

Comprehensive Halftoning of 3D Scenes

O. Veryovka and J. Buchanan[†]
Department of Computing Science
University of Alberta, Edmonton, Canada
[†]On leave at Electronic Arts,
Burnaby, British Columbia, Canada

Abstract

The display of images on binary output hardware requires a halftoning step. Conventional halftoning algorithms approximate image values independently from the image content and often introduce artificial texture that obscures fine details. The objective of this research is to adapt a halftoning technique to 3D scene information and thus to enhance the display of computer generated 3D scenes. Our approach is based on the control of halftoning texture by the combination of ordered dithering and error diffusion techniques. We extend our previous work and enable a user to specify the shape, scale, direction, and contrast of the halftoning texture using an external buffer. We control texture shape by constructing a dither matrix from an arbitrary image or a procedural texture. Texture direction and scale are adapted to the external information by the mapping function. Texture contrast and the accuracy of tone reproduction are varied across the image using the error diffusion process. We halftone images of 3D scenes by using the geometry, position, and illumination information to control the halftoning texture. Thus, the texture provides visual cues and can be used to enhance the viewer's comprehension of the display.

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 study deals with rendering of images on binary output devices such as black and white printers. For the display on binary hardware continuous tone images must be approximated by black and white regions. The process of binary approximation is called halftoning or dithering. Due to the limited resolution of the commonly used printers the halftoning manifests itself in the form of artificial texture that may obscure image details. Thus, the goal of conventional halftoning research[2, 9, 32] is to minimize texture artifacts [11, 33, 31, 24, 18, 16].

Our study is a departure from conventional halftoning. Instead of hiding the halftoning texture we use it to enhance the display of computer generated three-dimensional scenes. We enable the user to adapt texture shape, scale, direction, and contrast based on the scene information. The halftoning texture is thus used here to provide visual cues and to enhance viewer's comprehension of 3D scenes.

Comprehensive halftoning is motivated by traditional art techniques and is closely related to the non-photorealistic rendering (NPR) area of computer graphics. We review the current NPR results in the following section. Many NPR techniques are based on line rendering and approximate image tones using computationally expensive algorithms. The main contribution of our research is the extension of previous results in conventional halftoning to non-photorealistic rendering. This allows us to take advantage of computational efficiency and accuracy of traditional or-

dered dithering[2, 32, 33] and error diffusion[9, 16] halftoning approaches and apply them to approximate tone and to define the texture of the non-photorealistic display.

This research is based on our previous work[34] on textured halftoning. The main focus of the current investigation is the local control of texture features by external parameters stored as auxiliary images. We enable the user to vary texture features by procedurally generating a dither matrix and by adapting it to the control image (Section 3). We use a mapping function to construct a dither screen from the resulting dither matrices and to adapt texture direction and scale to the control image (Section 4). Image tones are approximated by the combination of ordered dithering with parameterized error diffusion. The use of error diffusion allows us to control the accuracy of tone reproduction and to vary texture contrast across the image (Section 5).

We apply these texture control techniques to enhance the display of 3D scenes. 3D information is stored in *g*-buffers[21] and is used to adapt texture shape, direction, scale, and contrast (Section 6). Our technique also enables a designer to provide additional visual cues and to tailor image display to the subjective information such as object importance.

2 Motivation and Current Results of Non-Photorealistic Rendering

The current research is motivated by the fundamental difference between photorealistic and artistic approaches to information display. Photorealistic image is either a mechanical recording of the reality or the result of modeling and approximation of geometry and illumination. An art work, on the other hand, is based on the author's interpretation of the world, that leads "*to selection of significant and suppression of non-essential*" [26] in the final image. Therefore, an artist eliminates the unwanted details that may obscure the display of the intended message.

Conventional halftoning takes a photorealistic approach and attempts to approximate continuous tones without taking into consideration the content of the image. Because texture is inevitably introduced by the halftoning method and is unrelated to the image it is often interferes with our understanding of the display (Figure 1, top). Thus, conventional halftoning is fundamentally different from artistic rendering, where the artist manipulates the display primitives — dots, lines, brush strokes — and adapts them to enhance our comprehension of the image.

The goal of our investigation is to tailor the halftoning texture to the 3D information and is thus to enhance the display with artistic elements providing visual cues of the form, position, and illumination of the scene. This goal is shared by the previous research in non-photorealistic rendering (NPR) and is best expressed by Strothotte et. al.[29]:

... to study alternative methods of rendering images to provide designers with the ability to express more than

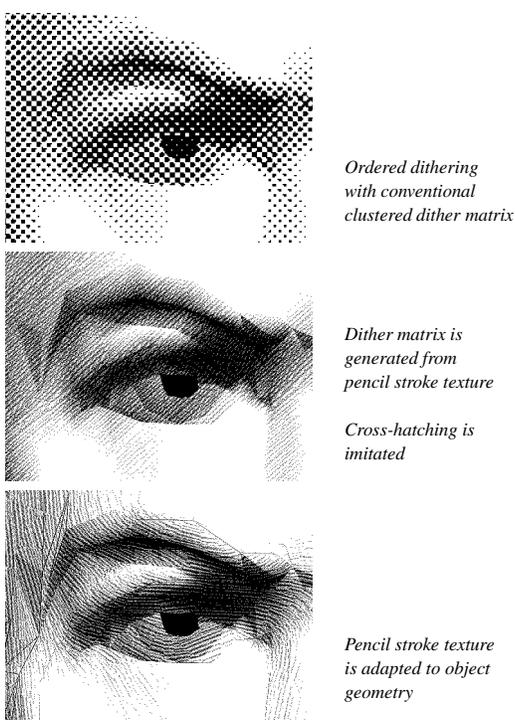


Figure 1: Conventional ordered dithering introduces strong artificial texture into halftoned images. The use of dither matrices generated from arbitrary textures enables us to emboss images with a desired texture and to imitate art media. In this work we extend the texture based halftoning and adapt the display to 3D information. These images are printed at 300 dpi.

just the geometry or the photometric qualities of their designs

Much of the effort in NPR is directed towards imitation of traditional art materials and techniques: *sumi-e* [27] and watercolor [5] paintings, pen and ink illustrations [22, 35, 36, 23], pencil drawings [30], engravings[20, 7, 8] and copper plates[14]. While the modeling of art materials allows designers to generate images that resemble traditional illustration styles, we share the view of Landsdown and Schoffield[13] and believe that it is more important to create rendering styles specific to computer generated artistic display. The examples of previous work in this area are expressive marks by Landsdown and Schoffield[13] and abstract image representations of Haerberly[10].

Beyond the creation of a pleasing display, the objective of non-photorealistic image synthesis is to improve the viewer’s comprehension of geometric objects. This problem was first addressed in the context of comprehensive wire-frame rendering by Appel et. al[1]. and improved by Kamada and Kawai[12]. Dolley and Cohen [6] further enhanced the treatment of line drawings by controlling the line style according to importance tags and illustration rules. Strothotte et. al[29]. 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[21] 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.

Even though many NPR algorithms generate black and white im-

ages they do not use the results of conventional halftoning. The issue of tone reproduction is either ignored [1, 12, 6, 29] or is addressed by an alternative often computationally expensive technique [14, 20, 7, 8, 22, 35]. In this work we follow a different approach and extend the conventional halftoning to non-photorealistic rendering. A similar approach was taken by Ostromoukhov and Hersch[17] who modified ordered dithering and enabled artists to design single screen elements. Buchanan[4] introduced a variety of halftoning textures by altering various parameters in a clustered error diffusion method. The halftoning techniques have been used in comprehensive rendering. Sloan[25] improved the display of image gradient by constructing a dither screen from directional dither matrices. Streit and Buchanan[28] developed an importance driven halftoning system and controlled the placement of graphics elements by an importance function.

In our previous work [34] we explored the property of the ordered dithering to shape the halftoning texture through the arrangement of thresholds in the dither matrix. We generated dither matrices by processing a texture image with the adaptive histogram equalization (AHE) algorithm. Our technique enables designers to emboss the halftoned image with a desired texture or to imitate traditional illustration styles (Figure 1, center). This research is the extension of the previous results to comprehensive rendering of 3D scenes. We start by developing algorithms to control the halftoning texture with external parameters. These parameters are stored in auxiliary images and define texture features at every pixel location. Further, we use auxiliary buffers that represent 3D and subjective user derived information to tailor texture features and thus to enhance image display (Figure 1, bottom).

3 Control of texture shape

In this section we discuss the control of texture shape by defining a dither matrix with the procedural textures.

Previous research in halftoning [2, 15, 33] identified the following essential properties of the threshold distribution in a dither matrix:

1. *Uniform distribution of threshold values.*
A dither screen should contain the same number of pixels with the same values. 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 regions of the image. The homogeneity property is satisfied automatically when a small dither matrix is tiled to generate a dither screen.

We demonstrated[34] that these distribution properties can be approximated by processing images with the adaptive histogram equalization algorithm. In this research we generalize the distribution properties of dither screens and present a procedural technique for the design of dither matrices. We start by expressing the property of uniform distribution of values. Similarly to the conventional dither matrices we satisfy the homogeneity property by constructing dither screens using nearly periodic extensions of the base procedural textures.

Let $\tau(s, t)$, $0 \leq \tau(s, t) \leq 1$ be a function defined on the unit box $[0, 1] \times [0, 1]$. We define $\tau_{a,b}$ to 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 and n , where $n > 1$ and $0 \leq i < n$.

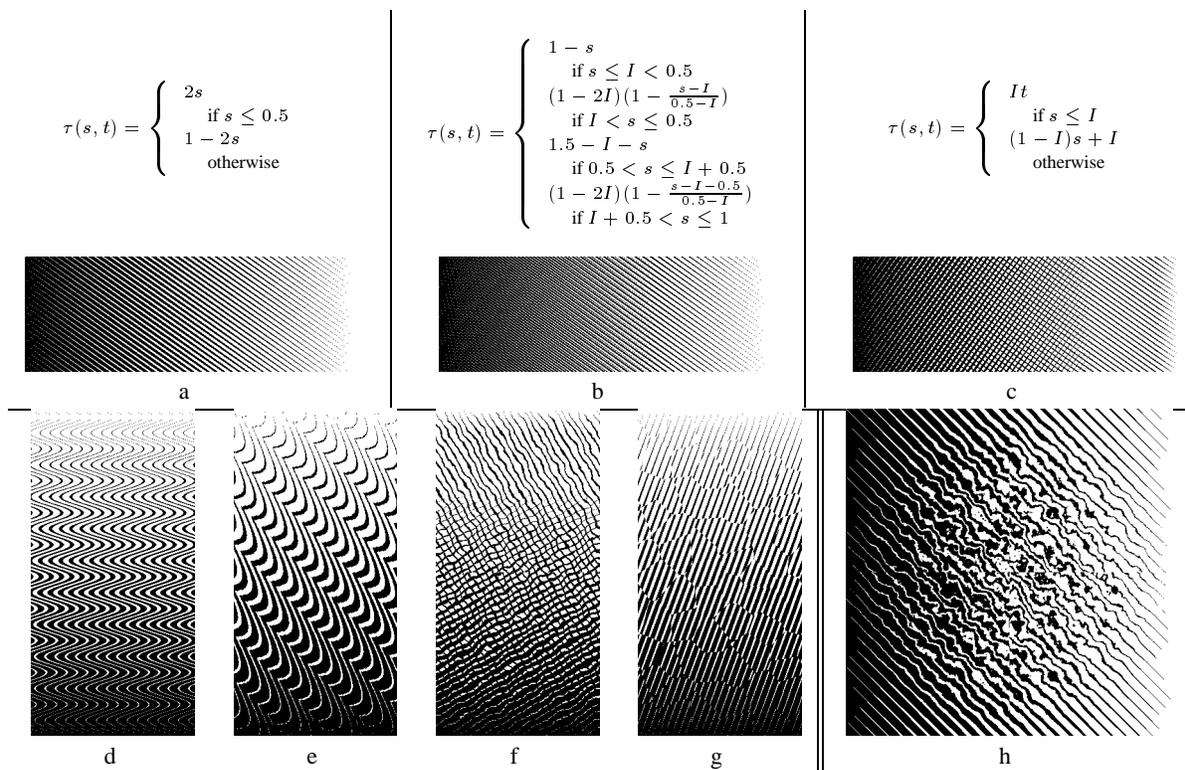


Figure 2: Procedural textures are based on combinations of scaled linear ramps (top). Texture variation is achieved by perturbing coordinates with displacement maps (bottom, d-g). Local texture shape is locally controlled by making the amplitude of the displacement to be determined by external parameter (e.g. elliptic ramp, image h).

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

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

In this work procedural dither matrices are based on a linear ramp function $\tau(s) = s$. Clearly, this function satisfies the requirement for uniform distribution of output values and thus can be used as a dithering matrix. Figure 2 demonstrates 1D (images a,b) and 2D (c) combinations of scaled ramps resulting in a variety of line and cross-hatching textures.

Unlike image based textures, procedural textures allow for local control of their shape. Texture variation across the images is achieved by perturbing the base τ function with displacement maps $D_s(s, t)$ and $D_t(s, t)$ as follows:

$$\mathcal{T}(s, t) = \tau(s + D_s(s, t), t + D_t(s, t)). \quad (2)$$

The requirement for homogeneous distribution of texture values in the dither screen limits displacement maps to smooth piece-wise linear continuous functions. Mathematical functions such as $\sin(s)$ or $|\sin(s)|$ (Figure 2, d and e respectively) satisfy this requirement and thus can be used as one dimensional displacement maps. Also, procedural Perlin's noise [19] and Worley's textures [37] are examples of two dimensional displacement maps (Figure 2, images f and g).

The use of displacement maps allows us to adapt the shape of the procedural halftoning texture to the external parameters. The external information is represented in the form of an image that controls

the amplitude of texture displacement for every pixel. Thus, the elliptic ramp increases the waviness of the line texture at the center (Figure 2, image h).

To summarize, we extended the previous research by generalizing the properties of the dither matrix and by constructing procedural textures that approximate these properties. We control the shape of procedural textures by a base function τ and by perturbing it with displacement maps. Thus, not only we defined the shape of the halftoning texture, we are also adapted this shape to the local parameters. In the rest of this paper, we use dither matrices generated by both procedural and the previous image based[34] approaches.

4 Texture Direction and Scale

In the previous section we presented a technique that generates dither matrices with desired textures. In conventional halftoning a dither matrix is tiled across the image to define a dither screen. We consider tiling to be only an example of a mapping function and use other mappings to control direction and scale of the halftoning texture. Our objective is to find a texture mapping function $\mathcal{M}(u, v)$ that matches direction and scale specified by auxiliary images.

The following observation is important in the definition of the mapping function. Continuous change of texture scale corresponds to non-zero derivatives 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, the control of direction and scale is dependent on each other and must be considered simultaneously. In this work we limit ourselves to the changes of scale and direction in the piece-wise constant fashion. Thus, direction and scale are controlled independently by piece-wise con-

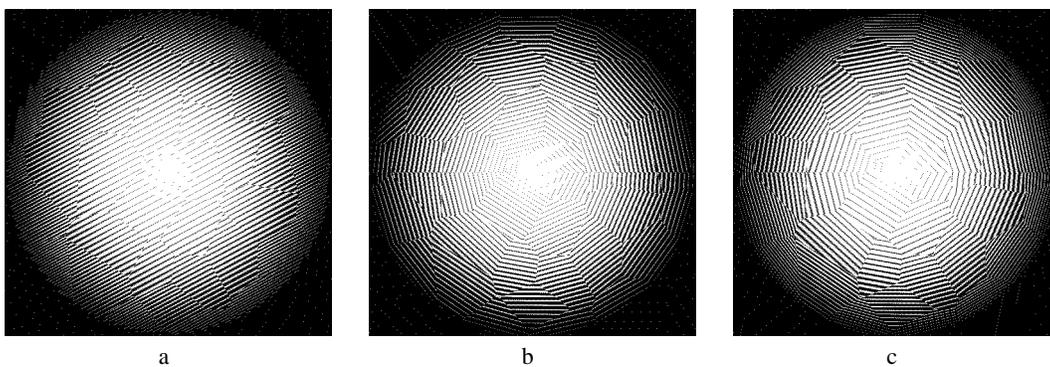


Figure 3: Texture scale and direction are tailored to the elliptic ramp in a piece-wise constant fashion. Texture scale in images *a* and *c* is proportional to the gray scale values of the ramp. Texture direction in *b* and *c* is perpendicular to the gradient of the ramp. The piece-wise constant parameters enable us to control both scale and direction independently. Unfortunately, the use of piece-wise parameters results in texture discontinuities.

stant auxiliary images.

Let $s(u, v)$ and $\theta(u, v)$ be the scale and direction specified by control images. Then, a dither matrix function $\mathcal{T}(s, t)$ is mapped to a dither screen $T(u, v)$ by the following mapping $M_{s,\theta}(u, v)$:

$$T(u, v) = M_{s,\theta} \circ \mathcal{T} = \begin{pmatrix} s(u, v)(\cos(\theta(u, v))u - \sin(\theta(u, v))v, \\ s(u, v)(\sin(\theta(u, v))u + \cos(\theta(u, v))v) \end{pmatrix} \quad (3)$$

Figure 3 demonstrates the use of the mapping function $M_{s,\theta}(u, v)$. Images *a* and *b* are controlled by a single auxiliary image that specified scale and direction respectively, whereas both scale and direction are controlled by two images in the example *c*. Because we used piece-wise constant control images the resulting texture varies its features in a step-like fashion with discontinuities on the boundary of constant regions.

5 Control of texture contrast

The main advantage of ordered dithering is that the distribution of pixels is defined a priori by the dither matrix. We use this property to introduce a desired texture. However, the threshold distribution properties do not guarantee the accurate approximation of image tone. Moreover, when a bad approximation for a pixel is generated there is no way to compensate for the error. Error diffusion techniques[9, 16], 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 context of photorealistic halftoning[3, 32].

In previous work[34] we used parameterized error diffusion to control the contrast and clustering properties of the halftoning texture. Here we extend the control of texture contrast to vary across the image. We vary the amount of error diffusion by an external parameter $\alpha_{i,j}$ where $\alpha_{i,j} \in [0, 1]$. Thus, the input pixel $g_{i,j}$ is set to 0 or 1 using the following formula:

$$b_{i,j} = \begin{cases} 0 & \text{if } g_{i,j} + \alpha_{i,j}E_{i,j} \leq t_{i,j} \\ 1 & \text{otherwise} \end{cases} \quad (4)$$

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, we implemented Floyd-Steinberg algorithm[9] with serpentine processing of rows.

The use of error diffusion does not only improve the accuracy of tone reproduction but also reduces texturing effects of ordered dithering. As the result, the maximal error diffusion with $\alpha = 1$ corresponds to the minimal texture contrast. By using external information to define values of the α parameter we control texture contrast and tone across the image. When these texture features are adapted to the external information they can provide additional visual cues in comprehensive halftoning (Figure 4).

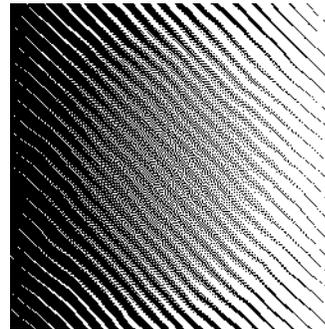


Figure 4: The amount of error diffusion is controlled by the elliptic ramp with maximum at the center. The resulting texture contrast is thus the weakest at the center. Also, tone approximation is directly dependent on the amount of error diffusion and is thus the highest at the center. This image is printed at 150 dpi.

6 Tailoring texture to 3D information

In the previous sections we presented a technique enabling us to tailor shape, direction, scale, and contrast of the halftoning texture to the auxiliary image. Here we apply this technique to 3D rendering and will adapt the halftoning texture to the geometry and illumination of the scene.

We follow an image based rendering approach proposed by Saito and Takahashi[21] and use *g*-buffers to store 3D information. In this research the following *g*-buffers are created by a raytracer as an intermediate step of the photorealistic rendering: *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.

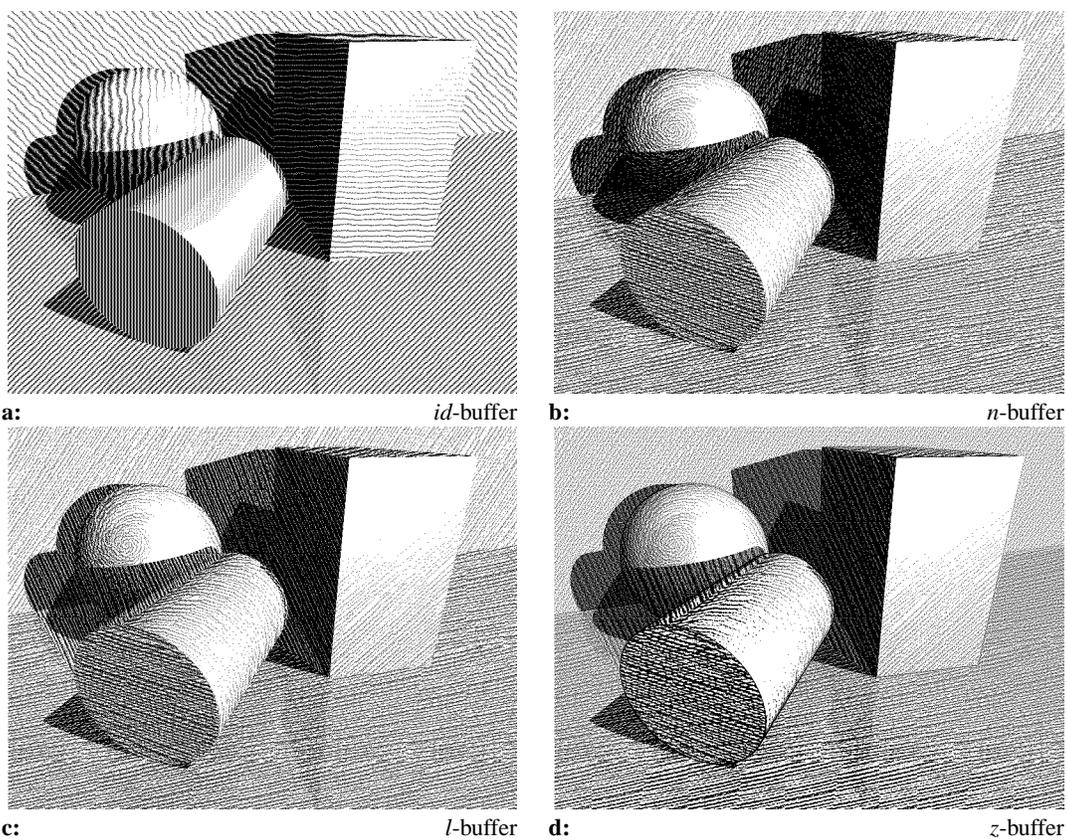


Figure 5: 3D information is stored in g buffers and is used to control the halftoning texture. The id -buffer identifies and segments objects of the scene. The n -buffer aligns texture direction with the direction of projected object normals. The l -buffer differentiates direct and diffused illuminated surfaces. The z -buffer stores depth information and enables us to enhance depiction of scene's depth and to delineate silhouettes and ridges of the objects. These images are printed at 300 dpi.

n -buffer Pixels of the n -buffer store surface information of the corresponding objects. Normals of the object are 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 diffuse light. Separate l -buffers are constructed for every light source in the scene.

z -buffer and its derivatives Distances between the camera and objects are stored in the z -buffer. Saito and Takahashi[21] pointed out that the first and second derivatives of the z -buffer represent silhouettes and ridges of the objects. Thus, we compute silhouettes z' -buffer and ridges z'' -buffer by differentiating the z -buffer.

We demonstrate the use of the g -buffers by halftoning a scene with the simple geometric objects: cylinder, sphere, and box. (Figure 5). The id -buffer controls texture scale, waviness, and direction of procedural texture in Figure 5a. 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 Haeberly[10] suggested aligned display texture to the projection of normals onto the viewing plane. Thus, we control the direction of the pencil stroke texture with the n -buffer in Figure 5b. 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 surfaces. We use the l -buffer to change the shape of the halftoning texture in the shadow of one of the light sources (Figure 5c). Thus, object surfaces that are not illuminated by the light source are unified by the halftoning texture.

Artist's often reduce the contrast of the background objects to enhance viewer's appreciation of the scene depth[26]. We have chosen to soften the halftoning texture of distant scene surfaces by increasing the amount of error diffusion proportionally to the z -buffer (Figure 5d). We also used the z -buffer derivatives to find silhouettes and ridges. Merging the z' and z'' -buffers into the dither screen results in the highlighting of object edges (Figure 5d).

The scene information used thus far is objective and is generated by the raytracing software. We extend the technique further by allowing a designer to tailor the display to the additional subjective information such object importance or completeness.

Image based rendering techniques lend themselves to the incorporation of the subjective information by the use of the importance maps. These importance maps represent the relative importance of image regions and are used in the combination with the g -buffers.

Consider the example of a foot in Figure 6. 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. The display of the most important bones is enhanced by the alignment of halftoning texture to the bone geometry stored in the n -buffer. The other bones are dis-



Importance image

The viewer's attention should be focused on the darkest two bones on the front.

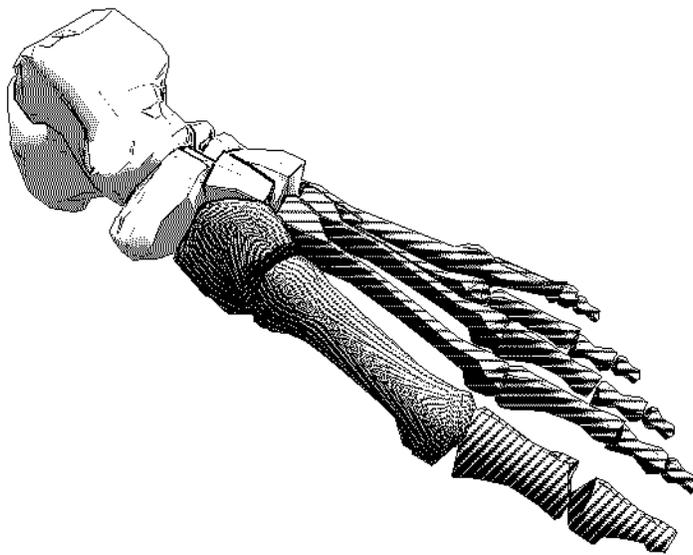


Figure 6: The importance map is used to control textures in the display of the three-dimensional model of the foot bones. The most important bones is halftoned with textures aligned by the n -buffer. The id -buffer controls the shape, scale, and direction of texture for the display of other bones.

played with 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 the application of the texture control techniques to the enhancement of the scene display. These techniques tailored texture features to the scene information stored in g -buffers and importance maps. Thus, our research enabled designers to use the halftoning texture for aid viewer's comprehension of the image.

7 Conclusion

Our research deals with rendering of three-dimensional scenes using binary primitives. Unlike the conventional halftoning techniques we followed a non-photorealistic approach and used the halftoning texture to enhance the display. In particular we adapted texture shape, direction, scale, and contrast to the 3D scene information and user specified importance maps.

Our work is based on the property of the ordered dithering algorithm to define the appearance of the halftoning texture. We developed techniques that generate dither screens with controlled texture features. We started by producing a dither matrix from an arbitrary image[34] or by a procedural texture. The resulting dither matrices were mapped to the dither screen taking into account the scale and direction information provided by the external parameters. We approximated image tones using the dithering screen in the combination with partial error diffusion. The local amount of error diffusion is controlled by an auxiliary image and effects texture contrast.

We used texture control techniques to enhance the display of 3D scenes by adapting texture features to the g -buffers and importance maps. The resulting halftoning texture is tailored to the image content and is thus aids in viewer's comprehension of the display.

References

[1] A. Appel, F. J. Rohlf, and A. J. Stein. The haloed line ef-

fect for hidden line elimination. In *Computer Graphics (SIGGRAPH '79 Proceedings)*, volume 13, pages 151–157, August 1979.

[2] B. E. Bayer. An optimum method for two-level rendition of continious tone pictures. *IEEE International Conference on Communications*, 1:26–11–26–15, June 1973.

[3] 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.

[4] J. W. Buchanan. Special effects with half-toning. In *Proceedings of Eurographics 96*, pages 97–108, 1996.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

[9] Robert W. Floyd and Louis Steinberg. An adaptive algorithm for spatial greyscale. *Proceedings of the Society for Information Display*, 17(2):75–77, 1976.

- [10] Paul E. Haeberli. Paint by numbers: Abstract image representations. In *Computer Graphics (SIGGRAPH '90 Proceedings)*, volume 24, pages 207–214, August 1990.
- [11] 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.
- [12] Tomihisa Kamada and Satoru Kawai. An enhanced treatment of hidden lines. *ACM Transactions on Graphics*, 6(4):308–323, October 1987.
- [13] J. Landsdown and S. Schofield. Expressive rendering: A review of nonphotorealistic techniques. *Computer Graphics and Applications*, 15(3):29–37, 1995.
- [14] W. Leister. Computer generated copper plates. *Computer Graphics Forum*, 13(1), 1994.
- [15] 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.
- [16] 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.
- [17] Victor Ostromoukhov and Roger D. Hersch. Artistic screening. In *SIGGRAPH 95 Conference Proceedings*, pages 219–228. Addison Wesley, August 1995.
- [18] 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.
- [19] Ken Perlin and Eric M. Hoffert. Hypertexture. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 253–262, July 1989.
- [20] Y. Pnueli and A. M. Bruckstein. DIG^i_d ürer – a digital engraving system. *The Visual Computer*, 10:277–292, 1994.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] 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.
- [26] Harold Speed. *The Practice and Science of Drawing*. Steeley, Service and Co. Ltd., London, 1917.
- [27] Steve Strassmann. Hairy brushes in computer-generated images. M.sc. thesis, MIT Media Laboratory, June 1986.
- [28] Lisa Streit and John Buchanan. Importance driven halftoning. In *Proceedings of Eurographics 98*, pages 207–217, August 1998.
- [29] 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.
- [30] Mario C. Sousa and John W. Buchanan. Computer-Generated Graphite Pencil Rendering of 3D Polygonal Models. *Submitted for publication*, 1999.
- [31] 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.
- [32] Robert Ulichney. *Digital Halftoning*. MIT Press, 1987.
- [33] 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.
- [34] Oleg Veryovka and John Buchanan. Halftoning with image based dither screens. Submitted for publication to *Graphics Interface 99*. The paper can be downloaded at <ftp://ftp.cs.ualberta.ca/pub/oleg/GI-99.ps.gz>.
- [35] 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.
- [36] 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.
- [37] 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.

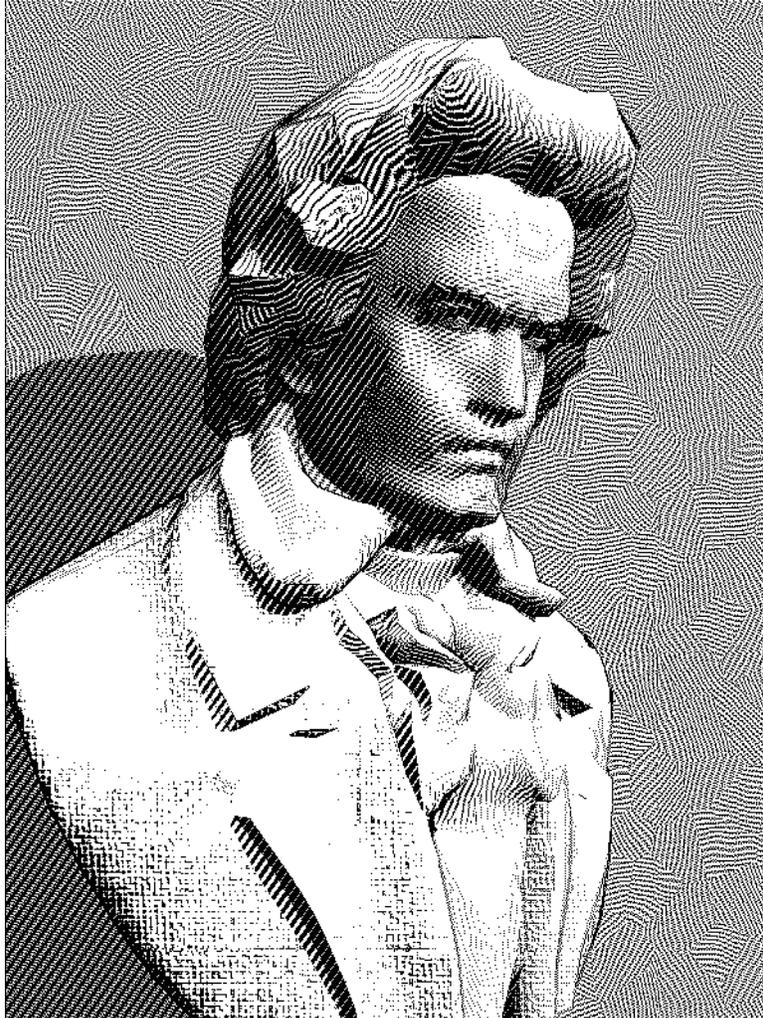


Figure 7: Both procedural and image based textures are used in halftoning of the model of the Beethoven's bust. Textures of the skin, hair, the shirt, and the tie are controlled by the n -buffer and are scaled using the id -buffer. The suite is displayed with the image based canvas texture. The background 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.