), 0 \leq \tau(s, t) \leq 1$ defined on the unit box $[0, 1] \times [0, 1]$.

Let $\tau_{a,b}$ be a set of points (s, t) such that: $a \leq \tau(s, t) \leq b$, where $a, b \in [0, 1]$ and $a < b$. Then, $|\tau_{a,b}|$ is a measure of a set $\tau_{a,b}$. A function has a uniform distribution of values if $|\tau_{\frac{i}{n}, \frac{i+1}{n}}|$ is constant for any integers $i, n, n > 1, 0 \leq i < n$.

For example, in the case of a one-dimensional and invertible function $\tau(s)$, the uniformity can be expressed as follows:

$$\forall i, n, n > 1, 0 \leq i < n, \int_{\frac{i}{n}}^{\frac{i+1}{n}} \tau^{-1}(t) dt = \text{const} \quad (4.3)$$

In this thesis procedural dither kernels are constructed as nearly periodic extensions of the $\tau(s, t)$ functions. Thus, similarly to traditional dither screens, such dither kernel has homogeneous distribution of output values. Examples of procedural dither kernels are found in the following section.

4.3.2 Procedural dither kernels

The requirement for uniform distribution of pixel values limits the set functions that can be used for construction of procedural dither textures. For example, a *linear ramp* function $\tau(s) = s$ can be used as a dither texture, while $\sin(s)$ function does not satisfy the uniform distribution property.

In this thesis the procedural textures are based on the linear ramp function. Halftoning with double sided ramp function:

$$\tau(s, t) = \begin{cases} 2s & \text{if } s \leq 0.5 \\ 1 - 2s & \text{otherwise} \end{cases} \quad (4.4)$$

approximates tone variations by parallel lines of variable width (Figure 4.3). Line density in dark regions can be changed by combining two scaled linear ramps:

$$\tau(s, t) = \begin{cases} 1 - s & \text{if } s \leq I < 0.5 \\ (1 - 2I)(1 - \frac{s-I}{0.5-I}) & \text{if } I < s \leq 0.5 \\ 1.5 - I - s & \text{if } 0.5 < s \leq I + 0.5 \\ (1 - 2I)(1 - \frac{s-I-0.5}{0.5-I}) & \text{if } I + 0.5 < s \leq 1 \end{cases} \quad (4.5)$$

where I determines the lightest intensity that is represented by high density lines.

The previous texture functions were one-dimensional and produced textures with features in only one direction. Two linear ramps can be combined in a two-dimensional manner. Thus, dark image regions are rendered with cross-hatched lines (Figure 4.3 right). The corresponding function is defined as follows:

$$\tau(s, t) = \begin{cases} It & \text{if } s \leq I \\ (1 - I)s + I & \text{otherwise} \end{cases} \quad (4.6)$$

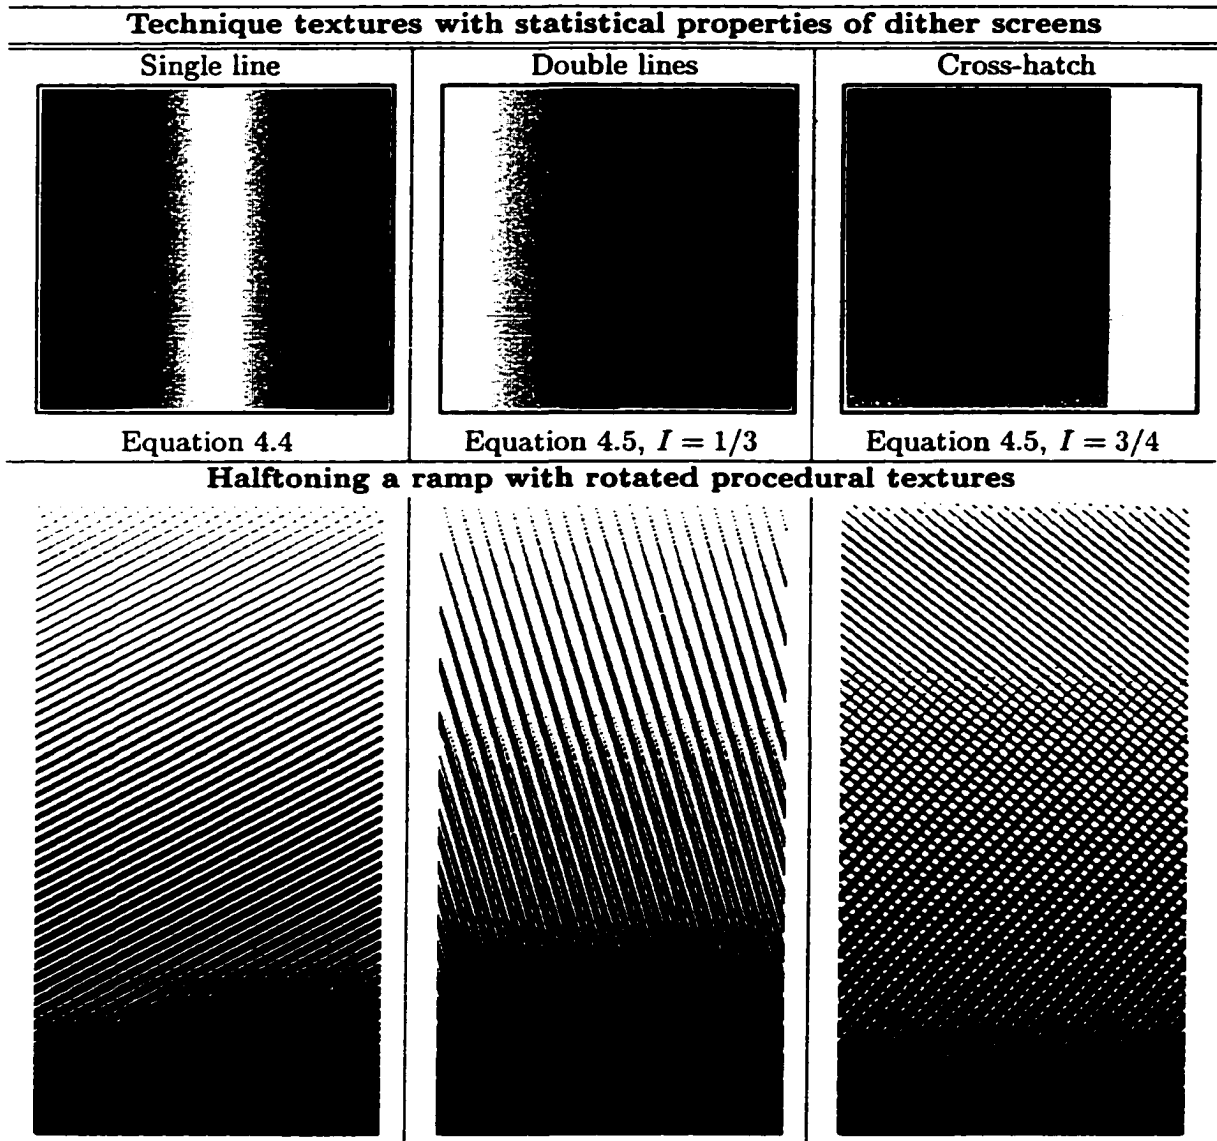


Figure 4.3: Procedural dither kernels constructed from a combination of linear ramps. Single line texture approximates image tone by changing line width. In addition to variable line width, double line texture changes line spacing in darker regions of the gray ramp (center). Cross-hatch texture approximates tones by adjusting line width in combination with lines in perpendicular direction.

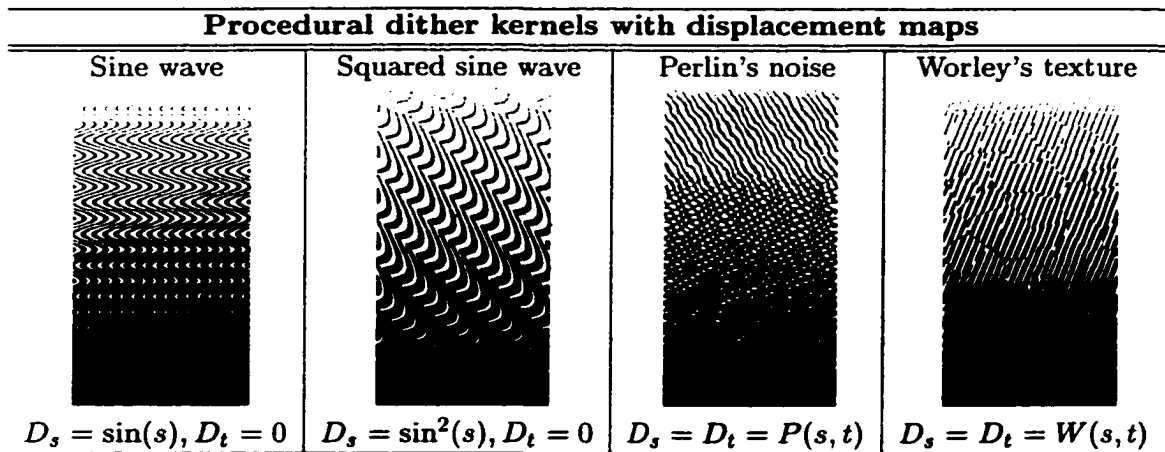


Figure 4.4: Variations of base textures (Figure 4.3) are generated by perturbing texture coordinates (s, t) with displacement functions. The displacement functions can be derived from simple mathematical functions such as $\sin(u)$, $\sin^2(u)$; or from procedural textures: Perlin's noise, Worley's textures.

where I determines the lightest intensity that is represented by the cross-hatched texture.

Examples of dither kernels \mathcal{T} reviewed thus far were constructed with a base $\tau(s, t)$ function. As a result, textures are the same in any region of the dither screen. Texture variation can be added into large dither kernels by combining a base function τ with displacement maps $D_s(s, t)$ and $D_t(s, t)$ as follows:

$$\mathcal{T}(s, t) = \tau(s + D_s(s, t), t + D_t(s, t)). \quad (4.7)$$

The requirement for homogeneous distribution of texture values in the dither kernel and the resulting dither screen limits displacement maps to smooth piece-wise linear continuous functions. Mathematical functions such as $\sin(s)$ or $\sin^2(s)$ satisfy this requirement and thus can be used as one dimensional displacement maps. Also, procedural Perlin's noise [PH89] and Worley's textures [Wor96] are examples of two dimensional displacement maps. Examples of textures produced by displacement of the previous base functions are presented in Figure 4.4.

4.3.3 Hybrid dither textures

The previous section discussed examples of procedural textures generated by simple linear ramps. Thus, the resulting halftoned images resemble line art techniques. Unfortunately, construction of procedural textures of more complicated shapes is a

challenging task. It is easier to use a photograph of a brick wall than to write a program that generates an image of this wall procedurally. Therefore, it is advantageous to combine procedural and image based textures into one dither kernel.

Motivation to use multiple textures for tone reproduction can be found in traditional illustration techniques. The need for multiple textures in the context of pen and ink illustrations was discussed by Winkenbach and Salesin [WS94]. They pointed out that textures in traditional illustration techniques serve a dual role: to present the information about the surface structure and to approximate shading of the same surface. For this purpose Winkenbach and Salesin developed procedural textures that combine long “shading” pen strokes with short pen stipples to illustrate materials (e.g. bricks).

I combine multiple textures into one dither kernel by a simple mathematical operation. Two dither kernels \mathcal{T}_1 and \mathcal{T}_2 can be *added* into one as follows:

$$\mathcal{T}(s, t) = \frac{\mathcal{T}_1(s, t) + \alpha \mathcal{T}_2(s, t)}{1 + \alpha} \quad (4.8)$$

Unfortunately, the addition of two textures does not preserve the overall homogeneous and uniform distribution of threshold values of the original textures¹. Thus, in general, additional post-processing of the resulting texture with the AHE algorithm may be needed. However, in practice, addition of textures with different scales of features produces acceptable dither kernels.

The other approach is to vary textures across a dynamic range of the image. Thus, two dither kernels \mathcal{T}_1 and \mathcal{T}_2 are *fused* together using the following formula:

$$\mathcal{T}(s, t) = \mathcal{T}_1(s, t) \star \mathcal{T}_2(s, t) = \begin{cases} \mathcal{T}_1(s, t) & \text{if } \mathcal{T}_1(s, t) \leq I \\ \mathcal{T}_2(s, t) & \text{otherwise} \end{cases} \quad (4.9)$$

where $0 < I < 1$ is the threshold intensity at which only one texture is visible. Similarly to texture addition, texture fusion may not preserve the essential properties of the dither matrix. Moreover, texture fusion is not a symmetric operation and the result is dependent on the ordering of the input textures. (Figure 4.5, top).

I found that merging multiple textures is essential for approximation of many traditional illustration techniques. For example painterly effects are achieved by combining oil brush strokes and canvas texture. Cross-hatching can be simulated by merging two fragments of a scanned pencil art work (Figure 4.5, center).

¹For example, textures $\mathcal{T}_1(s, t) = s$ and $\mathcal{T}_2(s, t) = 1 - s$ cancel themselves out when added with $\alpha = 1$.

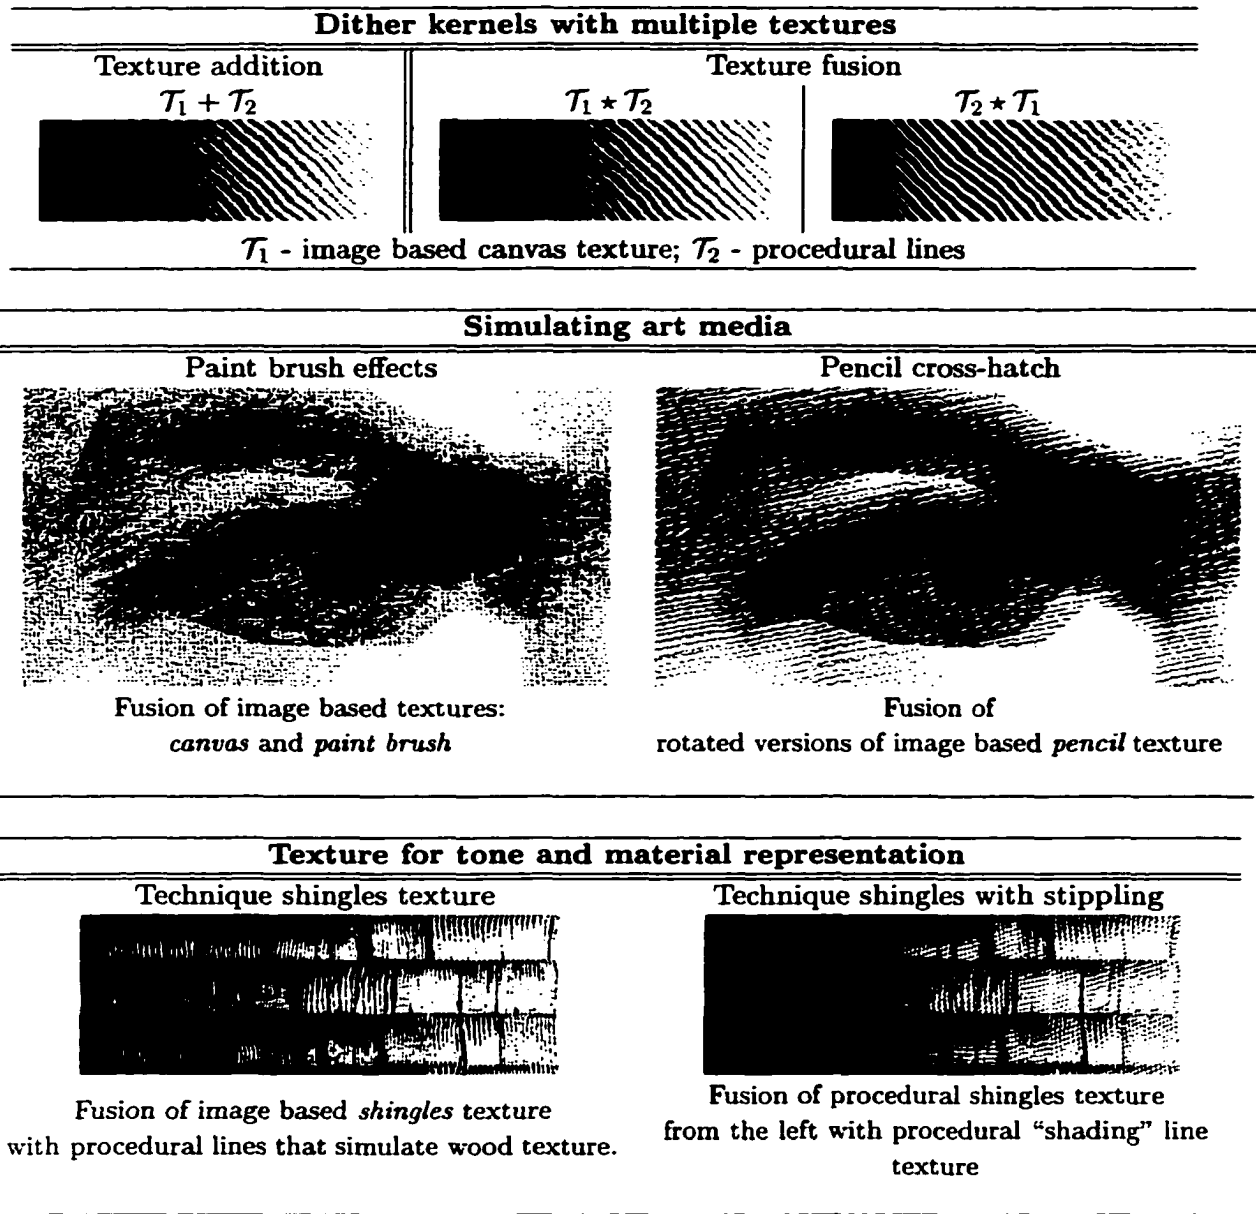


Figure 4.5: Texture addition and fusion operations combine multiple textures into one dither kernel. Thus, hybrid dithering textures are created (top). Possible applications of hybrid textures include simulation of art materials and techniques such as paint brush, and pencil cross hatching. (center). Similarly to traditional art techniques, hybrid textures enable us to represent tone and materials with different textures (bottom).

To summarize, hybrid texture techniques allows us to achieve effects similar to the results of the previous NPR research (e.g. pen and ink [WS94]). Unlike this previous research, the tone reproduction algorithm is independent from the texturing effect — the same ordered dither algorithm is applied regardless of the number of textures used in the threshold matrix. Moreover, the texture control approach is not limited to a particular illustration medium. For example, the same technique can be used to imitate pencil drawing and oil painting.

4.4 Control of texture contrast

Previous sections presented an algorithm that enables us to introduce a desired texture into the halftoned image. In this section I discuss a method to control contrast and dot clustering of this texture.

The main advantage of ordered dithering is that the distribution of pixels is defined a priori. Unfortunately when a bad approximation for a pixel is generated there is no way to compensate for the error. This means that the structure of the dither screen is clearly visible. Error diffusion techniques, on the other hand, attempt to compensate for quantization errors by propagating the error to unprocessed neighbor regions. The advantages of using error diffusion in combination with ordered dither were studied in the past [BHB83, Uli87] and were discussed in Section 2.2.2. This combination not only improves the accuracy of tone reproduction but also reduces texture artifacts of the ordered dither.

I use a partial error diffusion [EK91] method to control the contrast of the halftoning texture. The parameter α where $\alpha \in [0, 1]$ determines the influence of approximation errors diffused from the neighbor regions on the output value of the current pixel. The input pixel $g_{i,j}$ is set to 0 or 1 using the following formula:

$$b_{i,j} = \begin{cases} 0 & \text{if } g_{i,j} + \alpha E_{i,j} \leq t_{i,j} \\ 1 & \text{otherwise} \end{cases} \quad (4.10)$$

where $E_{i,j}$ is the sum of errors diffused into the current pixel and $t_{i,j}$ is the corresponding threshold value from the dither screen. While any error diffusion technique can be used in this application, I implemented Floyd-Steinberg algorithm [FS76] with serpentine processing of rows. Halftoning results for various values of the α parameter are presented in Figure 4.6. Maximal error diffusion corresponds to $\alpha = 1$ and pro-




Figure 4.6: A combination of ordered dither with parameterized error diffusion allows as to control the appearance of halftoning texture. By increasing the amount of errors diffused into a neighboring region (parameter α) I improve the accuracy of tone reproduction and reduce texturing effects. Dot clustering and improved preservation of texture features for large α is achieved by processing image by pixel blocks. However, the use of large (e.g. 4x4 blocks) in the error diffusion process may introduce block artifacts into the halftoning texture. Test images are printed at 150 dpi.

duces the softest textures. I found that in most cases $\alpha \in (0.2, 0.8)$ is a good choice for this parameter.

This error diffusion algorithm can be generalized to process images in pixel blocks. In this case a sum of quantization errors is computed over a block of pixels and is distributed to the neighbor blocks. The output pixel is determined using the accumulated error E weighted by the number of pixels in the processing block. It turns out that processing images in a block fashion leads to clustering of dots in the resulting halftoning. Dot clustering is often desired to compensate for dot overlap in laser printers. Also, I found that error diffusion over pixel blocks preserves small texture features for large values of α (Figure 4.6).

4.5 Discussion

This chapter presented the fundamentals of texture control in halftoned images. My approach is to use ordered dithering to define the shape of textures in the halftoning layer. The strength of the texture is controlled by parameters of the error diffusion process. Thus, the main goal of this research is to generate a dither matrix that produces a desired texture in the output image. The following approaches are developed:

- an image precessing technique that converts an image into a dither matrix;
- a procedural texture approach that defines a class of functions with the properties of a dither matrix;
- a hybrid approach enables us to combine multiple textures into one dither matrix.


Application of the texture control technique leads to introduction of a variety of special effects. Examples include embossing images with another image, texture, or text (Figure 4.7); approximation of traditional art media: oil brush painting, carving, etching, pencil drawing; imitation of art illustration techniques: dry brush painting, cross-hatching (Figure 4.5).

Texture control approach compares well with the current results in non-photorealistic rendering:

1. My approach is not limited to any particular style or art medium.
2. In other NPR techniques an artist must design the elements of the halftoning layer (procedural texture in pen and ink [WS94]; screen elements in artistic

Embossing images with textures

An image Mandril



Stone texture



Text SIGGRAPH



Figure 4.7: Texture control technique enables introduction of a variety of artistic effects. An image can be embossed by another image (top left), a texture (bottom left), or text (right).

screening [OH95]). In this approach the texture does not need to be designed by an artist, but can be extracted from an arbitrary image.

3. The use of different textures in the halftoning layer enables us to reproduce a number of effects generated by the previous techniques: Figure 4.5, cross hatching and multiple textures of pen and ink [WS94]; Figure 4.6, texturing with text of artistic screening [OH95].
4. The control of texture strength and variation of the accuracy of tone reproduction with error diffusion was not investigated before in the NPR literature.

Previous non-photorealistic rendering techniques have a number of advantages over the described texture control system. In particular, Ostromoukhov and Hersch [OH95] use multiple screen elements to represent different tone ranges. Therefore, their artistic rendering system allows for more variation of textures across a dynamic range of the image. Since, a single threshold screen is used in this research, texture variation is limited to a few strongly correlated textures.

Line based nature of pen and ink textures [WS94] allows designers to fine tune individual pen strokes. Procedural textures investigated in this thesis are area based and thus do not allow for such fine control.

The following chapter extends the use of texture control to interpretive rendering of photographs and three-dimensional scenes. I apply the image based approach to the design of dither screens to photorealistic halftoning in Chapter 6. It turns out that a dither screen generated from the image to be reproduced enhances the display of the original edges and textures.

Chapter 5

Texture Based Interpretive Halftoning

The research discussed thus far enabled the introduction of textures into halftoned images. In this chapter I explore the use of the halftoning texture for the enhancement of information display. My approach is based on the control of local texture features by external parameters. One approach is to store these parameters as control images representing external information including three-dimensional scene data or the user's interpretation of the display. The other approach is to control the halftoning texture through the scene description and to apply previous texture mapping techniques. Thus, this chapter presents image based and scene based approaches to the control of local texture features and applies these techniques to the halftoning of photographs and computer generated scenes.

5.1 Application of texture to interpretive halftoning.

Texture is an important channel of visual information. While texture has many visual dimensions this thesis deals with the control of the following texture features:

- **Tone**

Texture tone is defined as an average texture intensity in a region. In this research control of texture tone is fundamental to reproduction of continuous tone images. I use ordered dithering and error diffusion to approximate a target tone.

- **Shape**

Texture shape is the overall appearance of a particular texture. Examples include material textures: stone, wood, canvas, water; or art media textures: pencil

strokes, lines, dots. I define the shape of the halftoning texture by incorporating this texture into a dither matrix. The previous chapter discussed image based and procedural approaches to the dither matrix design.

- **Contrast**

I define texture contrast as the strength or visibility of texture. The control of texture contrast is based on the property of the error diffusion algorithm to ameliorate halftoning textures.

- **Scale**

Texture scale is the size of prominent texture features: the scale of bricks in the brick texture; or the distance between lines that represent the same tone in the case of line textures.

- **Direction**

Texture direction is the principal direction of texture features: direction of bricks in the brick texture; or the main direction of lines in the line textures.

The previous chapter presented the framework for the control of the above texture features. In particular, the halftoning texture $T(u, v)$ is a function composed of the kernel texture $\mathcal{T}(s, t)$ and the mapping function $\mathcal{M}(u, v)$. Thus, the texture features are controlled by either a kernel function, a mapping function, or the error diffusion process. That is:

- Texture shape is controlled by the texture kernel $\mathcal{T}(s, t)$.
- Texture scale and direction are controlled by the mapping function $\mathcal{M}(u, v)$.
- Tone and contrast are defined by the parameters of the error diffusion algorithm.

In this chapter I discuss tools and techniques that adjust local texture features according to the external information. The local control of texture features enables the use of the halftoning layer as an expressive mean in information display. The enhancement of visual display is facilitated by the user defined mapping between texture features and the information. I investigate texture control based on the following information:

2-D information All the information that can be derived from an image using some image processing technique: image gradient, contrast, edges, etc.

3-D information All the information that can be derived from a scene description. Examples include object identifiers, surface normals, distance from the viewer, scene illumination.

Subjective information Subjective information is defined by the user's understanding and analysis of an image or a scene and can be expressed as importance, completeness, focus.

I study two ways to represent information: as an image or as a scene.

In the image based approach information is stored in a set of auxiliary images. A display system is using pixel values of these images to control parameters of the rendering process. An image based approach was taken by Saito and Takahashi [ST90], Haeberly [Hae90], Landsdown and Schoffield [LS95]. I study the image based approach in the following section and demonstrate its use in the halftoning of photographs (Section 5.3). Further in the chapter, I focus my research on the rendering of 3-D scenes. I store the scene information in g -buffers [ST90] and use them to control the halftoning texture.

The scene based approach is limited to the rendering of three-dimensional scenes. The scene information is represented in the scene description language and includes object geometries and positions of objects, camera, and light sources. Examples of previous research that follows a scene based approach are comprehensive wire-frame display [ARS79, KK87, DC90] and sketch rendering system of Strothotte et. al. [SPR+94]. I extend the scene description to include control of the halftoning texture and apply previous texture mapping techniques to generate a dither screen. I discuss the advantages and disadvantages of the image and scene based approaches and combine these techniques in rendering of a single scene.

5.2 Image based control of texture features

In this section I focus my research on the techniques to control texture features by local parameters. Because these parameters are stored in auxiliary control images I refer to this approach as *image based*. The user of the halftoning system may store external information in the control images thus adapting texture features to this information. However, the texture control techniques are independent from the type of information being represented by the auxiliary image. Therefore, for the purposes of this discussion the auxiliary image is just a two-dimensional function $C(u, v)$ in the normalized coordinate space (u, v) .




Figure 5.1: The auxiliary image controls the amount of error diffusion resulting in variable texture contrast across the halftoning. For example, the control of texture contrast by the circular ramp with the maximum at the center produces soft textures in the center of the halftoning (left image). A “fog”-like effect is achieved by controlling the amount of error diffusion in the image on the right.

5.2.1 Texture tone and contrast

The halftoning system controls texture tone and contrast by using a combination of ordered dithering with error diffusion. I employ the property of the error diffusion process to reduce the contrast of the halftoning texture. Because the use of error diffusion improves approximation of the original image tones, tone and contrast of the texture are controlled simultaneously.

The local control of tone and contrast is achieved by making error diffusion dependent on the local parameter. As a result Equation 4.10 is modified to:

$$b_{i,j} = \begin{cases} 0 & \text{if } g_{i,j} + \alpha C_{(i,j)} E_{i,j} \leq t_{i,j} \\ 1 & \text{otherwise} \end{cases} \quad (5.1)$$

where $C_{i,j}$ is the corresponding value of the control image. An example of applying locally controlled error diffusion to halftoning with a line-based texture is found in Figure 5.1. The texture contrast is controlled by the circular ramp with the peak at the center. The resulting texture contrast is decreasing with the distance from the center. Also, when compared to the top and the bottom of the image, the same gray scale tone is best approximated in the center of the image.



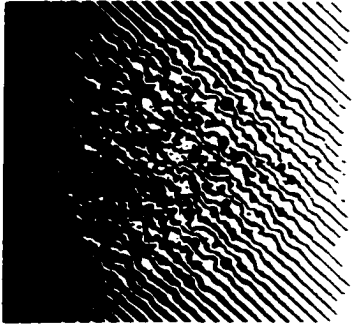
Image based control of texture shape	
Image base textures	Procedural textures
 <p>Auxiliary image</p> 	

Figure 5.2: Texture shape is controlled by modifying the dither kernel. In the case of the image based kernels texture shape cannot be continuously changed. An auxiliary image (top left) is thus a map that indicates what texture should be used for any given pixel. As a result, three textures — lines, leather, and rock — are used to halftone the lily image (bottom left). Procedural textures allow for smooth variation of their shape. The amplitude of texture distortion is determined by a circular ramp with maximum at the center. The waviness of the resulting dither texture is the highest in the center of the image (right).

5.2.2 Texture shape

The shape of the halftoning texture is defined by the arrangement of threshold values in the dither screen. The previous chapter presented image based and procedural approaches to the design of textured dither screens. The objective here is to control texture shape by an auxiliary image. However, the amount of control is dependent on the type of the dithering texture. Due to the fundamental limitations of the image based textures, their local control is limited. Unlike procedural textures, image based textures cannot be parameterized making impossible smooth changes of their shape.

The abrupt changes of texture shape are facilitated by piece-wise constant control images that define a particular texture shape from the set of available textures. This technique is applicable to both image based and procedural dither screens and results in the use of multiple textures for halftoning of a single image. Thus, a photograph of a lily is displayed with three different textures: rocks at the background, leather for the small leaves, and line texture for the flower (Figure 5.2, bottom left).




Figure 5.3: Definition of texture scale is not sufficient to control the appearance of texture. Thus, these images demonstrate three ways to scale the same texture from left to right. Moreover, the continuous change of texture scale leads to the change in texture direction making the control of scale and direction to be dependent on each other.

Procedural textures allow for smooth shape variations by making a dither kernel $\mathcal{T}(s, t)$ dependent on the auxiliary image. In the case of dither kernels produced by a combination of a base texture function $\tau(s, t)$ and two displacement maps (Equation 4.7) the auxiliary image controls the displacement amplitude as follows:

$$\mathcal{T}(s, t) = \tau(s + \mathcal{C}(s, t)D_s(s, t), t + \mathcal{C}(s, t)D_t(s, t)) \quad (5.2)$$

where $\mathcal{C}(s, t)$ is the control image value in the normalized coordinate space (s, t) . For example, a circular gray ramp with the maximum at the center varies the amplitude of Perlin's noise texture (Figure 5.2). The waviness of the line texture varies smoothly across the image and is the highest at the center.

5.2.3 Texture direction and scale

The control of texture direction and scale are facilitated by the mapping function $\mathcal{M}(u, v)$. Examples of mapping functions for constant scale and rotation are found in Section 4.1. The objective of this research is to control the mapping function locally. However, smooth changes of texture scale and direction are hard to achieve.

Consider the problem of continuously reducing the scale of the line texture from left to right. Figure 5.3 presents three different solutions to this problem. This example demonstrates the following problems:




Figure 5.4: Texture scale is defined by the mapping function (Equation 5.3) and is controlled by the auxiliary image. The control is limited to the locally constant texture scale produced by piece-wise constant auxiliary images (image on the left). The use of the mapping function to vary the scale continuously may result in undesirable texture deformations (image on the right).

- Definition of the local texture scale does not produce a unique solution and is not sufficient for the overall control of the mapping function.
- The continuous change of texture scale leads to the change in texture direction. Therefore, scale and direction are dependent and cannot be controlled by uncorrelated auxiliary images.

The dependence of the scale and direction can be explained mathematically as follows. The continuous change of the texture scale corresponds to the non-zero derivative of the mapping function. However, texture direction is also dependent on the derivatives of the mapping function and is found as $\tan^{-1}(\mathcal{M}'_u/\mathcal{M}'_v)$. Therefore, continuously changing scale leads to the change of texture direction.

The problem of defining a general mapping function that allows for continuous control of texture direction and scale is beyond the scope of this thesis. Thus, I limit the problem to the piece-wise constant control parameters resulting in texture discontinuities. The use of such parameters enables independent control of texture direction and scale (Figure 5.5, image on the right).

I control texture scale locally by modifying the scale mapping \mathcal{M}_s and augmenting

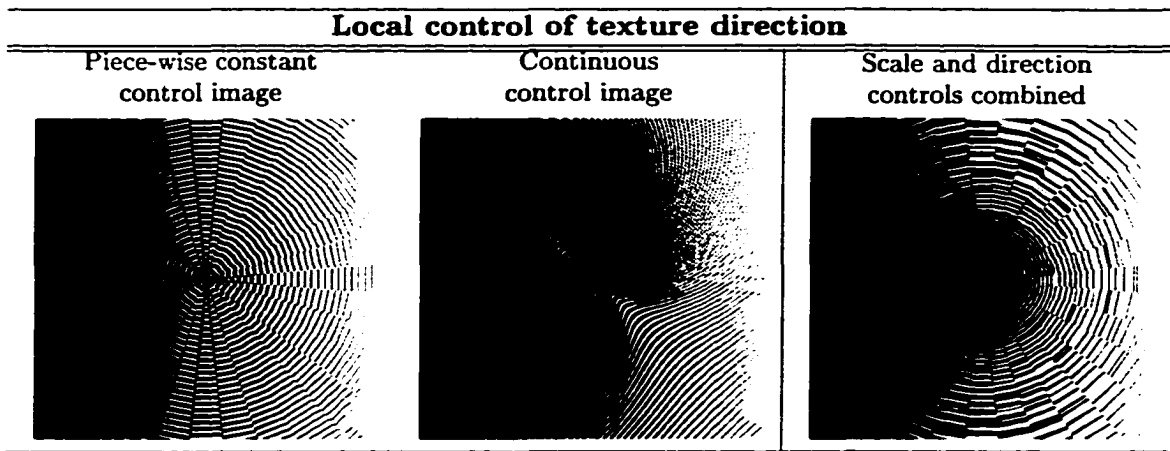


Figure 5.5: Texture scale is defined by the mapping function (Equation 5.4) and is controlled by the auxiliary image. The control is limited to the locally constant texture direction (image on the left) and fails to change the texture direction continuously (image in the center). While in general the control of scale and direction are dependent on each other, the use of piece-wise constant parameters enables the independent control of scale and direction. Thus, the circular change of texture direction is complemented by radial texture scaling (image on the right).

scale factors s_u and s_v with the corresponding value of the auxiliary image $C(u, v)$.

$$T(u, v) = \mathcal{M}_s(u, v) \circ \mathcal{T} = \mathcal{T}(C(u, v)s_u u, C(u, v)s_v v) \quad (5.3)$$

This local scale mapping $\mathcal{M}_s(u, v)$ works reliably only for the piecewise constant control images but may produce scaling artifacts in the case of continuously changing parameters. As an example, I attempted to control texture scale using a continuous circular ramp with the maximum in the center (Figure 5.4). The mapping function $\mathcal{M}_s(u, v)$ produced a desired result in the case of the line texture (center image), however, failed for the cross-hatched texture (right image). The problem was corrected by quantizing a ramp to 8 values thus resulting in a piece-wise constant auxiliary image (image on the left of Figure 5.4).

The angle of the rotation mapping is controlled by the corresponding value $C(u, v)$ of the auxiliary image:

$$T(u, v) = \mathcal{M}_{C(u, v)} \circ \mathcal{T} = \mathcal{T}(\cos(C(u, v))u \sin(C(u, v))v, \sin(C(u, v))u + \cos(C(u, v))v) \quad (5.4)$$

Similarly to the scale mapping, the rotation mapping works only for a piece wise constant function $C(u, v)$ and creates unpredictable results in the case of continuously

changing control parameters (Figure 5.5).

5.2.4 Summary of the image based texture control

The discussion above was focused on controlling texture features locally by the use of a two-dimensional parameter function. This function is represented by an image and thus the techniques are referred to as “image based” texture control.

Image based techniques allow us to control texture shape, tone, contrast, direction and scale. Control of texture features is achieved by making either a dither kernel function, or a mapping function, or the error diffusion process to be dependent on the parameter function. While in general texture control techniques should accommodate any type of parameter functions, control of texture shape in image based dither screens, and the mapping function work reliably only for piece-wise constant parameters.

5.3 Interpretive halftoning of two-dimensional images

Image based control of texture features enables us to adapt the display to the external information. In this section I apply the image based techniques to the halftoning of photographs. I use the standard test image “Lena” and gradually adapt texture to various information thus increasing the complexity of the resulting halftoning. Because a two-dimensional image is displayed, the available information is limited to two-dimensional information and subjective, user defined information.

The first dither screen is the procedural line texture rotated by 60 degrees. The halftoning was produced by a combination of ordered dithering with the partial error diffusion (Equation 4.10, $\alpha = 0.3$). The resulting image (Figure 5.6, top) exhibits strong line texture. However, the halftoning texture is not adapted to the displayed information and interferes with the reproduction of image details. The rest of this section is focused on adapting the halftoning texture for the enhancement of the image display.

One approach to enhancement of image details is to incorporate these details into the dithering screen. I produced such a dithering screen by processing the original photograph with the CLAHE algorithm (Section 4.2). The resulting image (Figure 5.6, bottom right) is merged with the previous procedural line texture using Equation 4.8. Incorporating image based features into the dither screen leads to increased



Dither screen is a procedural line texture rotated to 60 degrees. Halftoning with error diffusion and $\alpha = 0.3$.



Dither screen is adapted to the image by merging a procedural line texture with the image based threshold matrix above. The threshold matrix was produced by CLAHE $w = 4$, $c = 0.25\%$.

Figure 5.6: Interpretive halftoning of the “Lena” image. Part 1

image contrast and improved rendering of edges. (Figure 5.6). In the next examples I modify only the line texture and add the image based texture to generate the final dither screen.

The line texture used thus far is uniform across the image. The display of image features is enhanced if direction of the line texture is adapted to the image. I take the approach presented by Haerberli in [Hae90] and align texture direction to the direction of the image gradient. The image gradient is computed from a smoothed out version of the original image. Image smoothing is needed to produce a slow changing gradient image and to avoid the influence of the original image textures and noise. Directions of the image gradient are mapped linearly to gray scale intensities ($0 = 0$ and $\pi = 1$) and stored in the auxiliary image (Figure 5.7, top right). The resulting gradient image quantized to 16 intensity levels is used to control the texture mapping function $\mathcal{M}_{C(u,v)}$ (Section 5.2.3). Thus, the line texture is perpendicular to the direction of the local image gradient enhancing display of the image details (Figure 5.7, top).

Information used thus far was extracted from the photograph by image processing operations. Dither screens adapted to such information can be generated without intervention of the user. However, the subjective image information is user dependent and is the result of user's interpretation of the image. I demonstrate the use of subjective information by using an importance map found at the bottom right of Figure 5.7. According to the importance map, viewer's attention should be focused on Lena's face and shoulder. I attempt to achieve this goal by using the importance map to control texture scale and shape. Thus, texture scale is inversely proportional to the importance of the image regions and is the largest at the background (Figure 5.7, bottom). Similarly, line waviness produced by Perlin's noise (Equation 5.2) is the largest at the background (Figure 5.8, top).

Subjective information can be combined with the objective two-dimensional information into one control image. I combined the importance map with the magnitude of the image gradient to control the contrast of the halftoning texture (Figure 5.8, bottom). The use of the resulting auxiliary image varied the amount of error diffusion and reduced texture contrast at the face and at the background edges of the image.

Overall, this section demonstrated application of image based texture control techniques to interpretive halftoning. The display of a photograph was modified by mapping image derived information and subjective user interpretation to the shape, scale,



Line texture is aligned with the directions of image gradient and merged with image based threshold matrix. Image gradient is computed from a smoothed out version of the original image. To avoid directional artifacts gradient directions are quantized to 16 angles (image above).



Line texture is aligned with the image gradient and scaled according to the map above. Black represents no scaling and white corresponds to maximum magnification of 2.5. Thus, textures on the background are rougher than on the foreground.

Figure 5.7: Interpretive halftoning of the "Lena" image. Part 2



Line texture is perturbed with Perlin's noise. Noise amplitude and the texture scale are controlled by the image above.



The image is halftoned with locally adjusted error diffusion. The amount of error diffusion (parameter α) is controlled by the image above. Black corresponds to $\alpha = 0$, and white to $\alpha = 0.5$. Dither matrix is the procedural line texture rotated to the gradient with scale and noise control.

Figure 5.8: Interpretive halftoning of the "Lena" image. Part 3

direction, and contrast of the halftoning texture. The use of the photograph limits the amount of information available to the user and to the rendering system. Thus, interpretive halftoning of two-dimensional images heavily depends on user's interaction with the rendering system and user's analysis of the displayed image.

5.4 Interpretive Halftoning of 3-D Scenes

This section is devoted to the interpretive halftoning of 3-D scenes. The goal of this study is to enhance scene the display using halftoning texture. In particular, the texture is used to enhance the depiction of objects' geometries, their relative positions, and the scene illumination. Also, image display is adapted to the subjective user defined information to focus viewer's attention on curtain objects or their features.

The first approach is scene based and relies on the extension of the 3-D rendering system to produce a dither screen. The second approach is to store scene information in g -buffers [ST90] and to use these buffers to control image based interpretive halftoning. I discuss the advantages and limitations of both approaches and demonstrate the use of a combination of both techniques to the rendering of 3-D scenes.

5.4.1 Scene based approach to the halftoning of 3-D scenes

Most current computer graphics systems use a scene description language to define object geometries, their optical properties and locations. Photorealistic rendering systems display this scene by approximating light distribution in the environment. The scene based comprehensive rendering techniques use object descriptions to control graphics primitives in the resulting display [ARS79, KK87, SPR⁺94].

In this research I follow a hybrid approach. I extent Persistence of Vision raytracer (PoV) to produce both the photorealistic display of the scene and the dither screen for the halftoning of this scene. Thus, the scene description language is extended to define not only the object surfaces but also the halftoning texture of the final display. Because the traditional rendering system is used to generate a dither screen the direction and scale of the halftoning texture are controlled by the standard texture mapping techniques. In particular, I explored the use of two-dimensional and volume texture mapping methods.

2D texture mapping of dither matrices




Figure 5.9: The dither matrix is constructed by merging an image based brick texture with a procedural line texture (left). When this dither matrix is textured mapped onto the rotated plane the texture becomes scaled to reflect the perspective distortion (center image). However, scaling leads to aliasing of the line texture at the far end of the plane. Distortions of the line texture is avoided by projecting only the image based component of the dither matrix (image on the right). Thus the line texture is generated using the projected image based texture and then merged with it to produce the final projected dither matrix.

Two-dimensional texture mapping

Two-dimensional texture mapping is based on the assignment of two-dimensional texture coordinates to the surface points of the object. These texture coordinates are projected onto the viewing plane together with the object geometry. Thus texture scale and direction are adjusted to the curvature of the object's surface and to the angle between the surface and the viewing plane. I use this property of two-dimensional texture mapping to construct a dither screen for halftoning of a plane (Figure 5.4.1).

The plane in Figure 5.4.1 is illuminated from the top and is rendered with two positions of the camera: straight-on close (left image) and distant angled at 50° to the surface (images at the center and right). The halftoning texture is a combination of the image based brick texture with the procedural line texture. The line texture was generated using image based techniques and is controlled by auxiliary images correlated with the brick texture. The dither screen is produced by texture mapping the dither kernel to the surface of the plane.

The center image demonstrates the scaling and the change of direction of the halftoning texture according to the viewing perspective of the image. However, scaling of the line texture leads to texture aliasing and to the distortion of the overall texture appearance. Moreover, due to the scaling and texture filtering that is necessary to

reduce aliasing, the distribution of threshold values in the resulting dither screen does not satisfy the uniformity and homogeneity properties of good dither screens. This problem of line scaling and aliasing in the context of pen and ink rendering was cited by Winkenbach and Salesin [WS94, WS96] and was solved by the use of planar maps. My solution in the context of interpretive halftoning is presented on the right of Figure 5.4.1. In this case only the brick texture and the auxiliary images that control line texture were mapped to the plane. The line texture is produced using the projected auxiliary images and is merged with the projected brick texture. This approach to dither screen design provides visual cues of object's scale and rotation while avoiding scaling and aliasing of the line texture.

Overall, two-dimensional texture mapping is an effective approach to the design of dither screens that are aligned with the surfaces of displayed objects. However, this technique may lead to undesired scaling and aliasing of the halftoning texture. These artifacts are avoided in the case of procedural textures by limiting the texture mapping to the control images and by generating the final dither screen with the previous image based techniques.

Two-dimensional texture mapping requires the assignment of texture coordinates to the object surfaces. The problem of minimal distortion texture mapping is the objective of recent research in computer graphics [BVG91, LM98]. These current solutions are applicable to polygonal models and are computationally intensive. Thus, the use of two-dimensional texture mapping is limited to objects with the simple correspondence between texture coordinates and the surface.

Volume textures

The problem of texture coordinate assignment is avoided by the use of volume (or three-dimensional) textures [PH89, EMP⁺94]. Volume textures were used in comprehensive rendering by Saito and Takahashi [ST90] and for computer generated copper plates by Leister [Lei94]. Moreover, Haeberli [Rob97] used volume textures for 3-D halftoning. My research extends Haeberli's approach and applies the techniques for the dither screen design to volume textures.

Similarly to Leister [Lei94] I define volume textures as three-dimensional functions that depend on the distance from some plane. In my research I use one-dimensional procedural textures $\tau(s)$ (Equations 4.4,4.5) where s is the distance to the plane. The

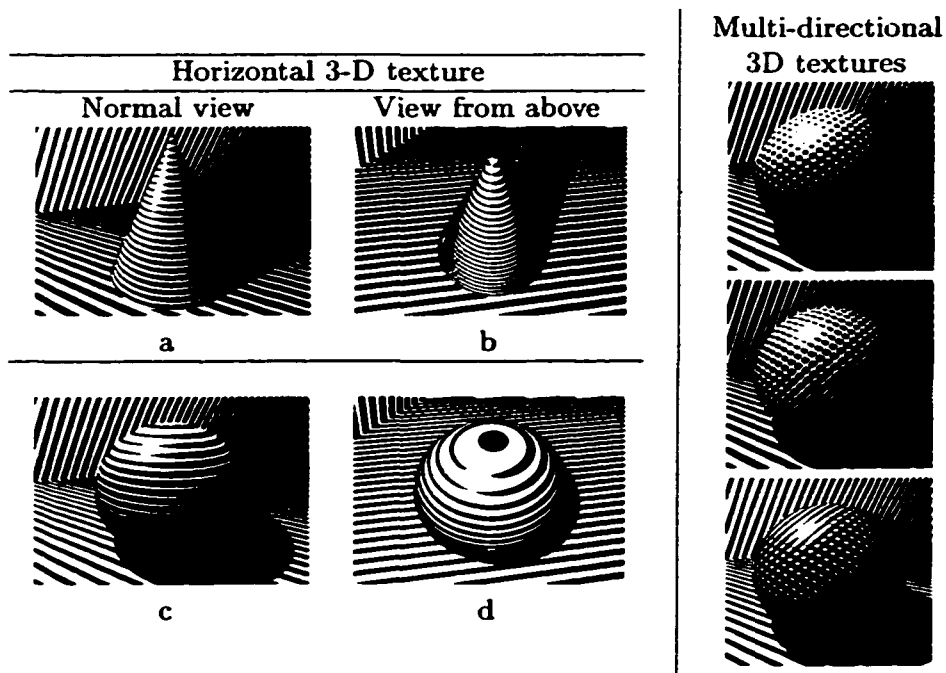


Figure 5.10: 3-D textures are used to produce a dither screen. The change in texture direction and scale reveals the shape of the object. However, unidirectional 3-D textures (e.g. horizontal texture at the left) may also create severe tone distortions (example d, black stop in the bright area of the sphere). Thus the alignment of 3-D textures is closely related to the curvature of the object and the camera position. Multi-directional are created by merging two unidirectional textures and are less dependent on the object curvature.

value of the threshold in the dither screen is thus depends on the texture function $\tau(s)$ and the distance of the corresponding surface point to the specified plane. Examples of dither screens produced by volume textures are presented in Figure 5.10. The dither screens for the sphere and cone are generated by a horizontally aligned texture. The use of this texture depicts the curvature of these objects and the perspective of the image. The volume texture on the scene's floor and the wall are scaled automatically providing a visual cue of the distance from the camera.

The use of dither screens with volume textures is a simple and powerful approach to interpretive halftoning. However, it is often hard to control the scale of the volume texture. Unlike two-dimensional texture mapping, the scale of the volume texture depends on the direction of the texture, on the position of the camera, and on the curvature of the object. Compare the examples *a,c* to *b,d* in Figure 5.10. The same volume textures are used to display the scene viewed by two different cameras. In the case of the sphere the change of the viewing position revealed the artifact of volume

texturing. Unlike the texture on the cone (image *b*) the lines on the surface of the sphere are spaced unevenly (image *d*). Moreover, the black spot on the top of the sphere is the result of the uneven scaling and the cause of shading distortions in the image. In general, texture scale increases when the surface becomes parallel to the plane that controls the texture (horizontal plane in this case). Thus, ideally the angle between the control plane and the object's surface projected onto the viewing plane should be nearly constant. However, such a direction is hard to find for an arbitrary object.

The dependency of the texture scale on the direction of the control plane can be reduced by the use of multi-directional textures. Two-directional textures are used in halftoning of spheres at the right of Figure 5.10. These dither screens are generated by two volume textures aligned at an angle to each other. The resulting halftoning approximates stippling and cross-hatching techniques of traditional art media and minimizes the non-uniform scaling of texture.



Discussion of the scene based approaches

Two-dimensional and volume texture mapping are two examples of the scene based approach to the interpretive halftoning of three-dimensional scenes. I extended the scene description language to include information about the halftoning texture. Moreover, the raytracing software was used to render both a photorealistic tone image and a dithering screen.

The advantage of the scene based approach is the ability to use previous texture mapping techniques to design the dither screen. Thus, the scale and direction of the halftoning texture are controlled automatically by the curvature of the object and its position in the scene. However, the use of traditional texture mapping also limits the control over these texture features. In particular, excessive texture scaling may lead to aliasing and modification of the threshold distribution.


5.4.2 Image based approach to the halftoning of 3-D scenes

Unlike the scene based approaches, the image based technique does not use a three-dimensional texture mapping to produce a dither screen. In the image based approach the dither screen is constructed by two-dimensional texture control techniques (Section 5.2). These techniques account for the three-dimensional scene information by the use



id-buffer

Unique object identifiers are stored in the *id*-buffer (above). The buffer is used to control texture waviness scale and direction resulting in the unique texture for display of every object (left image).



n-buffer

Direction of object normals projected on the viewing plane are stored in the *n*-buffer. The buffer controls direction of the "pencil" texture (left image).

Figure 5.11: Image based 3-D halftoning using *id*-buffer and *n*-buffer

of *g*-buffers [ST90]. In this research the *g*-buffers are created by the PoV raytracer as an intermediate step in the photorealistic rendering of the scene. The raytracer produces the following *g*-buffers:

- *id*-buffer

A unique identifier is assigned to every object in the scene. These identifiers are stored in the *id*-buffer to separate image regions that depict different objects of the scene.

- *n*-buffer

Pixels of the *n*-buffer store the surface information of the corresponding objects. In this research the normal of the object is projected onto the viewing plane. The direction of this projection is mapped to the gray scale values and is stored

in the n -buffer.

- **l -buffer**

The scene illumination information is stored in the l -buffer. This buffer differentiates image regions illuminated by direct and diffused light. Separate l -buffers are constructed for every light source in the scene.

- **z -buffer and its derivatives**


The distance between the camera and the object's surface is stored in the z -buffer. Saito and Takahashi [ST90] pointed out that the first and second derivatives of the z -buffer represent silhouettes and ridges of the objects. Thus, the z -buffer is produced by the raytracer directly. The silhouettes z' -buffer and the ridges z'' -buffer are computed afterwards by taking the derivatives of the saved z -buffer.

I demonstrate the use of the g -buffers by halftoning a scene with the simple geometric objects: cylinder, sphere, and box. (Figure 5.11, 5.12, 5.13). The id -buffer controls texture scale, waviness, and direction in the top image of Figure 5.11. Thus, every object is visually separated by different textures in the resulting halftoning. However, the use of the id -buffer alone does not enhance the display of object geometries. To achieve this goal Haeberli in [Hae90] suggested the alignment of display texture to the projection of the normals onto the viewing plane. Thus, I control the direction of the pencil stroke texture with the n -buffer in the bottom image on Figure 5.11. In this example the halftoning texture helps to differentiate between surfaces and reveals objects' curvature.

The viewer's understanding of the scene illumination can be enhanced by segmenting shadows and directly lit image regions. I use the l -buffer to change the shape of the halftoning texture in the shadow of one of the light sources (Figure 5.12, top). Thus, the object surfaces that are not illuminated by this light source are unified by the halftoning texture.


The fog effects were often used in computer graphics to reduce the contrast of the distant objects in the scene thus enhancing the appreciation of the scene's depth. Similarly, I reduce the contrast of the halftoning texture that displays the distant scene surfaces. This effect is achieved by increasing the α parameter of the error diffusion process proportionally to the z -buffer (Figure 5.12 bottom).

The other example of the use of the z -buffer is presented in Figure 5.13. In this case the first and second derivatives of the z -buffer are merged into the previous



l-buffer

The *l*-buffer identifies direct light (white and gray) and shadows (black) created by one of the light sources. The buffer modifies halftoning texture in the shadow thus unifying image areas with the similar illumination.



z-buffer

Distances from the viewing plane to scene objects are stored in the *z*-buffer. The buffer controls the amount of error diffusion producing softer textures on the distant surfaces and high contrast textures on the close surfaces.

Figure 5.12: Image based 3-D halftoning using *l*-buffer and *z*-buffer

dithering screen. Incorporating z' and z'' -buffers into the halftoning texture results in the enhancement of object silhouettes and ridges.

Image based rendering techniques lend themselves to the incorporation of the subjective information into the scene display. The subjective information is presented in the form of the importance map and is used in the combination with the previous *g*-buffers. Consider the example of the foot bones rendering (Figure 5.14). The image map assigns the highest importance values to the front two bones of the model. The display is adapted to this importance map by using different *g*-buffers to control the texture. Thus, the *n*-buffer controls direction of the halftoning texture enhancing the display of curvature of the most important bones. The other bones are displayed with




Figure 5.13: First and second derivatives of the z-buffer represent silhouette edges and internal ridges of the objects (top images). These buffers are added into the dither screen generated before to delineate objects and the background.




Figure 5.14: The importance map is used to control textures in the display of the three-dimensional model of the foot bones. The most important bones is halftoned with textures aligned by the n -buffer. The id -buffer controls the shape, scale, and direction of texture for the display of other bones.

textures controlled by the id -buffer. The least important bones are halftoned by the partial error diffusion algorithm with the constant threshold. The silhouettes of all bones are enhanced using the z' -buffer.

Overall, the examples above demonstrated enhancement of the scene display by the use of the g -buffers and importance maps. The g -buffers store three-dimensional scene information and are produced by the traditional rendering system. The importance map is drawn by the user and defines the relative importance of the scene objects. The g -buffers and the importance map are assigned to control features of the halftoning texture. The texture control is facilitated by the two-dimensional methods thus avoiding aliasing and scaling artifacts of the scene based approaches.

5.4.3 Discussion of the 3-D scene based halftoning

I presented two approaches to the enhancement of the binary display of three-dimensional scenes. Both approaches use three-dimensional information to control the halftoning texture.

The first approach is based on the scene description and extends traditional photorealistic rendering system to produce a dither screen. Parameters of the texture in this

case are controlled by the previous texture mapping techniques. The advantage of this approach is thus the automatic adjustment of the texture scale and direction to the position and to the curvature of the object's surface. However, the use of three-dimensional texture mapping limits the control of texture features and may result in aliasing.

The second approach is based on two-dimensional methods of texture control that are discussed in Section 5.2. The scene information is stored by the rendering system into g -buffers and is mapped to the shape, scale, direction, and contrast of the halftoning texture. Therefore, in the image based approach texture features are not controlled automatically by the scene, but are dependent on the user specified mapping. Moreover, the use of two-dimensional texture control allows us to avoid aliasing artifacts of the scene based approach.

The scene and image based approaches can be combined to generate a single dither screen. Thus, skin and eyes surfaces of the Beethoven's bust (Figure 5.15) are rendered with two-directional volume texture. The hair, shirt, and tie are halftoned with the procedural line texture aligned by the n -buffer and scaled by the id -buffer. The suit surface is displayed by the image based canvas texture. The appreciation of the scene illumination is enhanced by unifying shadow regions with the procedural line texture. The overall texture contrast is controlled by the z -buffer. Thus, in this example volume textures are combined with procedural and image based textures adapted to the g -buffers. Image based techniques are used to control the overall assignment of textures and the overall contrast of the display.

5.5 Discussion of the interpretive halftoning

This chapter was devoted to the use of the halftoning texture for the enhancement of information display. This goal is reached by adapting texture features to the displayed information. Thus, my research was focused on the techniques that control texture shape, scale, direction, and contrast according to the specified information.

I developed the control of texture using two approaches. The first approach is image based and tailors texture features using auxiliary images. This approach is suitable for halftoning of both photographs and computer generated three-dimensional scenes. The second approach is specific to the halftoning of 3-D scenes and is based on the




Figure 5.15: Both image and scene based techniques are used in the halftoning of the model of the Beethoven's bust. The skin and the eyes are rendered with the volume textures. Textures of the hair, the shirt, and the tie are controlled by the n -buffer and are scaled using the id -buffer. The suit is displayed with the image based canvas texture. The background is the line texture controlled by the Worley's texture. The silhouettes and ridges are highlighted by incorporating the z' and z'' -buffers into the dither screen. The shadow regions are identified by the l -buffer and are displayed with the line texture.

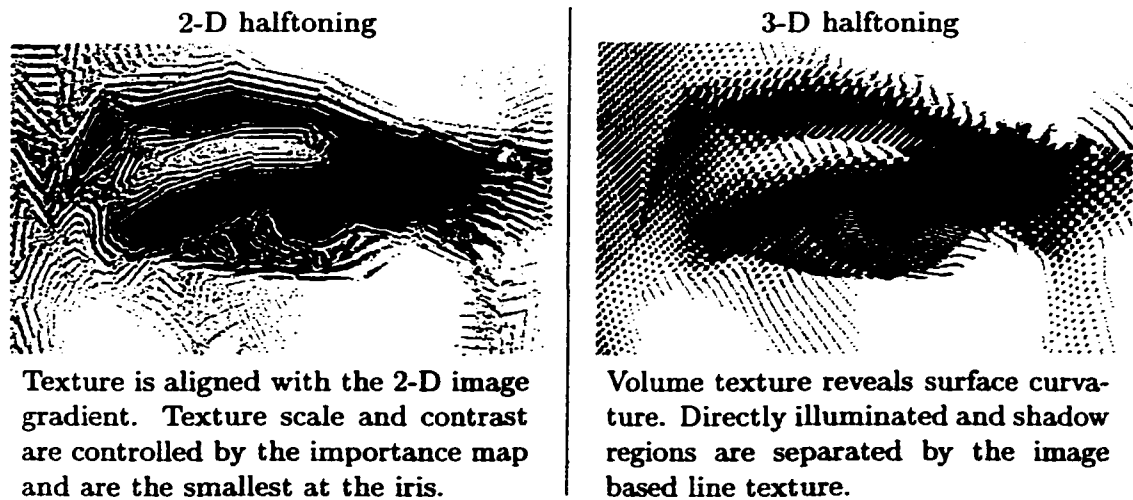


Figure 5.16: The same image is halftoned using different information to control the texture. 2-D halftoning is limited to the subjective user defined information — an importance map; and is meant to enhance the appreciation of the features already present in the image — image gradient in this example. The 3-D halftoning uses the scene information including object geometry and the illumination thus providing visual cues about external information.

scene description. In this case previous texture mapping algorithms are used to control the direction and scale of the halftoning texture.

The resulting halftoning system enables the user to adapt image halftoning to hers/his interpretation of the displayed information. The effectiveness of the produced rendering is dependent on the mapping of information to the features of the texture. The effectiveness of this mapping is the subject to the stylistic, aesthetic, cultural, or other evaluations that are beyond the scope of this thesis. Thus, the use of the external information and its mapping to the texture features are presented here only for the demonstration of the texture control techniques.

Chapter 6

Texture Control in Photorealistic Halftoning

In this chapter I apply the previously developed texture control technique to photorealistic halftoning. The goal of this investigation is to minimize the halftoning artifacts and to enhance the rendition of the original image details. In order to evaluate the results I develop a multiscale edge measure. This measure provides a tool for quantitative analysis of the halftoning texture and distortions of the original image details. I use the error measure to support my previous observations and to compare halftoned images.

6.1 Preservation of image details in halftoned images

The texture control technique (Chapter 4) enables us to define the halftoning texture. Thus far I have used this technique to introduce an external texture and to enhance information display with illustrative elements. In this chapter I take a different approach and use the texture control technique to enhance photorealistic image display. Moreover, the objective is to avoid the appearance of external textures and to preserve edges and textures present in the original image.

Unlike previous ordered dithering techniques, I do not attempt to find a perfect dither matrix suitable for the display of every image. My approach is to construct a dither screen that is specific to the reproduction of a particular image. Because I want to avoid the introduction of external textures, I construct a dither screen from the original image by processing it with the adaptive histogram equalization (AHE) algorithm (Section 4.2). The use of original image textures in the dither screen guarantees the

appearance of these textures in the halftoned image. Furthermore, I combine ordered dithering with the error diffusion algorithm to improve the reproduction of image tones and to control texture contrast. The error diffusion algorithm is controlled by two parameters. Parameter α defines the influence of approximation errors on the output of the current pixel and is discussed in relation to the thresholds used for ordered dithering (Section 6.1.1). The size of the approximation regions processed by the error diffusion algorithm effects the dot clustering properties in the resulting halftoning and is explored in Section 6.1.2.

6.1.1 Error diffusion with various thresholds

The effect of error diffusion in the combination of constant threshold and dither screens generated by AHE and CLAHE algorithms is presented in Figure 6.1. The halftoning with a constant threshold $\frac{1}{2}$ (left column) is equivalent to previous research in photorealistic tone reproduction. In particular, Floyd and Steinberg [FS76] use the constant threshold and distribute all the approximation errors to the neighbor pixels (the bottom left image). The use of the full error diffusion with the constant threshold results in the good approximation of image tones but “washes out” original image textures and edges. Eschbach and Knox [EK91] also use a constant threshold but reduce the amount of error diffusion. Partial error diffusion results in the overall increase of image contrast but does not preserve original image textures (the center image on the left, Figure 6.1).

In Chapter 4 I have shown that processing images with the adaptive histogram equalization (AHE) algorithm approximates threshold distribution properties of traditional dither screens. However, the close correspondence between the reproduced image and the dither screen may lead to the poor approximation of image tones. Clearly, because I use the dither screen generated from the image to be reproduced, the correlation between threshold values and the image is very high. This correlation explains the results of halftoning with the AHE and CLAHE dither screens (top row, center and right image, Figure 6.1). Even though these dither screens reproduce image details better than the simple thresholding (top left image), the gray scale values of the original image are poorly approximated. This flaw is resolved by the use of error diffusion.

Compare the halftoning with the CLAHE and AHE dither screens combined with full error diffusion (bottom row, center and right images, Figure 6.1). As expected,

**The effect of error diffusion
on halftoning with various dither screens.**



Figure 6.1: The effect of error diffusion depends on the thresholds used for halftoning. The left column corresponds to the constant threshold and demonstrates Floyd-Steinberg [FS76] (bottom image) and Eschbach-Knox [EK91] (center image) approaches. I use dither screens generated by the AHE and CLAHE algorithms to enhance the display of the original image details. However, due to correlation between these screens and the reproduced image, the tone reproduction is inaccurate (top row). The use of error diffusion improves tone reproduction and avoids the distortion of image details. The noise artifacts resulted from the enhancement of textures by the AHE algorithm and are controlled by clipping histogram values in the CLAHE algorithm (center column). These images are printed at 150 dpi.

in comparison to Floyd-Steinberg error diffusion (bottom left image), the use of the original image textures in these dither screens improved the display of image details. Unfortunately, the halftoning with the AHE dither screen suffers from noise amplification in the uniformly shaded areas. These noise artifacts are reduced by limiting the contrast enhancement and by clipping histogram values in the CLAHE algorithm.

6.1.2 Applications to the non-clustered and clustered halftoning

Examples in Figure 6.1 demonstrated image enhancement by the use of error diffusion with various dither screens. The other approach is to enhance image features independently from the halftoning algorithm. Thus, Jarvis *et al.* [JJN76] suggested a pre-processing step and convolved an image with an edge enhancement filter. This technique combined with the Floyd-Steinberg error diffusion reveals the edges and textures of the original image (Figure 6.2, top right image). However, this pre-processing step is a compromise between the display of edges and the accuracy of tone reproduction. Because a processed image is used for halftoning, the error diffusion algorithm is unable to approximate original gray scale values. Thus, shifts in the overall image intensity are visible in the regions with many details (e.g. hair, feathers on the hat). The halftoning with the CLAHE dither screen with full error diffusion avoids this artifact while enhancing the reproduction of image texture and edges (Figure 6.2, bottom left).

Overall, the use of image based dither screens produced by the CLAHE algorithm compare favorably to previous halftoning approaches. My approach avoids the local and global shifts in image tones present in partial error diffusion of Eschbach and Knox [EK91] and image pre-processing of Jarvis *et al.* [JJN76] (Figure 6.2). Also, the halftoning with image based dither screens results in the similar tone reproduction as Floyd-Steinberg error diffusion without the loss of original image textures.

The halftoning techniques discussed thus far did not take into account the properties of the output hardware. Because most hard-copy devices are unable to set isolated pixels reliably, halftoning with clustered dots is required (Section 2.2.1). Dot clustering is often achieved by the arrangement of thresholds in the dither screen (Figure 6.3, left). Unfortunately, the use of clustered screens leads to the introduction of strong artificial texture that destroys fine image details.

In this chapter I show how dot clustering is achieved by increasing the size of ap-

Reproduction of edges in non-clustered halftoning



Floyd-Steinberg
error diffusion [FS76]



Error diffusion
with Jarvis et. al.
edge enhancement [JNN76]



Partial error diffusion
Eschbach and Knox [EK91]
 $\alpha = 0.7$



CLAHE
4x4 regions, 50% clipping
full error diffusion

Figure 6.2: Floyd-Steinberg error diffusion smooths out edges and textures of the original image (top left). A pre-processing edge enhancement of Jarvis et. al. [JNN76] minimizes these artifacts of error diffusion (top right). However, the edge enhancement approach leads to local shifts of tone in the highly detailed areas of hair and the hat. The other technique is to reduce the amount of error diffusion thus increasing the overall contrast of the image (bottom left). My approach is to use a dither screen generated by the CLAHE algorithm and to control tone reproduction with the error diffusion algorithm. The resulting image (bottom right) is a close approximation of the original tones that avoids the loss of the original image edges and textures. The images are printed at 150 dpi.




Figure 6.3: Ordered clustered dithering is used in printing hardware. Unfortunately, this algorithm introduces artificial texture and distorts image details (left image). I achieve clustered halftoning by using an error diffusion that processes blocks of pixels at a time. This error diffusion in combination with the image based dither screening results in the alignment of dot clusters along image edges (right image). A dither screen is generated by processing an image with the CLAHE algorithm (33% clipping). Error diffusion processes 2x2 pixel blocks. These images are printed at 300 dpi.

proximation regions used in the error diffusion algorithm (Section 4.4). In a manner similar to non-clustered halftoning I combine error diffusion with the image based dither screen to enhance the display of edges and textures. Thus, 2x2 pixel blocks are used by the modified error diffusion process and resulted in sufficient dot clustering (Figure 6.3, right). In comparison to the traditional ordered dither my approach aligns edge clusters with the edges of the original image thus guaranteeing their display in the final halftoning. Moreover, the original image textures are enhanced.

The above discussion was based on subjective judgment of image quality and visual evaluation was the only argument in favor or against a certain technique. This analysis is dependent on the viewer and is hard to duplicate. Moreover, subjective analysis cannot be adapted to automatic selection of parameters of halftoning in photorealistic optimization based methods (Section 2.2.3) and in interpretive halftoning. In the rest of this chapter I develop a quantitative measure of the halftoning artifacts and distortions. Further, I use this measure to review the results presented thus far.

6.2 Multiscale analysis of halftoned images

Most researchers in quantitative analysis of images have used general models of image vision. An overview of the application of these models to halftoning was presented in Section 2.1. The halftoning process introduces several types of image distortions. For the analysis of a particular photorealistic halftoning algorithm it is important to measure each distortion type separately. Unfortunately, general error metrics are difficult to adapt to this task.

Because edges are important to our perception of visual information, edges should be preserved by both photorealistic and NPR halftoning techniques. The motivation for an edge based image quality research comes from the observation that in traditional illustration techniques such as pen-and-ink or pencil drawing, object edges are the first to be drawn by an artist. Thus, NPR techniques that attempt to imitate traditional art medium are likely to pay greater attention to edges than to other image features such as shading. None of the previous image quality measures could evaluate edge structures independently from the rest of the shading information.

Thus, I have concentrated my research on the analysis of edges in halftoned images. My work is based on multiscale edge analysis developed by Mallat et. al. [MH92, MZ92]. Their approach allows me to detect significant edges at various image

resolutions. I apply multiscale edge analysis to halftoned images and differentiate between two types of edges:

- **reproduced edges**

These edges are present in the original image and should be reproduced by the halftoning algorithm. Typical distortions of image edges include modification of edge shape and position, exaggeration or reduction of edge contrast.

- **artificial edges**

These edges are introduced by the halftoning algorithm and are often the result of tone quantization, abrupt changes of texture, and dot alignment.

The multiscale edge analysis algorithm is described in Section 6.2.1. Application of this algorithm to detection of image artifacts is discussed in Section 6.2.2. I extend the edge based measure to evaluation of rendered image edges in Section 6.2.3.

6.2.1 Multiscale analysis of edges.

The multiscale edge analysis approach is motivated by the physiological evidence that retinal receptive fields act like edge detectors of different sizes and directions. Mallat and Hwang [MH92] proposed a wavelet based technique for detection of significant image edges. Their multiscale edge characterization led to the development of image compression schemes, image noise reduction algorithms, and techniques for motion estimation [MZ92]. I use this multiscale analysis to detect significant edges in halftoned images. Two types of edges are identified: *reproduced* edges and *edge artifacts*. Reproduced edges are edges that were present in the original image. Edge artifacts are introduced by the halftoning algorithm. The halftoned image quality measure uses these two types of edges.

Characterization of multiscale edges.

In this section I review the multiscale edge characterization algorithm proposed by Mallat et. al. [MZ92, MH92].

Most multiscale edge detectors smooth an image at various scales and detect sharp variation points from their first- or second- order derivatives. Similarly, Mallat and Zhong [MZ92] construct a wavelet transform for multiscale edge detection. The wavelet transform $W(f)$ is based on a convolution of a signal f and a dilated *wavelet*

function at multiple scales. Various scales of the signal are derived by multiple application of the *smoothing* function. The smoothing function $\theta(x)$ used by Mallat and Zhong [MZ92] is a spline approximation of a Gaussian filter. If the wavelet functions ψ is the first-order derivative of $\theta(x)$, then the detection of extrema points of the wavelet transform is equivalent to Canny[Can86] edge detection. Thus, I will deal only with extrema points of the wavelet transform.

In the case of images, the extrema points are two-dimensional vectors $(W_x(f), W_y(f))$. The modulus of the wavelet $(|W_x(f)|^2 + |W_y(f)|^2)^{\frac{1}{2}}$ characterizes the strength of an edge. Edge direction at an extrema point can be computed as $\arg(W_x(f) + iW_y(f))$. Extrema points of the same scale are connected into *contour lines*. The two points are chained together if they have similar magnitude and direction. The edge length is equal to the number of points in the contour line.

The evolution of extrema values across multiple scales is the basis for characterization of edges. The *Lipschitz regularity* is the quantitative measure of edge types. Mallat and Hwang[MH92] proved that the Lipschitz regularity can be computed from the decay of coefficients across scales. For example, the Lipschitz regularity for isolated image dots (Dirac functions) is equal to -1 , similarly, the regularity of white noise is negative and close to 0.

In other words, the modulus of a wavelet transform at extrema points which correspond to noise decrease quickly with the reduction of scale. On the other hand, the extrema points that correspond to significant edges are found at a broad range of scales. Therefore, cross-correlation between edges at multiple scales allows us to define a measure of edge importance.

In this work I explore the use of multiscale edge characterization for the analysis of halftoned images. In the next section I show that Mallat and Zhong's [MZ92] algorithm is directly applicable to detection of various types of edges in halftoned images.

Significant edges of a halftoned image.

The analysis of halftoned images is based on detection of the visible edges. I call these edges *significant* edges. The significant edges are detected by the following procedure:

1. The Mallat's et. al. [MH92, MZ92] wavelet transform is applied to a halftoned image.

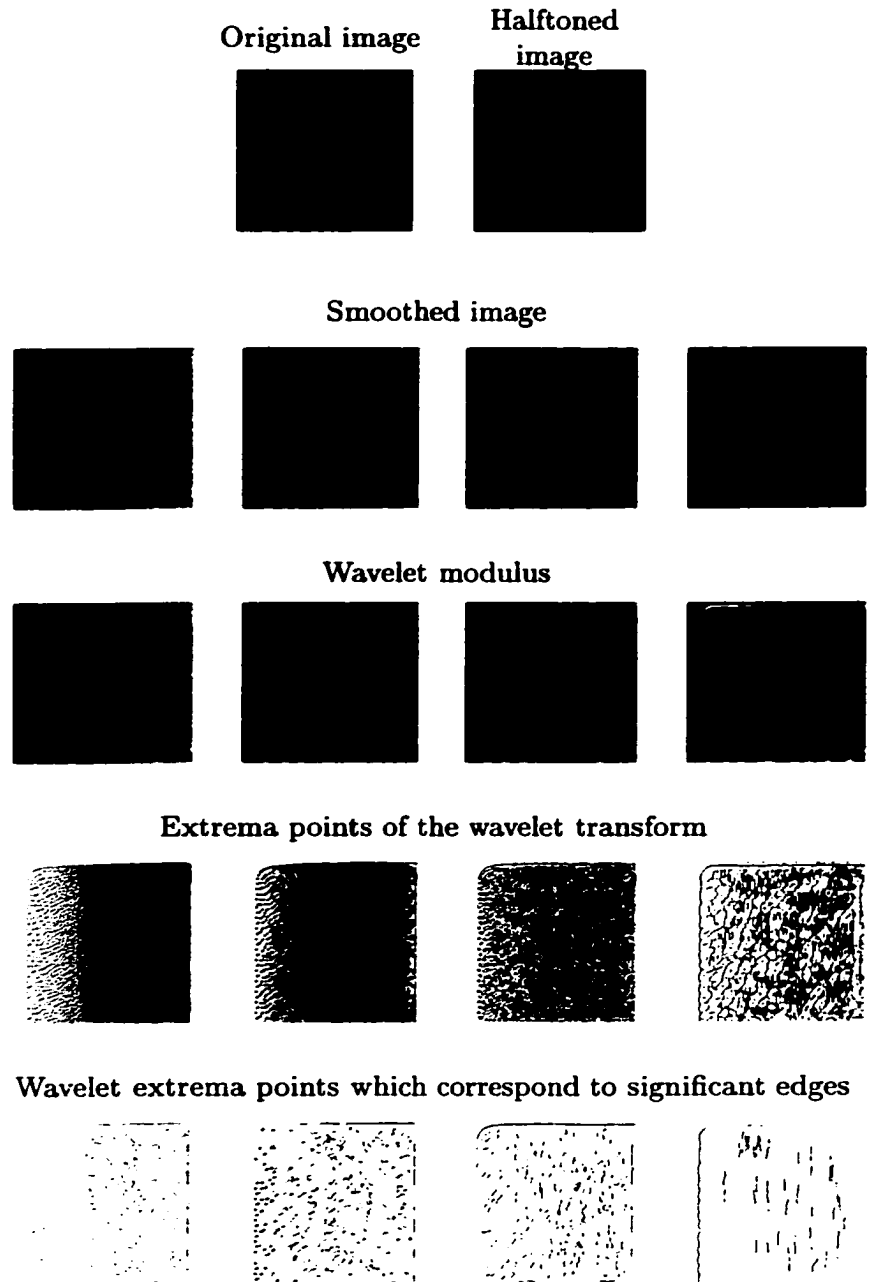


Figure 6.4: A ramp image is halftoned by Floyd-Steinberg error diffusion algorithm and printed at 300 dpi. Multiresolution analysis is used for detection of significant edges. Significant edges closely correspond to visible artifacts

2. Extrema points of the wavelet transform are identified and linked together to form edge contours.
3. Contour lines with the Lipschitz regularity of < -0.3 are removed. Thus, edges corresponding to high frequency noise are removed.
4. Extrema points of a fine scale that do not propagate to a coarser scale are eliminated. Thus, weak edges are removed.

I illustrate the detection of significant edges using a ramp image with intensities $[0, \frac{1}{4})$ halftoned with the Floyd-Steinberg error diffusion algorithm (Figure 6.4, top row). A wavelet smoothing and a wavelet decomposition to four dyadic scales are presented (Figure 6.4, rows 2 and 3). Notice, that at the coarse levels the smoothed out image becomes a closer approximation of the original. Extrema points of the wavelet transform are presented in Figure 6.4 (row 4). The results of detection of significant edges presented in Figure 6.4 (row 5). The significant edges are representative of the known weaknesses of the Floyd-Steinberg halftoning process. Thus, this technique is applicable to the evaluation of the halftoned images.

I define a numeric measure of image quality using significant edges. The first part of this numeric measure should be applied to images that depict edge artifacts. The second part of the measure considers edges in the original image and the manner in which they are reproduced. I present the measure of edge artifacts in Section 6.2.2 and the edge reproduction measure in Section 6.2.3.

6.2.2 Evaluation of edge artifacts.


Given a halftoned image I wish to measure the strength of edge artifacts. This measure is based on the detection of significant edges described in the previous section. I discuss this measure using a gray scale ramp image with intensity range of $[0, \frac{1}{4})$. I study the performance of three popular algorithms applied to this image. These algorithms introduce distinct types of artifacts (Figure 6.5):

- Clustered dot dither with 3×3 matrix produces strong quantization bands.
- Ordered dispersed dither with 8×8 matrix produces visible texture segmentation artifact.
- Floyd-Steinberg error diffusion suffers from snake-like dot alignment.

Clustered 3x3 dither



Dispersed 8x8 dither



Floyd-Steinberg error diffusion




Figure 6.5: Detection of edge artifacts produced by popular halftoning algorithms. Left column presents halftoned images printed at 150 dpi. The edge artifact detected at the 4-th dyadic scale are presented in the right column.

Table 6.1: Strength of edge artifacts across multiple scales

	ϵ^1	ϵ^2	ϵ^3	ϵ^4
Clustered dither	91.03	22.53	5.35	4.08
Dispersed dither	102.65	15.43	4.04	1.46
Floyd-Steinberg	109.47	14.12	2.82	1.45

Let $\mathcal{W}^i = (|W_x^i(f)|^2 + |W_y^i(f)|^2)^{\frac{1}{2}}$ be the magnitude of an extrema point of the wavelet transform at a scale i . The strength of an edge artifact is an average of magnitudes of edge points at a given scale:

$$\epsilon^i = \frac{1}{n^i} \sum_{k=1}^{n^i} \mathcal{W}_k^i \quad (6.1)$$

where, n^i is the number of edge points at a scale i . Numerical results for the tested halftoned images are presented in Table 6.1.

I believe that this ϵ^i measure of the artifact strength closely corresponds to visual observations of the halftoned images. Clustered dither produces the most visible quantization artifact. Also, the corresponding values ϵ^i are consistently large for coarse scales. Dispersed dither does not introduce strong quantization bands, though, sharp texture borders are visible. This is directly reflected by the numerical measures. Dispersed dither artifacts are weaker than the quantization bands. Error diffusion is often the preferred method of image display. The numerical measures of artifacts for this method are also the smallest. Notice, for the coarsest scale dispersed dither and error diffusion artifacts are similar in strength. Indeed, both algorithms, produce similar visual impression when viewed from a large distance.

It is also important to analyze the decay of ϵ^i values with the change of scales. The strength of artifacts decrease the fastest for the error diffusion algorithm. Thus, according to multiscale edge theory [MH92], these artifacts correspond to high frequency noise. Also, according to Ulichney's analysis [Uli88], the spectrum of error diffusion is approximately equal to the blue noise. On the other hand, ϵ^i decrease slowly for the clustered dot dither, due to the low frequency of the produced artifacts.

The presented measure of multiscale edge artifacts closely agrees with the previous research in perception of halftoned images and with visual observations.

6.2.3 Analysis of edge distortions in halftoning.

In the previous section I evaluated the edge artifacts introduced by the halftoning algorithms. There is an increasing interest in the accurate display of edges by clustered dither algorithms. In this section I introduce an edge distortion measure and test it using the following clustered dot algorithms (Figure 6.6):

- Ordered clustered dither with 4×4 matrix.
- Space filling curve clustered dither. Image is halftoned using segments of a Peano curve [VG91].
- Adaptive space filling curve halftoning. The algorithm ensures that segments of the Peano curve do not cross strong image edges [BV95, VG95]. This results in improved rendering of edges.

In order to assess reproduction of edges significant image discontinuities are detected by the technique in Section 6.2.1. Because the strongest edges are present at many resolutions it is sufficient to measure edge distortions only at some coarse scale. Moreover, I decided to account for the fact that human observers often cannot detect small shifts in the absolute position of an edge. Thus, small variations in edge positions are ignored.

Let $\mathcal{W}_{(a,b)}$ and $\bar{\mathcal{W}}_{(a,b)}$ be modulus of wavelet extrema at a point (a, b) of original and halftoned images correspondingly. The error at this point is defined as follows:


$$E_{(a,b)} = \min_{(\alpha,\beta)} |\bar{\mathcal{W}}_{(a,b)} - \mathcal{W}_{(\alpha,\beta)}| \quad (6.2)$$

where point (α, β) is in a close neighborhood of point (a, b) . The overall measure of edge distortion \mathcal{E} is an average of point distortions for a given scale.


In this work I measured distortions at a scale 4 and ignored edge shifts smaller than 3 pixels. The images of edge distortion errors and the numerical results of the tests are presented in Figure 6.6 in the right column. The numerical results closely correspond to my visual observations of rendered edges. Adaptive halftoning produces the smallest average distortion error. While, Peano curve dither introduces the strongest edge distortions.

I found that the signs of the distortion errors provide additional information in classification of distortion types. A negative $E_{(a,b)}$ represents edges smoothed out by the

Ordered clustered 4×4 dither (min=-4.37 max=7.23 , ave=1.09)



Hilbert curve (min=-4.87 max=3.04, ave=1.28)



Adaptive clustering (min=-2.33 max=3.8, ave=0.88)



Figure 6.6: Edge distortions produced by clustered dot halftoning algorithms. Halftoned images printed at 150 dpi are presented in the left column. The right column presents an image distortion errors for significant edges at a scale 4

halftoning process. On the other hand, a positive $E_{(a,b)}$ corresponds to strong edge artifact or an exaggerated halftoned edge. The numerical results reflect the visual observations. Compared with Peano curve dither adaptive clustering introduces less edge smoothing and also tends to overemphasize some edges. Ordered dither and Peano curve result in similar edge smoothing, while ordered dither may create strong artificial edges.

6.2.4 Summary

The main contribution of this research is the introduction of the edge based measures for evaluation of photorealistic halftoning algorithms.

In this study I was mainly concerned with the appearance of edges in halftoned images. Two types of edges were studied: real edges present in the original image; and edge artifacts, created by the halftoning algorithm. Image edges are detected using extrema points of the wavelet transform. Unimportant edge structures are eliminated using multiresolution edge analysis technique. I have derived two quantitative error measures to evaluate significant edges and artifacts. These measures are formulated in terms of modulus of extrema points of the wavelet transform. The numerical results closely correspond to the visual observations and previous research in quantitative analysis of halftoning.

This work is only the first step in the development of a general multiresolution edge based image quality metric. In the future, the following avenues should be explored:

- The general metric should include information about multiple scales with appropriate weights to reflect the effects of features at different sizes.
- In this work directional information was ignored. The edge direction will allow the measurement of not only the strength of an edge but also its shape.

In the following section I use the edge based image quality measures to evaluate texture control techniques presented in this dissertation.

6.3 Quantative evaluation of texture control techniques

Previous discussion of the texture control approaches was based on visual observation. By using the edge based image quality measures I can evaluate these approaches quantitatively and compare texture based techniques to the previous research in halftoning.

Note, that this evaluation must be made in conjunction with the user study since no existing quantitative measure can account for all the properties of the human visual system.




6.3.1 Evaluation of the AHE parameters

My texture control technique relies on the arrangement of thresholds in the dither screen to shape the halftoning texture. I investigated the use of the adaptive histogram equalization (AHE) algorithms to convert an arbitrary image into a dither screen with the desired distribution of threshold values (Section 4.2). Two parameters of the AHE algorithm control the scale of texture features in the resulting halftoning: the processing block size, and the clipping of histogram values (the CLAHE algorithm). I use the measure of halftoning artifacts (Section 6.2.2) to evaluate the influence of these parameters on the appearance of texture.

The motivation to use this measure comes from its ability to estimate edge artifacts introduced by the halftoning algorithm. Artificial edges of the texture control technique correspond mainly to the halftoning texture defined by the dither screen. Therefore, analysis of artificial edges provides a quantitative measure of the appearance of the halftoning texture.

I evaluated the influence of the AHE parameters on the resulting texture by measuring the edges in the halftoning of the ramp with the “wrinkled paper” texture (Figure 4.1). These results are presented in Figure 6.7. The block size parameter of the AHE algorithm allows us to select the size of texture features present in the resulting halftoning. This observation is supported by the measure of edge artifacts. The AHE algorithm with the 8x8 processing blocks produces stronger artifacts at the coarser resolutions than the AHE with the 4x4 blocks. However, the increase of the size of processing blocks eliminates texture features at fine scales. Thus, the edge artifacts at fine scales are higher for the 4x4 blocks than for the 8x8 blocks.

The contrast limited AHE (CLAHE) was introduced to preserve texture features at a wide range of scales and is achieved by the clipping of histogram values. Indeed, the edge artifacts that correspond to the CLAHE algorithms with 4x4 processing blocks are similar to those of the AHE algorithm with 4x4 blocks at fine scales and to the AHE with 8x8 blocks at coarser scales. Moreover, the 50% clipping produces coarser halftoning textures than the 75% clipping and thus the corresponding edge artifacts



AHE parameters:		The strength of edge artifacts				
Region size	Clipping	ϵ^1	ϵ^2	ϵ^3	ϵ^4	ϵ^5
4x4	no clipping (AHE)	97.40	42.28	27.23	18.27	9.53
8x8	no clipping (AHE)	97.55	46.45	37.61	25.63	11.82
4x4	75 % (CLAHE)	97.55	41.84	27.92	20.25	12.18
4x4	50 % (CLAHE)	97.27	43.29	30.87	23.21	15.07




Figure 6.7: The measure of edge artifacts is used for evaluation of the parameters of the AHE algorithm. The increase in the size of the processing blocks leads to the increase in the size of texture features in the resulting texture. Thus, magnitude of the edge artifacts at coarser scale is higher for the AHE with 8x8 blocks than for the 4x4 blocks. The use of the histogram clipping by the CLAHE algorithm results in the appearance of texture features at a wide range of scales. Thus, the corresponding edge artifacts for the CLAHE algorithm are similar to those of the AHE with the 4x4 blocks at high resolutions and to the AHE with the 8x8 blocks at coarser scales.

are stronger at low resolutions.

6.3.2 Evaluation of the error diffusion parameters

I control the contrast of the halftoning texture by combining ordered dithering with the error diffusion algorithm (Section 4.4). The reduction of texture contrast results in weakening of edge artifacts at coarse scales. Thus, the measure of edge artifacts can be used to assess the influence of error diffusion parameters on the halftoning texture.


The error diffusion algorithm is controlled by two parameters: α is the influence of approximation errors on the output for the current pixel (Equation 4.10); and the size of the approximation regions. I evaluate the effect of these parameters using the halftoning of the gray scale ramp with the “stone” texture (Figure 4.6). The fragment of the ramp that is used in this evaluation and the numerical values of the measure are presented in Figure 6.8.

The increase of the α parameter corresponds to the higher weight of approximation errors and leads to softer textures. Thus, the strength of artificial edges is decreasing for the coarse scales and increases at the fine scales. The $\alpha = 1$ corresponds to the full error diffusion and is the closest to the Floyd-Steinberg error diffusion algorithm with constant threshold. However, artificial edges of the Floyd-Steinberg algorithm are significantly smaller at coarser scales than those introduced by the halftoning with the dither screen. This effect is explained by the visibility of textures even for $\alpha = 1$ and proves that full error diffusion preserves texture features of the dither screen.

I control the dot clustering of the resulting halftoning by increasing the size of approximation regions. I compare artificial edges produced by error diffusion with $\alpha = 1$ and various sizes of the approximation regions: 1 pixel, 2x2 pixels, and 4x4 pixels (Figure 6.8). The clustering effect increases the strength of edges at coarser scales. Never the less, error diffusion with 4x4 approximation regions introduces less edges at coarser scale and thus less clustering than ordered dither without error diffusion $\alpha = 0$.

6.3.3 Evaluation of photorealistic halftoning results

Previously, I used the multiscale edge measure to assess artificial edges introduced by halftoning with different parameters. Here I apply multiscale analysis to evaluate the accuracy of image reproduction in photorealistic halftoning. In particular, I use



Error diffusion parameters:		The strength of edge artifacts				
α	Region size	ϵ^1	ϵ^2	ϵ^3	ϵ^4	ϵ^5
0.0	1	97.20	39.28	17.95	8.12	4.31
0.5	1	100.69	26.16	10.57	4.99	3.69
1.0	1	104.66	21.19	5.46	2.10	2.95
1.0	2x2	97.832	37.78	12.12	3.71	3.20
1.0	4x4	97.92	39.44	18.07	6.16	3.55
Floyd-Steinberg with constant threshold						
		109.47	14.12	2.82	1.45	1.02

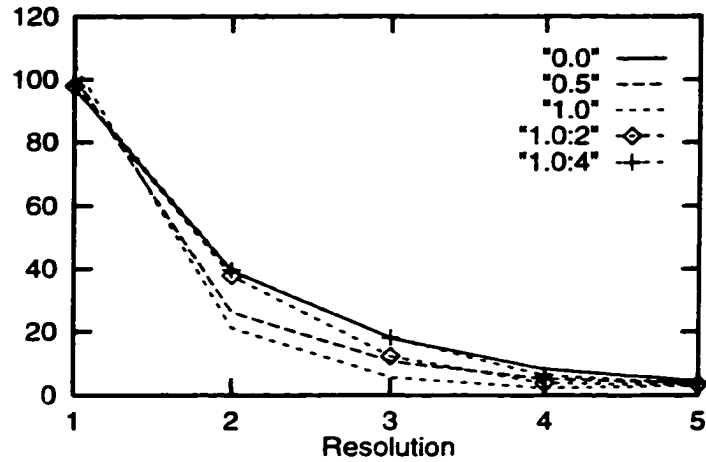


Figure 6.8: The α parameter of the error diffusion and the size of the approximation region influence the contrast of the halftoning texture. The increase in α corresponds to softer textures and reduces the strength of artificial edges at coarse scales. However, the increase of the approximation regions leads to texture clustering and stronger texture edges for the same parameter α .

Table 6.2: Edge distortions of the photograph “Lena”

Algorithm	Edge distortions				
	ϵ^1	ϵ^2	ϵ^3	ϵ^4	ϵ^5
Floyd-Steinberg	101.5	3.4	0.42	0.28	0.28
Jarvis edge enhancement	76.4	4.9	0.69	0.41	0.33
Partial error diffusion	87.2	3.4	0.45	0.31	0.30
CLAHE 4x4 regions 50 & clipping	93.37	3.7	0.42	0.29	0.27

the edge distortion measure (Section 6.2.3) to compare the halftoning results of the photograph “Lena”. The goal of this study is to evaluate the distortion of textures and edges produced by different halftoning approaches (Figure 6.2):

- Floyd-Steinberg error diffusion [FS76];
- Jarvis et. al. [JJN76] edge enhancement combined with error diffusion;
- The partial error diffusion of Eschbach and Knox [EK91] ($\alpha = 0.7$);
- Halftoning with the CLAHE generated dither screen and full error diffusion, developed in this dissertation.

I compare the gray scale photograph with the halftoning results by detecting significant edges at multiple scales and by matching these edges to each other. I use 5 scales for this evaluation and the edges are matched in the 3x3 neighborhood. The numeric results are summarized in Table 6.2.

Floyd-Steinberg error diffusion produced the smallest distortions at scales 2, 3, and 4, while edge enhancement of Jarvis et. al. generated the highest distortion of edges at all scales except the first. These quantitative measures support the previous observation that the pre-processing edge enhancement leads to shifts in overall tone of image regions with many edges (feathers and hair). Moreover, the Jarvis et. al. approach tends to overemphasize edges and amplifies image noise.

Partial error diffusion of Eschbach and Knox [EK91] increases the overall contrast of the halftoned image and thus the corresponding edge distortion measures are somewhat higher than the distortion of Floyd-Steinberg approach. Unlike the edge enhancement technique, partial error diffusion does not overemphasize edges and produces smaller edge distortions.

I enhance reproduction of image edges and textures using the dither screen generated from the Lena image by the CLAHE algorithm. I combine ordered dithering with

full error diffusion to guarantee the accurate approximation of image tones. The edge distortion measure supports my observation that texture control technique compares favorably against previous approaches. Numeric measures demonstrate that the use of the CLAHE generated dither screens produce less edge distortion than the edge enhancement and partial error diffusion. Furthermore, the edge distortion of my technique is close to the distortions of Floyd-Steinberg approach and are smaller at the finest and the coarsest scale. Therefore, my texture control technique results in the accurate approximation of image tones and preserves image edges at multiple scales.

I believe that the close correspondence between my texture control technique and Floyd-Steinberg error diffusion motivates further research in the field of image distortion measures. Even though we often prefer images with clear defined edges and textures the numeric results produced by the current measures are negligible close. Moreover, my previous study [BSV98] concluded the preference of images with sharp edges is subjective and depends on the viewer and the displayed image.

6.3.4 Summary of the quantitative results

I used multiscale edge analysis to support visual evaluation of texture control parameters developed in this thesis. In particular, the measure of edge artifacts provided numerical estimates of the scale and contrast of halftoning texture.

The scale of texture features is controlled by the parameters of the AHE algorithm. Numerical measures proved that increasing the size of processing blocks enables us to select texture features of a limited scale. Moreover, the use of histogram clipping allows us to preserve texture features of a wide range of scales.

I control texture contrast using the error diffusion process. Numerical measures demonstrated the reduction of coarse scale textures with the increase of the amount of error diffusion. Also, the use of larger approximation regions creates clustered dithering with strong edge artifacts at coarse scales.

I used the edge distortion measures to compare the halftoning of the photograph "Lena". Again, the numerical results supported visual observations and proved that the use of my texture control technique results in a good approximation of image tones with enhanced rendering of image textures and edges.

6.4 Discussion of texture control in photorealistic halftoning

The challenge of photorealistic halftoning is to approximate a continuous image without introduction of strong artifacts. Previous research in ordered dithering attempts to define a single dither matrix that is suitable for reproduction of every image. I took a different approach and applied the image based texture control technique (Section 4.2) to generate a dither screen for the halftoning of a particular image. A dither screen is thus created using the CLAHE algorithm. Because the resulting screen contains textures and edges of the original image these features are guaranteed to appear in the halftoned display.

I combine ordered dithering with the error diffusion process to improve the accuracy of tone reproduction. The full error diffusion is required to minimize the effect of the correlation between the threshold values in the dither screen and values of the reproduced image. Thus, the resulting display is the close approximation of the original image tones that minimizes distortions of image edges and textures. Moreover, I control the clustering properties of the halftoned image by adjusting the size of approximation regions of error diffusion. Thus the resulting image is less subjected to the dot gain of the current hard copy devices and is the alternative to the clustered ordered dither.

In order to compare the halftoned images I developed multiscale edge based measures of image distortion. I use the wavelet decomposition algorithm to identify significant image edges and to quantify their strength at multiple resolutions. I tested the edge based measures by analyzing the previous halftoning approaches and found a close correspondence between the visual observation and the numerical results.

I compared the texture control technique to the previous halftoning approaches. In particular, my technique is an improvement over the Floyd-Steinberg approach [FS76] in reproduction of image details without the loss of the tone approximation accuracy. Moreover, the accuracy of tone approximation is far superior to the accuracy of the pre-processing edge enhancement of Jarvis et. al. [JJN76] and partial error diffusion of Eschbach and Knox [EK91]. These conclusions are based on both visual observation and the values of numerical edge distortion measures.

The developed measure of edge distortion is only a first step in the research of the gen-

eral image quality measures. My experiments showed that the current measure is not sensitive enough to differentiate between the edge reproduction of the Floyd-Steinberg and the texture control techniques. Possibly the use of directional information and the incorporation of contrast sensitivity functions will improve the accuracy of numeric evaluation.

Chapter 7

Conclusion

Rendering of continuous tone images on binary hardware requires a halftoning step. A halftoning process approximates continuous tones by an arrangement of binary display primitives and is often associated with the introduction of additional texture. Conventional halftoning research attempts to minimize texture artifacts. In this thesis I took a different approach and used the halftoning texture to enhance image display. I achieved this goal by controlling the halftoning texture and by adapting it to the image information. Thus, the control of texture in binary displays is the main contribution of this thesis.

This research is motivated by artistic rendering and is closely related to the non-photorealistic area of computer graphics. The goal of artistic rendering is to express information using limited means of the display medium. The objective of NPR research is to develop tools that allow an artist to render expressive images using computers. My research is specific to binary displays. It enables an artist to enhance image rendering through the manipulation of halftoning texture.

My approach to texture control is based on previous halftoning results and allows us to define texture shape, direction, scale, and contrast. Image tones are approximated using a combination of ordered dithering and error diffusion. I use the property of ordered dithering to control the appearance of the halftoning texture. Similarly to previous research, my dither screens are created from dither matrices. The essential properties of a dither matrix are identified and approximated using image processing and procedural approaches. Further, unlike conventional halftoning, I combine dither matrices into a dither screen using mapping functions. These functions control texture scale and direction, and account for objective and subjective image information. I

studied image based and scene based mapping functions and demonstrated their use in rendering photographs and 3-D scenes. Once a dither screen is constructed, I use it to modulate thresholds of a parameterized error diffusion algorithm. Error diffusion parameters control the accuracy of tone reproduction and the contrast of the halftoning texture.

The control of halftoning texture is applicable to both photorealistic and artistic rendering of images. Therefore, my research is a contribution to both areas of computer graphics.

7.1 Contributions to photorealistic halftoning

The texture control techniques developed in this thesis extend and generalize previous results of conventional halftoning.

- The essential distribution properties of a good dither matrix are identified and approximated. It is proven that contrast limited histogram equalization algorithm transforms an image based texture into a dither matrix. Thus, dither matrices with desired shape can be generated automatically from texture images. Moreover, a procedural approach to dither matrix design was investigated.
- It is discovered that the use of image based dither screens in ordered dithering improves the fidelity of photorealistic halftoning. In particular, I generated dither screens from the images to be reproduced and used them to modulate thresholds of the error diffusion algorithm. As a result, the reproduction of original edges and textures was enhanced.
- I developed and studied a multiscale edge based image quality measure. This measure allows us to quantify the artifacts of the halftoning image including quantization bands, abrupt texture changes, and dot alignment. Also, this measure enables us to evaluate distortions of edges and textures of the original images.

7.2 Contributions to non-photorealistic rendering

This research is a new approach to expressive display of images. Instead of developing a new rendering system, I extended and adapted to the task conventional halftoning algorithm. The following are the advantages of my approach:

- The use of conventional halftoning results in computationally efficient and accurate tone approximation in a non-photorealistic style.
- The texture control technique is not limited to any style of rendering. Rendering styles resembling traditional art materials and techniques may be generated by designing a dither matrix with the desired properties. Unlike previous extensions of halftoning to NPR [OH95, Buc96], the design of dither matrices is a well controlled process and does not require manual input from an artist.
- I developed the control of texture shape, direction, scale, and contrast and allow an artist to map any of these features to the chosen type of information. Most of the previous non-photorealistic binary rendering systems were limited to the control of texture direction [SC93, Elb95, WS96] and none were considering texture contrast.
- The algorithm is independent from the image source. Therefore, both photographs and photorealistic computer generated images can be used as an input to the halftoning system. The system can account for any type of information either derived by an image processing operation, supplied by the scene description, or defined by an artist.

7.3 Limitations of the technique

Even though my research contributed to photorealistic and artistic rendering of images it has a number of limitations and disadvantages.

Most previous NPR techniques control the individual placement of graphics primitives. For example, lines and dots in pen and ink rendering [SWHS97] can be moved and rotated one by one to match the desired tone and direction. The control of fine texture features — e.g. one brush stroke — is impossible in my system. Moreover, the use of error diffusion often destroys the shape of small texture features. Therefore, when compared to artistic screening of Ostromoukhov and Hersch [OH95], the system does not produce clearly defined texture elements (e.g. text).

The control of texture scale and direction is limited to piece-wise constant parameters. Continuous changes of these texture features must be correlated and require additional investigation.

The use of image generated dither screens in photorealistic halftoning often leads to

noise amplification in the resulting images. Optimization based halftoning techniques often result in a better approximation of image features and are able to avoid strong artificial texturing.

The multiscale edge based measure is only a first step in the development of a reliable image quality metric. The current measure does not take into account many known properties of the human visual system including contrast sensitivity function, texture orientation and viewing conditions.

Bibliography

- [AA92] Mostafa Analoui and Jan P. Allebach. Model based halftoning using direct binary search. In Rogowitz [Rog92], pages 96–108.
- [AR93] Jan P. Allebach and Bernice E. Rogowitz, editors. *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display IV*, volume 1913, San Jose, California, February 1993. SPIE.
- [ARS79] A. Appel, F. J. Rohlf, and A. J. Stein. The haloed line effect for hidden line elimination. In *Computer Graphics (SIGGRAPH '79 Proceedings)*, volume 13, pages 151–157, August 1979.
- [Bay73] B. E. Bayer. An optimum method for two-level rendition of continuous tone pictures. *IEEE International Conference on Communications*, 1:26–11–26–15, June 1973.
- [BHB83] C. Billotet-Hoffman and O. Bryngdahl. On the error diffusion technique for electronic halftoning. *Proceedings of the Society for Information Display*, 24(3):253–258, 1983.
- [BSV98] John W. Buchanan, Lisa Streit, and Oleg Veryovka. Edge enhancement issues in half-toning. In *Graphics Interface 98*, June 1998. Accepted for publication.
- [Buc96] J. W. Buchanan. Special effects with half-toning. In *Proceedings of Eurographics 96*, pages 97–108, 1996.
- [BV95] J. W. Buchanan and O. Verevka. Edge preservation with space-filling curve half-toning. In W. Davis and P. Prusinkiewicz, editors, *Graphics Interface '95*, pages 75–82, May 1995.
- [BVG91] Chakib Bennis, Jean-Marc Vézien, Gérard Iglésias, and André Gagalowicz. Piecewise surface flattening for non-distorted texture mapping. In Thomas W. Sederberg, editor, *Computer Graphics (SIGGRAPH '91 Proceedings)*, volume 25, pages 237–246, July 1991.
- [Can86] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:679–698, 1986.
- [CAS+97] Cassidy J. Curtis, Sean E. Anderson, Joshua E. Seims, Kurt W. Fleischer, and David H. Salesin. Computer-generated watercolor. In *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 421–430, August 1997.
- [DC90] Debra Dooley and Michael Cohen. Automatic illustration of 3D geometric models: Lines. In Rich Riesenfeld and Carlo Sequin, editors, *Computer Graphics (1990 Symposium on Interactive 3D Graphics)*, volume 24, pages 77–82, March 1990.
- [EK91] Reiner Eschbach and Keith T. Knox. Error-diffusion algorithm with edge enhancement. *Journal of the Optical Society of America A*, 8(12):1844–1850, December 1991.
- [Elb95] Gershon Elber. Line art rendering via a coverage of isoparametric curves. *IEEE Transactions on Visualization and Computer Graphics*, 1(3):231–239, September 1995. ISSN 1077-2626.

- [Elb98] Gershon Elber. Line art illustrations of parametric and implicit forms. *IEEE Transactions on Visualization and Computer Graphics*, 4(1), January - March 1998. ISSN 1077-2626.
- [EMP⁺94] David Ebert, Kent Musgrave, Darwyn Peachey, Ken Perlin, and Worley. *Texturing and Modeling: A Procedural Approach*. Academic Press, October 1994. ISBN 0-12-228760-6.
- [Fan94] Zhigang Fan. Error diffusion with a more symmetric error distribution. In *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display V*, volume 2179, pages 150–158, February 1994.
- [FS76] Robert W. Floyd and Louis Steinberg. An adaptive algorithm for spatial greyscale. *Proceedings of the Society for Information Display*, 17(2):75–77, 1976.
- [FvDFH90] James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics, Principles and Practice, Second Edition*. Addison-Wesley, Reading, Massachusetts, 1990. Overview of research to date.
- [Gau92] J. M. Gauch. Investigations of image contrast space defined by variations on histogram equalization. *CVGIP: Graphical Models and Image Processing*, 4(54):269–280, July 1992.
- [GMY97] A. Gaddipatti, R. Machiraju, and R. Yagel. Steering image generation with wavelet based perceptual metric. *Computer Graphics Forum*, 16(3):241–252, August 1997. Proceedings of Eurographics '97. ISSN 1067-7055.
- [GRS93] Robert Geist, Robert Reynolds, and Darrell Suggs. A markovian framework for digital halftoning. *ACM Transactions on Graphics*, 12(2):136–159, April 1993.
- [Hae90] Paul E. Haeberli. Paint by numbers: Abstract image representations. In *Computer Graphics (SIGGRAPH '90 Proceedings)*, volume 24, pages 207–214, August 1990.
- [IFP96] Victoria Interrante, Henry Fuchs, and Stephen Pizer. Illustrating transparent surfaces with curvature-directed strokes. In *IEEE Visualization '96*. IEEE, October 1996. ISBN 0-89791-864-9.
- [IFP97] Victoria Interrante, Henry Fuchs, and Stephen M. Pizer. Conveying the 3D shape of smoothly curving transparent surfaces via texture. *IEEE Transactions on Visualization and Computer Graphics*, 3(2), April–June 1997. ISSN 1077-2626.
- [Int97] Victoria L. Interrante. Illustrating surface shape in volume data via principal direction-driven 3D line integral convolution. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 109–116. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.
- [Jai89] A. K. Jain. *Fundamentals of Digital Image Processing*. Prentice Hall, 1989.
- [JJN76] J. F. Jarvis, C. N. Judice, and W. H. Ninke. A survey of techniques for the display of continuous tone pictures on bilevel displays. *Comp. Graphics and Image Processing*, 5:13–40, March 1976.
- [KK87] Tomihisa Kamada and Satoru Kawai. An enhanced treatment of hidden lines. *ACM Transactions on Graphics*, 6(4):308–323, October 1987.
- [Knu87] Donald E. Knuth. Digital halftones by dot diffusion. *ACM Transactions on Graphics*, 6(4):245–273, 1987.
- [Lei94] W. Leister. Computer generated copper plates. *Computer Graphics Forum*, 13(1), 1994.

- [LM98] Bruno Lévy and Jean-Laurent Mallet. Non-distorted texture mapping for sheared triangulated meshes. In Michael Cohen, editor, *SIGGRAPH 98 Conference Proceedings*, Annual Conference Series, pages 343–352. ACM SIGGRAPH, Addison Wesley, July 1998. ISBN 0-89791-999-8.
- [LS95] J. Landsdown and S. Schofield. Expressive rendering: A review of nonphotorealistic techniques. *Computer Graphics and Applications*, 15(3):29–37, 1995.
- [MA92] Jeffrey B. Mulligan and Albert J. Ahurmada, Jr. Principled halftoning based on human vision models. In Bernice E. Rogowitz, editor, *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display III*, volume 1666, pages 109–121, San Jose, California, February 1992. SPIE.
- [Mar79] D. Marr. Representing and computing visual information. In D. H. Winston and R. H. Brown, editors, *Artificial Intelligence: An MIT Perspective*, pages 18–80. MIT Press, Cambridge, MA, 1979.
- [MB95] Theophano Mitsa and Paul Brathwaite. Wavelets as a tool for the construction of a halftone screen. In Bernice E. Rogowitz and Jan P. Allebach, editors, *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display VI*, volume 2411, pages 228–238, San Jose, California, February 1995. SPIE.
- [MH92] S. Mallat and W. L. Hwang. Singularity detection and processing with wavelets. In *IEEE Transactions on Information Theory*, volume 38, pages 617–643, Mar 1992.
- [Mit93] Theophano Mitsa. Evaluation of halftone techniques using psychovisual testing and quantitative quality measures. [AR93], pages 177–187.
- [ML80] S. Matsumoto and B. Liu. Analytical fidelity measures in the characterization of halftone processes. *Journal of the Optical Society of America*, 70(10):1248–1254, 1980.
- [MP91] Theophano Mitsa and Kevin J. Parker. Digital halftoning using a blue noise mask. In *ICASSP 91: 1991 International Conference on Acoustics, Speech, and Signal Processing*, volume 2, pages 2809–2812, Toronto, Canada, May 1991. IEEE.
- [MP92] Theophano Mitsa and Kevin J. Parker. Digital halftoning using a blue-noise mask. *Journal of the Optical Society of America A*, 9(11):1920–1929, November 1992.
- [MV93] Theophano Mitsa and Krishna Lata Varkur. Evaluation of contrast sensitivity functions for the formulation of quality measures incorporated in halftoning algorithms. In *ICASSP-93: 1993 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 5, pages 301–304, Minneapolis, Minnesota, April 1993. IEEE.
- [MVA93] Theophano Mitsa, Krishna Lata Varkur, and Jennifer R. Alford. Frequency-channel-based visual models as quantitative quality measures in halftoning. In Allebach and Rogowitz [AR93], pages 390–401.
- [MZ92] S. Mallat and S. Zhong. Characterization of signals from multiscale edges. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 14, pages 710–732, Jul 1992.
- [NL96] Avi C. Naiman and David T. W. Lam. Error diffusion: Wavefront traversal and contrast considerations. In Wayne A. Davis and Richard Bartels, editors, *Graphics Interface '96*, pages 78–86. Canadian Information Processing Society, Canadian Human-Computer Communications Society, May 1996. ISBN 0-9695338-5-3.
- [OH95] Victor Ostromoukhov and Roger D. Hersch. Artistic screening. In *SIGGRAPH 95 Conference Proceedings*, pages 219–228. Addison Wesley, August 1995.

- [OHA94] Victor Ostromoukhov, Roger D. Hersch, and Isaac Amidror. Rotated dispersion dither: a new technique for digital halftoning. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24-29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 123-130. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
- [PB94] Y. Pnueli and A. M. Bruckstein. DIGⁱ_dürer – a digital engraving system. *The Visual Computer*, 10:277-292, 1994.
- [PH89] Ken Perlin and Eric M. Hoffert. Hypertexture. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 253-262, July 1989.
- [PN92] Thrasyvoulos N. Pappas and David L. Neuhoff. Least-squares model-based halftoning. In Rogowitz [Rog92], pages 165-175.
- [Rob97] Barbara Robertson. Different strokes. *Computer Graphics World*, pages 26-35, 1997.
- [Rog92] Bernice E. Rogowitz, editor. *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display III*, volume 1666, San Jose, California, February 1992. SPIE.
- [SABS94] Michael P. Salisbury, Sean E. Anderson, Ronen Barzel, and David H. Salesin. Interactive pen-and-ink illustration. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24-29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 101-108. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
- [SB98] Lisa Streit and John Buchanan. Importance driven halftoning. In *Proceedings of Eurographics 98*, August 1998. To appear.
- [SC93] K. R. Sloan and Dan Campbell. Texture as information. In Jan P. Allebach and Bernice E. Rogowitz, editors, *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display IV*, volume 1913, pages 344-354, San Jose, California, February 1993. SPIE.
- [Sny85] Harry L. Snyder. The visual system: Capabilities and limitations. In L.E. Tannas, Jr., editor, *Flat Panel Displays and CRT's*, pages 54-69. Van Nostrand Reinhold, New York, 1985.
- [Sou98] Mario Costa Sousa. Computer generated pencil sketching and rendering of three-dimensional models, 1998. Ph.D. Thesis Proposal.
- [SP94] Mark A. Schulze and Thrasyvoulos N. Pappas. Blue noise and model-based halftoning. In *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display V*, volume 2179, pages 182-194, February 1994.
- [Spe17] Harold Speed. *The Practice and Science of Drawing*. Steeley, Service and Co. Ltd., London, 1917.
- [SPR⁺94] T. Strothotte, B. Preim, A. Raab, J. Schumann, and D. R. Forsey. How to render frames and influence people. *Computer Graphics Forum*, 13(3):455-466, 1994. Eurographics '94 Conference issue.
- [ST90] Takafumi Saito and Tokiichiro Takahashi. Comprehensible rendering of 3-D shapes. In Forest Baskett, editor, *Computer Graphics (SIGGRAPH '90 Proceedings)*, volume 24, pages 197-206, August 1990.
- [Str86] Steve Strassmann. Hairy brushes in computer-generated images. M.sc. thesis, MIT Media Laboratory, June 1986.
- [Stu79] Peter Stucki. *Image Processing for Document Reproduction*, pages 177-218. Plenum, 1979.

- [SWHS97] Michael P. Salisbury, Michael T. Wong, John F. Hughes, and David H. Salesin. Orientable textures for image-based pen-and-ink illustration. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, pages 401–406, August 1997.
- [TM94] Stefan Thurnhofer and Sanjit K. Mitra. Nonlinear detail enhancement of error-diffused images. In *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display V*, volume 2179, February 1994.
- [ULi87] Robert Ulichney. *Digital Halftoning*. MIT Press, 1987.
- [ULi88] R. A. Ulichney. Dithering with blue noise. *Proc. IEEE*, 76(1), January 1988.
- [ULi93] Robert Ulichney. The void-and-cluster method for dither array generation. In *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display IV*, volume 1913, pages 332–343, February 1993.
- [VFB98] Oleg Veryovka, Alain Fournier, and John Buchanan. Edge enhancement issues in half-toning. In *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display VIII*, pages 461–472, January 1998.
- [VG91] Luiz Velho and Jonas de Miranda Gomes. Digital halftoning with space filling curves. In Thomas W. Sederberg, editor, *Computer Graphics (SIGGRAPH '91 Proceedings)*, volume 25, pages 81–90, July 1991.
- [VG95] Luiz Velho and Jonas Gomes. Stochastic screening dithering with adaptive clustering. In Robert Cook, editor, *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pages 273–276. ACM SIGGRAPH, Addison Wesley, August 1995. held in Los Angeles, California, 06-11 August 1995.
- [WN82] I.H. Witten and R.M. Neal. Using peano curves for bilevel display of continuous-tone images. *IEEE Computer Graphics and Applications*, 2(3):47–52, May 1982.
- [Wor96] Steven P. Worley. A cellular texture basis function. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 291–294. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.
- [WS94] Georges Winkenbach and David H. Salesin. Computer-generated pen-and-ink illustration. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 91–100. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
- [WS96] Georges Winkenbach and David H. Salesin. Rendering parametric surfaces in pen and ink. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 469–476. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.
- [ZLE92] Avideh Zakhor, Steve Lin, and Farokh Eskafi. A new class of b/w halftoning algorithms. In Bernice E. Rogowitz, editor, *Proceedings, SPIE—The International Society for Optical Engineering: Human Vision, Visual Processing, and Digital Display III*, volume 1666, pages 122–133, San Jose, California, February 1992. SPIE.
- [ZLE93] Avideh Zakhor, Steve Lin, and Farokh Eskafi. A new class of b/w halftoning algorithms. *IEEE Transactions on Image Processing*, 2(4):499–509, October 1993.
- [Zui94] Karel Zuiderveld. Contrast limited adaptive histogram equalization. In Paul Heckbert, editor, *Graphics Gems IV*, pages 474–485. Academic Press, Boston, 1994.