

Printing Reflectance Functions

TOM MALZBENDER and RAMIN SAMADANI

Hewlett-Packard Laboratories

STEVEN SCHER and ADAM CRUME

University of California Santa Cruz

DOUGLAS DUNN

3M

and

JAMES DAVIS

University of California Santa Cruz

The reflectance function of a scene point captures the appearance of that point as a function of lighting direction. We present an approach to printing the reflectance functions of an object or scene so that its appearance is modified correctly as a function of the lighting conditions when viewing the print. For example, such a “photograph” of a statue printed with our approach appears to cast shadows to the right when the “photograph” is illuminated from the left. Viewing the same print with lighting from the right will cause the statue’s shadows to be cast to the left. Beyond shadows, all effects due to the lighting variation, such as Lambertian shading, specularity, and inter-reflection can be reproduced. We achieve this ability by geometrically and photometrically controlling specular highlights on the surface of the print. For a particular viewpoint, arbitrary reflectance functions can be built up at each pixel by controlling only the specular highlights and avoiding significant diffuse reflections. Our initial binary prototype uses halftoning to approximate continuous grayscale reflectance functions.

Categories and Subject Descriptors: B.4.2 [Input/Output and Data Communications]: Input/Output Devices—*Data terminals and printers; image display*; I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture

General Terms: Design, Experimentation

Partial funding for this work was provided from NSF no. CCF-0746690 and LANL ISSDM.

Authors’ addresses: T. Malzbender (corresponding author) and R. Samadani, Hewlett-Packard Laboratories; email: {tom.malzbender, ramin.samadani}@hp.com; S. Scher, A. Crume, University of California at Santa Cruz; email: sscher@cs.ucsc.edu, acrume@ucsc.edu; D. Dunn, 3M; email: dsdunn@mmm.com; J. Davis, University of California at Santa Cruz; email: davis@cs.ucsc.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2012 ACM 0730-0301/2012/05-ART20 \$10.00

DOI 10.1145/2167076.2167078

<http://doi.acm.org/10.1145/2167076.2167078>

Additional Key Words and Phrases: Reflectance function, reflectance field, BRDF

ACM Reference Format:

Malzbender, T., Samadani, R., Scher, S., Crume, A., Dunn, D., and Davis, J. 2012. Printing reflectance functions. *ACM Trans. Graph.* 31, 3, Article 20 (May 2012), 11 pages.

DOI = 10.1145/2167076.2167078

<http://doi.acm.org/10.1145/2167076.2167078>

1. INTRODUCTION

An object’s appearance changes when lit from different directions. Shadows are cast and specular highlights shine. Even diffuse surfaces change appearance, providing a valuable perceptual cue to object shape and material. Unfortunately, when printed on paper, this variability is lost. A photographic print represents just one appearance, regardless of the ambient lighting when the print is viewed.

A real scene’s appearance at each point depends upon the incident illumination, the surface Bi-directional Reflectance Distribution Function (BRDF) [Nicodemus et al. 1977], and the local surface orientation. Typical printers cannot print images which react to incident illumination because they use inks with a limited range of BRDFs and print on paper which is flat, having a single orientation everywhere.

Simply expanding the range of specularity of the printer’s inks is insufficient. Surface orientation plays an important role in controlling the observed intensity of reflected light. To correctly represent this interaction with light, the local surface orientation of the paper needs to match that of the original object.

A conceptually simple solution would be to orient the local surface normal of each pixel appropriately, and to print the object’s surface BRDF onto this pixel. However, this requires changing the physical shape of the underlying paper at the time of printing. This is an expensive process that requires equipment far more complex than just depositing ink onto a surface.

We introduce a method of printing reflectance functions [Debevec et al. 2000] which makes use of paper with a static microgeometry structure, shown in Figure 1(b). The paper consists of a hexagonal array of spherical depressions that have a specularly reflective surface. By selectively printing opaque or partially opaque ink on portions of that surface, we can control whether a specific incident lighting direction returns a specular highlight, and to what degree. Because both the paper and ink are designed to minimize diffuse reflections, the specular return fully controls the appearance of this “reflectance paper” and is used to control its appearance as a




Fig. 1. (a) We introduce a method for printing reflectance functions, images which correctly respond to the direction from which they are illuminated. (b) A microgeometry reflective substrate returns a specular highlight unless opaque ink is deposited on the point corresponding to a particular incident illumination direction. (c) Our prototype is shown displaying different images in response to two different illumination conditions. Image used with the permission of Cultural Heritage Imaging.

function of lighting direction. Note that this scheme gives sufficient expressive power to specify two dimensions of the 4D BRDF. We choose to specify, for a single fixed viewing direction, how much light will be reflected from each incident lighting direction, thereby building up an arbitrary reflectance function.

The primary contribution of this work is a method for printing reflectance functions using existing printers and special paper. We support this contribution with a ray traced simulation and analysis of errors, as well as with an initial prototype implementation.

2. RELATED WORK

Several research groups have demonstrated passive physical surfaces and devices to generate controlled reflectance functions. Fuchs et al. [2008] share our goal of generating the reflectance field of a scene, and built an optical assembly that produces lighting-dependent imagery using a lenticular array. Light entering a lenslet from behind the print will pass through a different location depending on its entering angle. A more complex optical assembly is presented that has both view and lighting dependence, albeit at low resolution. These methods generate a back-lit image as opposed to our printing process, which generates controlled reflectance functions designed to be viewed in a reflective manner, like a conventional print.

Matusik et al. [2009] use a printing process, as we do, to generate images with a varying BRDF by linear blends of a set of inks and foils. In their work, BRDFs are associated with normals perpendicular to the surface of the paper, and thus do not afford enough control to mimic the reflectance functions of an arbitrary 3D scene.

Alexa et al. [2010] share one of our primary goals, namely to develop a mechanism to generate arbitrary images as a function of varying lighting direction by controlling the reflectance function of a surface. In their work, diffuse reflection is used as a basis for generating images, which results in the ability to specify two samples of the reflectance function in practice, and thus two approximate output images. By using only specular reflections to build up arbitrary reflectance functions and a 32×32 printed matte overlaying each dimple, our design can achieve greater independence between many more images.

Regg et al. [2010] also build up an image using only specular highlights, as we do. Their goal is not to mimic the reflection functions of a real scene, but to create a set of points with a controlled

3D percept. A curved scratch on acrylic generates highlights at two separate locations along the scratch, for the viewpoints associated with each eye. A set of such curves are etched under computer control, each yielding a perceived 3D point location to build up percept of the desired 3D shape.

Several researchers have also explored the possibility of milling a homogenous material's surface to achieve a desired effect. The milling induces a variation of the material's appearance by setting the orientation of each surface patch to reflect more or less light in desired directions. Such methods can be used to automate the design of bas relief in an approach that preserves important 3D depth cues [Weyrich et al. 2007]. Weyrich et al. [2009] mill a custom array of microfacets to control the overall BRDF of a patch of aluminum, effectively treating it as a single pixel. In addition, the authors validate their ability to generate a complex BRDF by viewing reflected light off that patch on a secondary surface. Our method extends their ability to control reflectance from a single large pixel to a 2D image and differs in that our pixels are three orders of magnitude smaller in area, and are produced with a printer rather than a milling machine.

Lastly, Nayar et al. [2004] present an active approach to varying the appearance of a photograph under lighting changes by measuring the 2D illumination field, and then generating correspondingly lit images. We strive to achieve similar functionality with a passive surface.

3. REFLECTANCE FUNCTIONS

2D reflectance functions can be measured easily for real scenes by taking photographs lit from many different lighting directions, and a wide variety of devices have been constructed to do so [Debevec et al. 2000; Malzbender et al. 2001; Peers et al. 2006; Wenger et al. 2005]. In principle this function can contain high-frequency components from specular highlights and hard shadows. However, in practice even low-order approximations are sufficient for a variety of presentation and evaluation purposes [Hawkins et al. 2001; Ramamoorthi et al. 2001]. Manipulations of reflectance functions have proved useful especially in both cinema applications [Debevec et al. 2000; Peers et al. 2007] and archeological applications [Freeth et al. 2006; Mudge et al. 2006].

Figure 2 shows several examples of reflectance functions for individual pixels, parameterized by the incoming illumination angle.




Fig. 2. The shape of a reflectance function (bottom) is dependent on the orientation and shadowing of the underlying surface. For example, the indicated pixels of the almonds on the right are chosen on opposite edges, so respond most brightly to light from opposing angles.

On the left is a set of colored leaves, with the reflectance of specific marked pixels shown below. On the right are pixels chosen on opposite edges of almonds. Note that shadows are present in both cases at extreme lighting angles, and this shows up as dark regions in the reflectance function. Although the reflectance functions are very smooth, they are oriented differently at different pixels, and this is sufficient to provide a convincing perception of object shape when the incident light angle is changed. Flat paper printed with standard inks can only represent centered isotropic functions similar to the red pixel in this example. Our design can represent arbitrary functions. Note that for the photographically collected reflectance functions we employ, global illumination effects such as shadows and interreflections are present at the time of image capture, and thus are represented in the reflectance function. Whether such secondary effects are modeled is dependent on the generation of the reflectance data that is input to our system, whether captured or simulated.

4. REFLECTANCE PAPER

We designed our reflectance paper substrate with spherical dimples, which we metalized to produce a mirrored finish. Since the surface is entirely specular, light incident on the paper from some particular direction will strike the entire surface, but only light that strikes points with a corresponding particular orientation will be reflected toward the viewer. Opaque ink covering this point on the dimple prevents the dimple from being illuminated from this direction, as seen in Figure 1(b).

Note that achieving sensitivity to the full hemisphere of possible lighting directions requires a reflecting surface consisting of only 1/4th spheres, not hemispheres. To simplify the process of printing on a geometrically complex surface even further, we constructed media consisting of 1/5th spheres, subtending approximately 70 degrees, which will respond to 140 degrees of the possible 180 degrees of incoming lighting directions. Lighting outside of this range will not change the appearance of the print measurably.

Our initial prototype considers only binary, black and white printing using fully opaque ink. With this constraint, levels of gray may be approximated through dithering, as discussed in Section 9. To achieve true continuous tone one could vary the thickness or transparency of the ink to modulate the reflectivity instead of providing a binary opaque mask as we did.

We have described this design using an orthographic directional light source for simplicity of description. The design itself is theoretically correct and functions correctly under arbitrary




Fig. 3. To avoid artifacts due to ink specularly we print reflectance functions on an overlaid transparency instead of directly on the metalized dimples themselves.

illumination environments due to the linearity of light transport. This includes area lights and near field point illuminants, as well as more general lighting environments.

Our reflectance paper builds arbitrary reflectance functions from specular highlights off a dimpled substrate. This is a dramatically different mechanism from conventional paper, which uses primarily diffuse reflections off paper and ink. The question arises how brightness compares between the two cases. For conventional white paper, assuming ideal diffuse reflection, all incoming light energy from any direction is distributed uniformly across possible viewing directions. This is also the case in our dimpled reflective substrate, assuming ideal specular reflections. Across a single dimple, our facet microgeometry distributes all incoming light uniformly across viewing direction. In both cases ink serves to attenuate the brightness spatially. We therefore expect reflectance paper prints to have comparable brightness to prints on conventional white paper.

Contrast ratio in our design is limited by the performance of the inks. In particular, any diffuse reflections introduced by either the dimpled metallic substrate or the inks is distributed across all viewing directions, and thus serves to decrease the contrast ratio of our passive “display”.

5. AVOIDING INK SPECULARITY

One design deposits ink directly on the dimples, blocking the specular highlights of the reflective substrate, without introducing diffuse contributions that would lower the contrast ratio. Dry electrophotographic toner, used in laserjet printers, or Liquid ElectroPhotographic (LEP) inks, used in commercial printing, are unfortunately quite specular themselves. We were able to reliably deposit LEP inks successfully in a controlled manner on our dimpled substrate, even on large 2.3mm diameter dimples, but the specularly of the ink itself causes specular returns from lighting directions the ink is intended to block.

To achieve a more matte finish, matte topcoats could be applied, but as these operate by introducing microfacets that spread the surface reflections to a wider range of reflection angles; they create new problems. Although effective at reducing specular highlights off the ink, this also distributes the energy from the uncovered specular region in a diffuse manner, which unacceptably lowers the contrast ratio of our reflectance paper.

We have developed another approach to avoid the specular return from the ink. Instead of printing directly on the metalized dimples themselves, we use a standard laser printer to print a mask on a transparent film that is placed on top of the dimpled surface, as shown in Figure 3. In this approach, specular highlights off the inks on the transparency film are in the reflection direction given by the normal of the film, not the normals associated with the shape of the metalized dimple. These unwanted specular highlights are generally never directed towards the viewer, and therefore avoided. This allows us to avoid the difficulty of inventing inks that do not have significant specular return.




Fig. 4. The geometry of a single metalized dimple with an overlaid transparency.

6. GEOMETRY OF THE SPECULAR DIMPLES AND TRANSPARENCY FILM

The geometry for a single dimple is shown in Figure 4. A right-handed coordinate system with origin o is placed at the center of the sphere. The dimple shape is a portion of the lower hemisphere, satisfying

$$x^2 + y^2 + z^2 = r^2, \quad (1)$$

$$z \leq z_0, \quad (2)$$

with the top of the grid of dimples, where the transparency film may be placed, at $z = z_0$. The value $z_0 = -\cos(\eta/2)$ is determined by the angle η defining the proportion of hemisphere formed by o and two opposite points on the dimple at the grid surface ($z = z_0$). Using the implicit representation of the sphere, $F(x, y, z) \equiv x^2 + y^2 + z^2 - r^2 = 0$, the (upward pointing) normal \mathbf{n} at a point $\mathbf{s} = (x, y, z)^T$ on the dimple is given by

$$\mathbf{n} = \frac{-\nabla F(\mathbf{s})}{\|\nabla F(\mathbf{s})\|} = \frac{-\mathbf{s}}{r} = \left(\frac{-x}{r}, \frac{-y}{r}, \frac{-z}{r} \right)^T. \quad (3)$$

To simplify notation for the subsequent derivations we set $r = 1$ without loss of generality. For a given view direction \mathbf{v} , any point \mathbf{s} on the specular dimple surface responds to a single lighting direction, $\mathbf{l}(\mathbf{s})$. The direction \mathbf{l} corresponds to a specular reflection of \mathbf{v} through the normal vector \mathbf{n} at \mathbf{s} . In other words,

$$\mathbf{l}(\mathbf{s}) = (2\mathbf{nn}^T - \mathbf{I})\mathbf{v} = (2\mathbf{ss}^T - \mathbf{I})\mathbf{v}, \quad (4)$$

where, with $r = 1$, we obtain the right-hand term in Eq. (4) by substituting $\mathbf{n} = -\mathbf{s}$ from Eq. (3). In Eq. (4) Matrix \mathbf{I} is the 3×3 identity matrix, and the vectors \mathbf{l} , \mathbf{v} , and \mathbf{s} are unit vectors, with the convention that \mathbf{v} and \mathbf{l} point towards the viewer and the light, respectively. For use in the subsequent discussion, we define matrix $\mathbf{M} \equiv 2\mathbf{ss}^T - \mathbf{I}$. The components of $\mathbf{l}(\mathbf{s}) = (l_x, l_y, l_z)^T$, for a general view direction $\mathbf{v} = (v_x, v_y, v_z)^T$, are given by the following.

$$\begin{pmatrix} l_x \\ l_y \\ l_z \end{pmatrix} = \begin{pmatrix} 2x^2 - 1 & 2xy & 2xz \\ 2xy & 2y^2 - 1 & 2yz \\ 2xz & 2yz & 2z^2 - 1 \end{pmatrix} \begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix} \quad (5)$$

Since each point \mathbf{s} on the dimple responds to a single lighting direction, each individual dimple samples the *lighting space* or reflectance function at one spatial location, and the set of dimples on the hexagonal grid provide a *spatial sampling* of the reflectance function.

Our prints are designed to be viewed from the direction normal to the top surface of the grid of dimples. In the coordinate system shown in Figure 4, the normal viewing direction is $\mathbf{v} = (0, 0, 1)^T$. When this view direction is substituted into Eq. (5), the resulting lighting directions, as a function of the dimple point coordinates, are given by

$$\mathbf{l}(\mathbf{s}) = \begin{pmatrix} l_x \\ l_y \\ l_z \end{pmatrix} = \begin{pmatrix} 2xz \\ 2yz \\ 2z^2 - 1 \end{pmatrix}. \quad (6)$$

This equation implies rotational symmetry for the normal viewing direction. Rotating any point on the dimple around the z axis rotates the corresponding lighting direction by the same amount around the z axis. This rotational symmetry means any vertical slice through the origin, for example $x = 0$, has all the information needed for analyzing the system.

7. VIEWPOINT TRANSFORMATION

Our system and prototypes are designed for a fixed view direction normal to the top surface of the grid of dimples. Real-world, near field viewing implies that the viewing direction of some dimples will be slightly away from normal. Empirically we find that changing the view direction degrades the experience very gracefully. This section confirms this empirical finding by analytically showing that a rotation of view direction has the effect of smoothly warping the lighting space function. Using Eq. (4) we see that a rotation of viewpoint $\mathbf{w} = R(\theta)\mathbf{v}$, where $R(\theta)$ is the rotation matrix (for example, around the x axis), results in a modification of lighting space given by $\mathbf{l}_R = MR(\theta)\mathbf{v}$. For the normal view direction, the components of lighting direction \mathbf{l} are given by Eq. (6). Applying a rotation around the x axis to the normal view direction results in transformed lighting directions given by

$$\mathbf{l}_R(\mathbf{s}) = \begin{pmatrix} 2xz \cos(\theta) + 2xy \sin(\theta) \\ 2yz \cos(\theta) + \sin(\theta)(2y^2 - 1) \\ \cos(\theta)(2z^2 - 1) + 2yz \sin(\theta) \end{pmatrix}. \quad (7)$$

One can show that for points where $x = 0$, $\mathbf{l}_R = R(-\theta)\mathbf{l} = MR(\theta)\mathbf{v}$ so that rotation by θ of the viewpoint results in a rotation in the opposite direction for the lighting space function (for those points where $x = 0$). The off-axis points where $x \neq 0$ result in the smooth transformations of x, y, z shown in Eq. (7).

8. ECHO IN LIGHTING SPACE FOR PRINTED FILMS

When printing on films placed on the top surface of the grid of dimples, we project the reflectance function from the dimple surface onto the transparency. For view direction normal to the transparency film surface, the remapping of reflectance values simply applies the reflectance value of the point on the dimple with the same x and y coordinates as the point on the transparency. Modifying the transparency of a point $\mathbf{p} = (x_0, y_0, z_0)$ on the film affects the transmission of light from exactly two points on the dimple, \mathbf{s}_1 and \mathbf{s}_2 , and therefore two lighting directions. By convention, we set the transparency of point \mathbf{p} to control the lighting direction labeled \mathbf{l}_1 in Figure 5, but because of the geometry, this constrains our control




Fig. 5. A single ink patch on a transparent overlay is effective at attenuating incoming light from direction I_1 , but inadvertently also attenuates a single other light direction, that from I_2 .

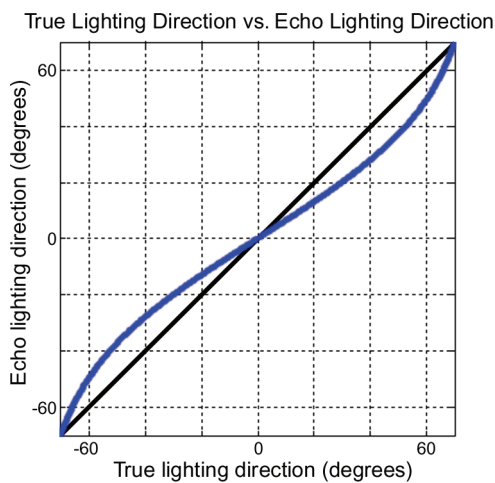


Fig. 6. The maximum offset between a desired, blocked light ray, I_1 and its “echo”, I_2 is at worst 12.9 degrees in our roughly 1/5th sphere prototype.

over lighting direction labeled I_2 , which we call the echo direction. In the dithering section we will show an adaptive, data-dependent algorithm to reduce the effects of this echo constraint.

The two points on the dimple satisfy independent conditions

$$\mathbf{s}_1 = \mathbf{p} - k_1 \mathbf{v} \quad (8)$$

$$\mathbf{s}_2 = \mathbf{p} - k_2 \mathbf{l}_2 \quad (9)$$

for constants k_1 and k_2 that determine the intersection points on the dimple. Intuitively, s_1 is found by *looking* from viewpoint \mathbf{v} through point \mathbf{p} to the intersecting point s_1 on the dimple. The second point, s_2 , is determined indirectly by the condition that the blocked lighting direction I_2 goes from the intersecting point s_2 on the dimple to the point \mathbf{p} . Thus a single point on the transparency designed to block lighting direction I_1 also blocks lighting direction I_2 , as shown in Figure 5. Note that this echo occurs in lighting space, within a dimple, not spatially across dimples. Figure 6 plots the angular deviation between the primary and echo directions. For the dimples used in our prototype, which subtend 70 degrees, the maximum angle between the primary and echo lighting angles is 12.9 degrees. We control the echo by designing the dither pattern to take them into account.



Fig. 7. Simply thresholding the grayscale reflectance function (top) to black and white (bottom) produces images which relight correctly, but with the expected quantized binary response.

9. DITHERING

Our current method of printing is limited to black ink, with no access to true grayscale inks. Simply thresholding the reflectance function to black and white, as was done in our initial prototype, produces images which relight correctly, but appear metallic, as shown in Figure 7. This section describes halftoning techniques for better rendering on the reflectance paper system of Figure 3. Halftoning techniques are well developed for traditional printing [Arce 2008; Ulichney 1987] and have been employed in reflectance functions [Matusik et al. 2009]. The reflectance paper halftoning problem is interesting because the reflectance function has statistics different from those of traditional images and the halftoning may be used to reduce the effects of the echo shown in Figure 5.

9.1 Spatial and Lighting Space Dither

Within each dimple in the printed film of Figure 3 are samples from different lighting directions, and the set of dimples arranged spatially in a hexagonal grid sample the reflectance field spatially. The film encoding reflectance functions can be halftoned, represented by binary reflectance values that preserve local grayscale average values. These are 4D printed functions, however, and care must be taken in choosing the halftone patterns. Halftoning can occur either spatially across dimples, or in the lighting domain within a single dimple, or jointly. Halftoning spatially uses a single threshold within each dimple, but the thresholds vary between the different dimples (according to a predetermined dither matrix). This results in grayscale reflectance functions when averaged spatially across dimples. When an area light source is present, halftoning in the lighting domain by varying thresholds within a dimple may improve the spatial resolution of the rendered grayscale values by averaging within a single dimple. This only occurs for area light sources. For a point light source, halftoning in the lighting domain within a single dimple has no effect since only a single ray is reflected from each dimple.

We implement spatial halftoning by using dispersed dot dither matrices designed for hexagonal grids. The dither matrices specify spatially varying thresholds that we apply across the dimples. Within a single dimple, a single threshold is used for all lighting angles, but across dimples, the threshold is spatially varying and periodic, as with traditional image halftones. We implemented a small hex mask with 3 threshold values, and a larger one with 27 threshold values [Ulichney 1987].




Fig. 8. Comparison of grayscale, black and white thresholding, and dithering, calculated at the dot pitch of our manufacturing process.




Fig. 9. Comparison of the attenuation masks produced by (b) binary thresholding and (c) an ordered dither pattern for (a) the original reflectance function.

Figure 8 compares grayscale, thresholding, and dithering simulated for a single lighting direction. The images are shown at the resolution of our physical prototype, illuminated from a single lighting angle.

Figure 9 illustrates the difference between masks produced through black and white thresholding and spatial dithering. The left image of the figure shows a portion of the original reflectance of a synthetically generated sphere. The center image in the figure shows the results of thresholding and the right image the results of halftoning (yellow borders outline the individual dimples). The hexagonal grid structure is aligned with our underlying dimples, at a 1mm scale. The thresholding shown in the center of the figure does not preserve grayscale because of the fixed threshold values but the halftone version on the right results in grayscale preservation by dithering across space with varying threshold values, producing the observed dot size differences.

9.2 Echo

For the ideal dither function printed directly on the dimples, there is no echo and each lighting direction could independently take values of 0 (opaque) or 1 (transparent). When printing on the transparency film, however, the geometry of Figure 5 shows that exactly two lighting directions are controlled by each single point on the printed film. The plot of Figure 6 shows the mapping relationship in angle between the two lighting directions controlled by point p , and this mapping determines which lighting directions interact. The plotted function is circularly symmetric so the single vertical slice shown represents the entire function. Whether an error actually occurs due to the echo interaction depends in a simple nonlinear manner on whether the lighting function dither matrix values are transparent or opaque.

To reduce the “echo” in lighting space, we modify the results of the standard halftone, which does not take echo into account, in a lighting-function-adaptive iterative second stage. We approximate the geometry of Figure 5 by using point samples, spatially quantized pixel locations on the printing film, and nearest pixel intersection

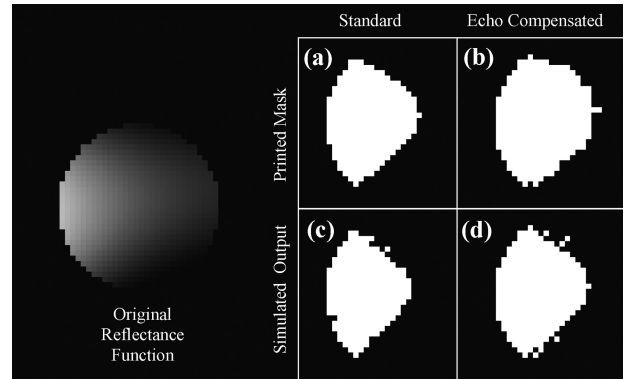


Fig. 10. Printed masks and simulated output for all lighting directions for a single dimple. (a) Standard halftone reflectance function for this dimple. (b) Echo compensated halftone, blocking fewer lighting directions. (c) Simulation output of halftone pattern in (a). (d) Simulation output for halftone pattern in (b). In both (c) and (d) the simulation shows that echoes erode the desired reflectance function. The echo compensated halftone (b) is designed so that (d) approximates (a) more closely than (c) does.

rules for determining which pixels positions interact together for the “echo” and proceed to reduce the echo. The appendix provides pseudocode for the two-stage halftoning algorithm that first generates a standard halftone and then reduces the echo in the second iterative stage.

A transparent primary halftone value does not *block* the echo halftone value, so that no change in the primary halftone value is needed. Only for those interacting pixels where the ideal dithering solution would result in primary lighting direction opaque and echo lighting direction transparent, we examine errors against the ideal first-stage halftone values and flip the dither matrix element to transparent when this lowers the error. This process is iterated (very few iterations are needed in practice), and the comparison is to the ground-truth dither that would be possible without “echo” when printing directly on the dimples.

In the binary search (DBS) halftoning of Allebach [2001], a halftone postprocessing step of halftone value flips and swaps is used in conjunction with a printer and vision model to iteratively improve halftone quality. Our approach differs in motivation; we aim to reduce “echo”, and in details, but at a conceptual level there are some similarities between our approach and DBS.

Figures 10 and 11 illustrate the effect of accounting for the echo interactions in the halftone design. Figure 10 maps a single dimple’s appearance when lit from each lighting direction, while Figure 11 shows an entire image lit from a single direction.

Figure 10 shows the original grayscale reflectance function, along with printed masks and the simulated outputs. The binary image in 10(a) shows the standard “halftone” for this dimple, determined by thresholding the grayscale with a fixed threshold (for this dimple). 10(b) shows the echo compensated halftone pattern that results by applying the second stage of the algorithm shown in the Appendix. In the bottom row of Figure 10 the outputs of simulations that model the echo interactions are shown. Figure 10(c) shows the simulation results when we apply the standard halftone, and Figure 10(d) shows the simulation results when echo compensation is applied. Figure 10(d) approximates Figure 10(a) more closely than does Figure 10(c), with 9 pixels different as opposed to 36.

Figure 11 illustrates the effect of accounting for the echo for a fixed single lighting direction for the full image. In the top row,




Fig. 11. Printed masks and simulated outputs for a fixed viewpoint and fixed lighting direction. (a) Standard half-tone. (b) Echo compensated half-tone. (c) Simulation output of half-tone pattern in (a). (d) Simulation output for half-tone pattern in (b). As in Figure 10, the simulation in (d) approximates (a) more closely than the simulation in (c) does.

Figure 11(a) shows the standard half-tone for this viewpoint, and Figure 11(b) shows the echo compensated half-tone for this viewpoint. In the bottom row of Figure 11 the outputs of simulations that model the echo interactions are shown. Figure 11(c) shows the simulation results for the standard half-tone and Figure 11(d) shows the simulation results for the half-tone designed with echo compensation. Figure 11(d) approximates Figure 11(a) much more closely than does Figure 11(c).

10. SIMULATION

We have built both an interactive viewer which shows the function we intend to print on the transparency, and a full ray traced simulation of the microgeometry. The interactive viewer calculates the amount of light which we expect the viewer to observe as a function of incoming light angle. This is the value obtained by evaluating the reflectance function, and represents ground truth. The viewer allows real-time interaction with the incident illumination, as well as switching between color, grayscale, black and white, and several versions of dithering.

The ray traced simulator allows us to verify that the microgeometry behaves as intended, and to investigate what effect printing on a transparency mask may have. The simulator uses a 10×10 grid of ink applied to either the dimple or transparency, approximately one dimple per rendered pixel, and 100 supersamples per pixel. We verified that sampling was sufficient by rendering a few comparisons with 2500 supersamples. Black/white dithering is not used, thus simulated ink is grayscale, in which case it ideally attenuates light without introducing unwanted scattering or specular contribution. The color example attenuates light separately in each




Fig. 12. Using an overlay mask is expected to introduce some errors. Using a ray traced simulation of our microgeometry we analyze this error, finding that the overlay has only a modest effect, when compared against ground truth. The residual RMS error of 3.2% is due to using a finite specular lobe width in the simulation. The error images in the bottom row have been contrast enhanced to make them visible when printed. Image used with the permission of Cultural Heritage Imaging.



Fig. 13. To validate our design using an overlay across multiple lighting directions we simulate the effect of moving a light across a 90 degree arc of illumination angles. Five frames from the resulting sequence are shown here. The average RMS error across all frames is 5.0%, and nowhere does the error exceed 6.3%, remaining relatively consistent across the full range of lighting angles.

of the RGB channels. The underlying surface is specular, but not a perfect mirror. The simulation accounts for multiple reflections between the transparency and surface, as well as shadows, although we do not expect these to be relevant, due to the specific design of our dimples.

In order to evaluate the effect of printing on a transparency, as opposed to directly on the dimples, we compare ray traced simulations. Figure 12 shows ray traced images. Absolute error images are also shown. Error is evaluated as the RMS difference from the reflectance image we intended to produce. As expected, printing on a transparency has a greater error than printing directly on dimples, however, the error is within the bounds predicted by our analysis, and the images themselves appear similar.

In order to verify that our design correctly reproduces variable illumination, we simulate a light moving from left to right over the surface. Figure 13 shows the intended ground-truth effect compared against the simulated results from our design, including printing on a transparency layer. The average RMS error over all frames is 5.0%, small enough that we consider this a validation that our design is functional.




Fig. 14. Our manufacturing process is validated by using a profilometer to measure a cross-section of our mirrored substrate. Dimples have a depth on the order of 50 microns, and deviations are below the resolution of our profilometer.




Fig. 15. (top) Our physical prototype consists of an array of mirrored spherical depressions covered with a mask printed on a transparency. (bottom left) Close-up view of the mirrored geometry inside the indicated box. (bottom right) Close-up view of the printed mask.

11. PROTOTYPE

We have created several prototype implementations of reflectance paper. We produced two sets of dimpled substrates, one with a hexagonal array of 2.3mm diameter dimples and another higher-resolution set with 1.0mm hexagonally arranged dimples. In both cases the substrates were produced by using a commercial array of microlenses as a mold master. These substrates were then metalized using silver or aluminum vacuum sputtered on the substrates to produce a mirror-like reflective surface. We validated that the shape of the dimples were spherical cross-sections subtending 70 degrees by using a profilometer, as shown in Figure 14. We were successful at registering the reflectance data with the substrate geometry and printing directly on the dimples with an Indigo commercial printer, but found the specularity of the inks far too high to effectively block specular highlights from the dimples themselves, as described in Section 5.

ACM Transactions on Graphics, Vol. 31, No. 3, Article 20, Publication date: May 2012.




Fig. 16. Photographs of our physical prototype under three different illumination conditions are compared against the intended black and white print from the same lighting angle.

To avoid artifacts due to ink specularity, we printed our binary reflectance pattern on a transparency using a standard laser printer, which was manually registered with the underlying reflective dimpled substrate. The reflectance function that defines the transparency mask was printed at a resolution of 32×32 samples per dimple, approximately the resolution limit of the printer. The lighting dome that captured the original real-world photographs samples far fewer lighting directions, so we used a polynomial approximation to generate a continuous reflectance function prior to resampling for printing. Figure 15 shows the physical prototype as well as macroviews of the dimples and printed transparency.

In order to evaluate whether our prototype faithfully represents the “printed” reflectance function, we illuminated it from several different angles, using a light bulb to approximate a point source.




Fig. 17. We have produced several prototypes. (left) Black and white masks simply threshold the measured reflectance function. (middle) Introducing spatial dithering produces much better grayscale reproduction, at the cost of resolution. (right) Dithering in the lighting domain is capable of producing grayscale at the full spatial resolution of our mask, but requires area lights for best effect.




Fig. 18. A single print with spatial dithering illuminated from a point source at many different lighting angles. Note that there is full 2D control over the hemisphere of incident lighting angles.




Fig. 19. A shell printed with spatial dithering, illuminated with several different lighting conditions.

Figure 16 shows a thresholded black/white print illuminated from three different directions. The figure also shows the ground-truth binary reflectance function from approximately the same lighting configuration. Note that the actual printed reflectance paper has appearance very close to the intended print. Note also that the illumination appears to come from different angles, as expected.

We have explored several methods for binarization of reflectance functions, spanning black and white thresholding, spatial dithering, and dithering in the lighting domain. We printed examples of each, and compared the appearance of the physical prototypes, as shown in Figure 17. As expected, spatial dithering obtains graylevels at the cost of spatial resolution and works best with low-frequency

spatial content. Lighting domain dithering obtains graylevels at the cost of lighting domain resolution, and works best with low-frequency lighting. We thus use area light sources to illuminate this prototype.

Figure 18 shows a print with spatial dithering illuminated from a point source at many different lighting angles. Note that there is full 2D control over the hemisphere of incident lighting angles. When viewing the physical prototype, the appearance varies smoothly as the light source moves. Representing different objects only requires printing a new transparency and aligning it to the dimpled substrate. Figure 19 shows a shell illuminated from several lighting conditions.

12. LIMITATIONS

The early physical prototypes shown in this work have many limitations: they lack color, the ink and transparency both have a strong specular component, and the mask and substrate are aligned by hand. We believe that design iterations and better engineering will remedy many of these deficiencies.

In our current implementation, incident illumination from the viewpoint will theoretically generate a specular return from the ink, negating the desired masking effect. Since the light source is rarely extremely close to the viewer, this has not been observed in practice.

Our design is optimized for a single viewing direction. This limitation prevents extreme viewpoint changes. However, in our design, viewpoint changes are approximately equivalent to rotation of the lighting environment, and we have noticed that the appearance is well preserved from nearby positions in practice.

13. CONCLUSIONS AND FUTURE WORK

We have developed a method for printing reflectance functions using a standard printer and specially prepared paper. We have produced several physical prototypes, and validated that they properly display reflectance images. We have also developed a method for producing custom halftones which takes into account our design using a transparent overlay.

Future work on developing less reflective black inks and reducing the size of the microgeometry would be useful. A full-color extension is also possible by having triads of reflective red, green, and blue specular divots on the media. Much like colored Christmas ornaments, these would still be mirrored, producing negligible diffuse return, but reflective in only a predetermined color band. The relative proportion of each primary would then be selectively controlled by spatially varying the transmissivity of the ink placed over this

ALGORITHM 1: Halftone with Echo Reduction.

```

1: for each dimple do
2:   get the constant Threshold value for this dimple
3:   for each Lighting Direction do {generate Standard halftone}
4:     if Reflection > Threshold then
5:       Halftone = Transparent
6:     else
7:       Halftone = Opaque
8:     end if
9:   end for{each Lighting Direction}
10:  for i = 1 to K do {Echo reduction iterations}
11:    for each Lighting Direction do
12:      if Halftone = Opaque then
13:        Compute quantized Echo Direction
14:        Retrieve HalftoneEcho value from the Echo Direction
15:        if HalftoneEcho = Transparent then
16:          if ErrorTransparent < ErrorOpaque then
17:            Halftone = Transparent
18:          end if
19:        end if
20:      end if
21:    end for{each Lighting Direction}
22:  end for{Echo reduction iterations}
23: end for{each dimple}

```

medium. Using conventional colored inks overlying a metalized mirrored substrate would be ineffective, as this would introduce unwanted diffuse contributions having the effect of lowering the contrast ratio of our prints.

A more speculative future direction is reversing the geometry so that the lighting angle is fixed and the images change in response to changing viewpoint. In this case we might encode animations or other effects, rather than illumination variation, and we envision such effects being potentially useful on billboards.

APPENDIX

For completeness we provide pseudocode for the spatial halftoning algorithm described in Section 9. This pseudocode includes the first stage labeled Standard halftone (described in Section 9.1), and a second stage called Echo Reduction Iteration (described in Section 9.2). The thresholds used for the dimples are periodic on the hexagonal grid of the reflection paper dimples, but for each individual dimple there is a constant threshold whose value is determined by the index of the dimple.

ACKNOWLEDGMENTS

We thank Mark Mudge and Carla Schroer of Cultural Heritage Imaging for capturing the multilight Tutela images used as source material to produce Figures 1(c), 7–9, 11, and 16–19. We thank Gary Gibson and Udi Chatow, both from HP Labs, for help fabricating early prototypes. We thank Susan Noe and Bob Smithson, from 3M, for their support with HP Labs / 3M collaborations. Also thanks to Bob Ulichney for discussions about halftoning.

REFERENCES

- ALEXA, M. AND MATUSIK, W. 2010. Reliefs as images. *ACM Trans. Graph.* 29, 4.
- ALLEBACH, J. 2001. DBS: Retrospective and future directions. In *Proceedings of the Color Imaging: Device Independent Color, Color Hardcopy, and Graphic Arts VI Conference (SPIE)*. Vol. 4300, 23–26.
- ARCE, G. 2008. *Modern Digital Halftoning 2nd Ed.* CRC Press.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Proceedings of the SIGGRAPH'00 Conference*. 145–156.
- FREETH, T., BITSAKIS, Y., MOUSSAS, X., SEIRADAKIS, J., TSELIKAS, A., MANGOU, H., ZAFEIROPOULOU, M., HADLAND, R., BATE, D., RAMSEY, A., ALLEN, M., CRAWLEY, A., HOCKLEY, P., MALZBENDER, T., GELB, D., AMBRISCO, W., AND EDMUNDS, M. 2006. Decoding the ancient Greek astronomical calculator known as the antikythera mechanism. *Nature* 444, 587–591.
- FUCHS, M., RASKAR, R., SEIDEL, H.-P., AND LENSCH, H. 2008. Towards passive 6D reflectance function displays. *ACM Trans. Graph.* 27, 3.
- HAWKINS, T., COHEN, J., AND DEBEVEC, P. 2001. A photometric approach to digitizing cultural artifacts. In *Proceedings of the Conference on Virtual Reality, Archeology, and Cultural Heritage (VAST'01)*.
- MALZBENDER, T., GELB, D., AND WOLTERS, H. 2001. Polynomial texture maps. In *Proceedings of the SIGGRAPH'01 Conference*. 519–528.
- MATUSIK, W., AJDIN, B., GU, J., LAWRENCE, J., LENSCH, H., PELLACINI, F., AND RUSINKIEWICZ, S. 2009. Printing spatially-varying reflectance. *ACM Trans. Graph.* 28, 5.

- MUDGE, M., MALZBENDER, T., SCHROER, C., AND LUM, M. 2006. New reflection transformation imaging methods for rock art and multiple-viewpoint display. In *Proceedings of the 7th International Symposium on Virtual Reality, Archeology, and Cultural Heritage (VAST'06)*.
- NAYAR, S., BELHUMEUR, P., AND BOULT, T. 2004. Lighting sensitive display. *ACM Trans. Graph.* 23, 4.
- NICODEMUS, F., RICHMOND, J., AND HSAI, J. 1977. Geometrical considerations and nomenclature for reflectance. U.S. Department of Commerce, National Bureau of Standards.
- PADFIELD, J., SAUNDERS, D., AND MALZBENDER, T. 2005. Polynomial texture mapping: A new tool for examining the surface of paintings. ICOM Committee for Conservation.
- PEERS, P., HAWKINS, T., AND DEBEVEC, P. 2006. A reflective light stage. Tech. rep. ICT-TR-04.2006, ICT (USC Institute for Creative Technologies).
- PEERS, P., TAMURA, N., MATUSIK, W., AND DEBEVEC, P. 2007. Post-Production facial performance relighting using reflectance transfer. *ACM Trans. Graph.* 26, 3.
- RAMAMOORTHI, R. AND HANRAHAN, P. 2001. An efficient representation for irradiance environment maps. In *Proceedings of the SIGGRAPH'01 Conference*.
- REGG, C., RUSINKIEWICZ, S., MATUSIK, W., AND GROSS, M. 2010. Computational highlight holography. *ACM Trans. Graph.* 29, 6.
- ULICHNEY, R. 1987. *Digital Halftoning*. MIT Press.
- WENGER, A., GARDNER, A., TCHOU, C., UNGER, J., HAWKINS, T., AND DEBEVEC, P. 2005. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Trans. Graph.* 24, 3.
- WEYRICH, T., DENG, J., BARNES, C., RUSINKIEWICZ, S., AND FINKELSTEIN, A. 2007. Digital bas-relief from 3D scenes. *ACM Trans. Graph.* 26, 5.
- WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM Trans. Graph.* 28, 5.

Received June 2011; revised November 2011; accepted November 2011