

Jiaping Wang  
Xin Tong  
John Snyder  
Yanyun Chen  
Baining Guo  
Heung-Yeung Shum

# Capturing and rendering geometry details for BTF-mapped surfaces

---

Published online: 31 August 2005  
© Springer-Verlag 2005

---

J. Wang\* (✉)  
Institute of Computing Technology,  
Chinese Academy of Science  
Graduated School of Chinese Academy of  
Science  
e\_boris2002@hotmail.com

X. Tong · Y. Chen · B. Guo · H.-Y. Shum  
Microsoft Research Asia  
{xtong, yachen, bainguo,  
hshum}@microsoft.com

J. Snyder  
Microsoft Research  
johnsny@microsoft.com

\* This work was done while J. Wang was  
visiting Microsoft Research Asia.

**Abstract** Bidirectional texture functions, or BTFs, accurately model reflectance variation at a fine (meso-) scale as a function of lighting and viewing direction. BTFs also capture view-dependent visibility variation, also called masking or parallax, but only within surface contours. Mesostructure detail is neglected at silhouettes, so BTF-mapped objects retain the coarse shape of the underlying model.

We augment BTF rendering to obtain approximate mesoscale silhouettes. Our new representation, the 4D mesostructure distance function (MDF), tabulates the displacement from a reference frame where a ray first intersects the mesoscale geometry beneath as a function of ray direction and ray position along that reference plane. Given an MDF, the mesostructure silhouette can be rendered with a per-pixel depth peel-

ing process on graphics hardware, while shading and local parallax are handled by the BTF. Our approach allows real-time rendering, handles complex, non-height-field mesostructure, requires that no additional geometry be sent to the rasterizer other than the mesh triangles, is more compact than textured visibility representations used previously, and, for the first time, can be easily measured from physical samples. We also adapt the algorithm to capture detailed shadows cast both by and onto BTF-mapped surfaces. We demonstrate the efficiency of our algorithm on a variety of BTF data, including real data acquired using our BTF-MDF measurement system.

**Keywords** Bidirectional texture functions · Reflectance and shading models · Rendering · Shadow algorithms · Texture mapping


---

## 1 Introduction

Real-world textures arise from both spatially varying surface reflectance and mesostructure, which represents fine-scale geometric details [14]. The 6D bidirectional texture function (BTF) captures these variations [6] and has been exploited to measure real materials and map their properties onto arbitrary geometric models [23, 25–27]. Rendering BTF-mapped surfaces yields rich visual effects such as fine-scale shading, shadows, interreflections, and occlusions, but it omits detailed silhouettes. This is because the

BTF only provides radiance measurements as a function of lighting and viewing direction, leaving mesostructure visibility implicit. Silhouettes are thus identical to that of the underlying mapped surface, as shown in Fig. 1a.

To solve this problem, we build on recent methods that explicitly tabulate local visibility. The view-dependent displacement map (VDM) [29] precomputes the mesostructure visibility as a function of texture position, view direction, and curvature. The result is stored as a 5D table and allows real-time rendering with graphics hardware. The generalized displacement map (GDM) [30] also stores a 5D table of displacements, but as



**Fig. 1.** **a** Existing BTF rendering techniques produce unrealistic silhouettes and shadow boundaries. **b** Our technique generates detailed silhouettes and shadow boundaries for BTF-mapped objects

a function of ray direction and 3D position in a volume. Both representations rely on synthetic visibility data, and it is not clear how either can be acquired from real materials. Moreover, the 5D table required strains the hardware memory system, making the approach less practical. Most fundamentally, neither method was combined with BTF rendering but instead used traditional normal-mapped shading and two-pass shadowing. Using a BTF efficiently provides local shadowing without a shadow buffer, as well as other more general effects such as diffuse interreflections and spatially varying reflectance, which are impractical with traditional shading techniques.

To add silhouette visibility to BTF rendering, we propose a new, image-based method that represents mesostructure as a 4D *mesostructure distance function* (MDF). It is used as a companion function for a given BTF and records, for each BTF viewing direction, the displacement along that direction of the mesostructure from the reference plane. Now we can determine whether the mesostructure is visible at the silhouette by rendering the front surface and the corresponding back surfaces and testing whether the depth between them exceeds the sum of the displacements indexed from the MDF at these two surfaces. Since object surfaces self-occlude, we use a multipass depth peeling approach to ensure that we can get the front surface and its corresponding back surfaces in silhouette rendering. Moreover, depth peeling ensures that we see through closer layers where the view ray misses the mesostructure to more distant ones where it may hit. The result provides visually important silhouette effects and can be extended to shadows, as shown in Fig. 1b.

The MDF can be seen as a simplified VDM without its curvature parameter. However, this simplification is not as straightforward as it may appear. The VDM

rendering method *relies* on a curvature parameter to obtain silhouette visibility effects. In other words, setting the curvature parameter to 0 and using the VDM rendering method [29] results in silhouettes from the underlying coarse model, just as with BTF rendering. It is only by using both our MDF representation and rendering method that we obtain mesoscale silhouette effects.

Our primary contribution is to add detailed silhouettes to BTF rendering. We show how our MDF representation can be obtained by measuring real samples at the same time the BTF is acquired. We also demonstrate realistic shadowing as well as silhouette effects, including, for the first time, shadowing *onto* surfaces mapped with mesoscale geometry. Our results demonstrate an efficient solution to simulate all significant visual effects caused by geometric details on BTF-mapped surfaces.

## 2 Related work

A number of techniques measure BTFs from real materials [5, 6, 8, 9, 23, 31]. Such BTFs can be mapped or synthesized onto geometric models (e.g., [16, 27]). BTFs have also been used for rendering complex, non-height-field mesostructures such as fabric weaves [23] and fur [8]. The polynomial texture map [17] may be regarded as a special form of BTF with a fixed view.

Despite the high-dimensional data required, BTF-mapped surfaces can be rendered in real time [15, 20, 23, 26, 28]. These methods decompose the 6D BTF data into a combination of lower-dimensional textures easily loaded and rendered on graphics hardware. Sloan et al. [25] apply the precomputed radiance transfer (PRT) technique

to BTF-mapped surfaces. All these approaches ignore the mesostructure silhouette.

For height-field mesostructures [1, 3], many techniques exist to achieve shading and parallax effects.

Horizon mapping generates mesoscale shadows on bump maps [19]. It can be accelerated using graphics hardware [24], and the representation can be captured from real samples [22]. Precomputed visibility [10] also renders mesoscale shadows and interreflections on bump-mapped surfaces. These methods neglect visibility effects entirely (e.g., parallax from bumps) both within surface contours and at silhouettes. They also lack the generality and realism of full BTF shading.

The parallax map [12] captures the local parallax of a height field in real time. Most recently, Policarpo et al. [21] have proposed a technique to render the self-occlusions, interpenetrations, and shadows of relief-mapped surfaces in real time. As is the case with BTF techniques, the underlying object’s silhouette is unchanged in these approaches. Shading is computed traditionally; i.e., is non-BTF.

Hirche et al. [11] propose an algorithm to render a displaced surface on a per-pixel basis. Each triangle of an underlying mesh is extruded to form a prism. The intersection of the ray and extruded triangle is computed at each pixel of the prism faces. (A similar approach is also employed by the GDM method and must also rasterize prism faces rather than just mesh triangles.) Although the technique can render the detailed silhouette of a displaced surface in real time, it is limited to height-field mesostructures and so precludes many real-world examples including the grass and weave textures in our results. We also demonstrate an acquisition system that measures both the BTF and MDF (visibility) simultaneously and a more realistic, BTF-based rendering method.

### 3 MDF measurement

For a given BTF, the companion MDF records the distance of the mesostructure surface from a reference surface along a BTF viewing direction. As illustrated in Fig. 2, the MDF is a 4D function  $d = D(x, y, \theta, \phi)$ , where  $(x, y)$  is the texture coordinate and the unit vector  $(\theta, \phi)$  is a sample viewing direction of the BTF expressed in spher-




Fig. 2. Definition of MDF

ical coordinates. For a given direction  $(\theta, \phi)$ , each reference surface point  $p = p(x, y)$  projects to a mesostructure point  $p'$  along the direction  $(\theta, \phi)$ , and  $D(x, y, \theta, \phi)$  is defined as the distance from  $p$  to  $p'$ . To facilitate the mesostructure visibility computation, we define both BTF and MDF on a reference surface that lies above the top of the mesostructure details.

Our MDF measurement device is a simple laser scan unit consisting of a laser source and a controller built with LEGO MindStorms Robotics Invention System 2.0. The laser scan unit is added to an existing BTF measurement system to measure a companion MDF for each captured BTF. Specifically, we generate a displacement image  $D_v(x, y) = D(x, y, v)$  for every viewing direction  $v$  of the BTF.

#### 3.1 System setup

MDF measurement is closely coupled with BTF measurement. As illustrated in Fig. 3a, our BTF measurement device mainly consists of a turntable and two concentric arcs, one for light sources and the other for cameras. The stationary camera arc shown on the left holds eight Dragonfly cameras, each providing  $1024 \times 768$  24-bit color images at 15 Hz. The cameras are calibrated as in [32]. The light arc on the right rotates horizontally, driven by a stepping motor. Eight halogen lamps are mounted on the light arc. The light sources are manually calibrated as in [2]. The turntable in the middle holds material samples to be measured and is driven by another stepping motor. All the lamps, cameras, and stepping motors are controlled by a PC.

BTF capture is a fully automatic process controlled by a PC. After rotating the light arc and turntable to their desired sampling positions, the PC sequentially turns the lamps on/off. With each lamp being turned on, every camera records an image. The captured images are rectified and then clipped to form the output BTF data.

Figure 3b shows the laser scan unit, which contains a laser strip generator and a controller. A miniature structured light diode laser serves as the light source, with a  $60^\circ$  line lens placed in front of the laser to project a laser strip




Fig. 3. **a** Our BTF–MDF measurement system. **b** Laser scan unit


on the turntable. The controller includes a robust gear-driven system that rotates the laser emitter vertically so that the laser strip scans across the turntable. We adjust the line lens direction to ensure that the laser scanning direction on the turntable is orthogonal to the laser strip. The laser unit is mounted on the light arc as shown in Fig. 3a.

In principle we could build a separate laser range scanner for measuring the MDF. By combining MDF with BTF acquisition, we use the same cameras for measuring both the BTF and MDF and avoid cross-registration problems. Our laser unit is easy to build and add to an existing BTF device.

### 3.2 Capturing process

The MDF measurement process involves three steps. First we project laser strips onto the material sample to be measured. As a laser strip scans across the sample, each camera records an image sequence. For a given camera with viewing direction  $v$ , the laser strip is often partly occluded from the camera by the mesostructure being measured. For this reason it is necessary to scan the material sample using laser strips projected from multiple source locations so that each mesostructure point visible from the camera is scanned at least once. Thus for each  $v$  we capture a number of laser scan image sequences, one for each laser source location. Then we apply space–time analysis to recover a displacement image for each captured image sequence. Finally, for every viewing direction  $v$  we merge all displacement images of  $v$  to generate  $D_v(x, y)$ .

*Scan image sequences.* For each viewing direction  $v$  we scan the material sample with the laser from 26 source locations as illustrated in Fig. 4. Since the laser emitter is far from the turntable, we regard all the laser planes (a laser plane is a plane containing the laser strip and laser source point) in one scan as being parallel. Interreflections on the mesostructure surface can produce noise in the laser scan images. We resolve this problem by using a quick shutter speed so that the interreflections are underexposed and can be filtered out by postprocessing. Alternatively




**Fig. 4.** The laser source locations are marked by orange dots. Currently we simply remount the laser unit on the light arc manually (5 times per capturing session). Alternatively, we can mount several identical units along the light arc or use a mobile unit

we can spray the material sample with matte paint before MDF measurement. Before recording laser scan images for each laser source location, we first compute the angle between the laser plane and the reference plane (the turntable plane). To do this we project the laser strip onto two planes with different heights and calculate the offset of the laser strip in the two planes from two recorded laser scan images. From this offset we derive the angle between the laser plane and the reference plane. Once the angle is found, laser scan images of the material sample are recorded automatically. Specifically, we rotate the light arc and turntable to the desired sampling position. As the laser strip scans across the material sample, an image sequence is recorded by each BTF device camera at 7.5 Hz.

An MDF capturing session produces a set of laser scan image sequences. The image sequence for laser source direction  $l$  and viewing direction  $v$  is  $\{I_{v,l}(x, y, t_i)\}$ .

*Initial displacement images.* From each laser scan image sequence  $\{I_{v,l}(x, y, t_i)\}$  we compute the displacement image  $D_{v,l}(x, y)$  using space–time analysis [4]. A similar method has been used for shadow scanning [2]. Before the space–time analysis, we compute the laser plane for every image  $I_{v,l}(x, y, t_i)$  in two steps. The first step is to detect the laser line in the image. Using the bounding box of the material sample, we divide  $I_{v,l}(x, y, t_i)$  into two parts: the reference region showing only the turntable and the sample region where the material sample is imaged. The bounding box is manually specified on the reference plane before capturing. The laser strip should be a line in the reference region since it is flat. As illustrated in Fig. 5a, to find this line, we search the reference region for candidate pixels on the line along the laser scan direction. Since the candidate pixels are by far the brightest pixels in the image, it is easy to reject outliers using a pre-



**Fig. 5a,b.** MDF capturing process. **a** Laser line detection. *Bottom green curve:* Intensity of pixels on vertical line. *Peak point:* candidate pixel on laser line. **b** Space–time analysis. *Bottom red curve:* intensity variation of a pixel in whole scanning sequence. *Peak point:* time at which this pixel images a mesostructure point

scribed threshold. For noise reduction, we Gaussian-filter the image before the search. After finding all candidate pixels, we fit a line to them. In the second step, we calculate the laser plane for the image  $I_{v,l}(x, y, t_i)$  by projecting the laser line from  $I_{v,l}(x, y, t_i)$  to the 3D reference plane.

Once the laser planes of all the images are found, the space-time analysis proceeds as follows. As illustrated in Fig. 5b, for each pixel  $(x_s, y_s)$  in the sample region, we search the temporal sequence  $\{I_{v,l}(x_s, y_s, t_i)\}$  to find the pixel  $(x_s, y_s, t_{i_0})$  that has the brightest  $I_{v,l}(x_s, y_s, t_{i_0})$ . If  $I_{v,l}(x_s, y_s, t_{i_0})$  is greater than the prescribed threshold, then the pixel  $(x_s, y_s, t_{i_0})$  images a mesostructure point  $p$ . We find  $p$  by projecting the pixel  $(x_s, y_s, t_{i_0})$  onto the laser plane of the image  $I_{v,l}(x, y, t_{i_0})$ . We then compute the displacement of  $p$  from the reference plane and record the result in  $D_{v,l}(x_s, y_s)$ . If the  $I_{v,l}(x_s, y_s, t_{i_0})$  is less than the prescribed threshold, the pixel  $(x_s, y_s, t_{i_0})$  corresponds to an occluded mesostructure point. In this case, we set  $D_{v,l}(x_s, y_s) = -1$ .

*Merging displacement images.* For every viewing direction  $v$ , we merge all displacement images  $D_{v,l}(x, y)$  to generate  $D_v(x, y)$ :

$$D_v(x, y) = \frac{1}{\sum w_i} \sum w_i D_{v,l_i}(x, y) \quad (D_{v,l_i}(x, y) \neq -1),$$

where the  $w_i$  is the weight of the scan result with laser direction  $l_i$ . Based on the error analysis in [2], we compute the weight for each scan  $D_{v,l}(x, y)$  as the volume of the tetrahedron formed by the laser source position, the camera location, and two points on the laser line passing through the turntable center.

## 4 Rendering BTF-mapped surfaces with MDF

After the BTF is mapped or synthesized onto arbitrary surfaces [27], the associated MDF texture can be mapped onto the surface with the same texture coordinates. The MDF data mapped onto surface are then used for rendering the mesostructure silhouette and mesostructure shadow silhouette, and also for enhancing shadow casting for a BTF-mapped surface. All these computations can be performed in the pixel shader and accelerated with graphics hardware.

### 4.1 Visibility computation

With the MDF we can render mesostructure silhouettes as well as mesostructure shadow silhouettes. These two types of silhouettes are intimately related in that both are the results of visibility changes caused by mesostructures. For silhouettes the visibility is that from the viewpoint; for shadow silhouettes the visibility is that from the light




Fig. 6. Determining mesostructure silhouettes using MDF

source. The visibility computation algorithm discussed here is limited to closed objects.

Figure 6 illustrates the algorithm for determining mesostructure silhouettes. For simplicity we consider the silhouette determination problem in 2D (we can reduce the original 3D problem to 2D by performing silhouette determination for each scanline of the rendered image). Consider the viewing ray  $r_0$  that enters the object surface at point  $p_f$  on a front face and leaves at point  $p_b$  on a back face. With MDF value  $D(x_f, y_f, v_f)$  on  $p_f$ , the offset from  $p_f$  to the mesostructure surface along  $r_0$  is computed as  $d_f = D(x_f, y_f, v_f)s_f$ , where  $s_f$  is the texture scale at  $p_f$ .  $(x_f, y_f)$  is the texture coordinate at  $p_f$ , and  $v_f$  is the viewing direction of  $r_0$  computed in the local coordinate frame at  $p_f$ . Similarly, the distance from point  $p_b$  to the mesostructure surface is  $d_b = D(x_b, y_b, v_b)s_b$ . As illustrated in Fig. 6, if  $d_f + d_b \leq \|p_f - p_b\|$ , the ray will intersect the mapped mesostructure at some point. Otherwise, the ray will pass through the mesostructure without intersection. If the ray intersects the object multiple times, we do the same computation for each pair of  $p_f$  and  $p_b$  to determine the visibility of the mesostructure mapped on the surface.

To handle the mesostructure silhouette with graphics hardware, we use the depth peeling technique [7, 18] to find  $p_f$  and  $p_b$  for each interval inside the objects. As illustrated in Fig. 7, before peeling we initialize the depth buffer and stencil buffer. We also need two extra depth buffers for peeling,  $Z_{peel\_now}$  and  $Z_{peel\_next}$ , which store the minimum depth for the current and next peeling pass, respectively. In our algorithm, we execute the depth peeling from front to back. In each peeling pass, we first render the back faces. The  $Z_{peel\_now}$  and the depth buffer  $Z$  work together to obtain the pixels  $p_b$  of the back faces in the current peel layer, whose depths are stored in  $Z_{peel\_next}$ . We then offset their depths with the MDF value and store the result  $(z_{p_b} - d_b)$  in the depth buffer. After that, we disable depth buffer writing and render all front faces. For pixels  $p_f$  that pass the peeling depth test, we compare their offset depths  $z_{p_f} + d_f$  with  $z_{p_b} - d_b$  stored in the depth buffer. If  $z_{p_f} + d_f \leq z_{p_b} - d_b$ , the ray intersects with the mesostructure detail between  $p_f$  and  $p_b$ . Otherwise, the ray passes through this peeling layer without intersection. For pixels whose rays intersect with mesostructure details, we shade  $p_f$  with BTF data accord-

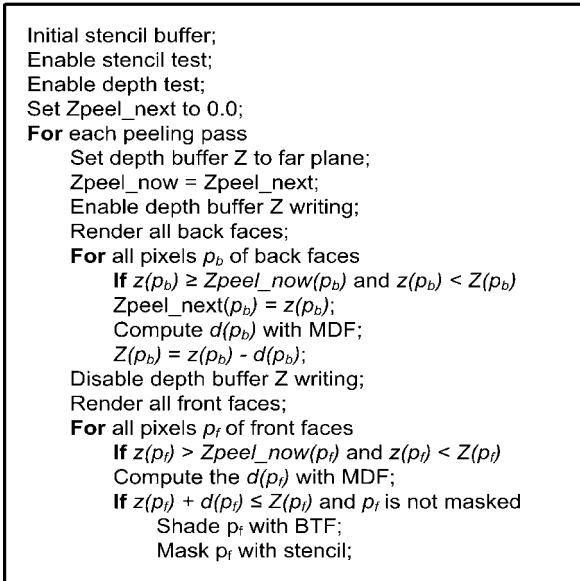


Fig. 7. Pseudocode of silhouette rendering

ing to the viewing and lighting conditions. After that, we mask the shaded pixel using the stencil buffer so that it will not be corrupted with the surfaces in the following peeling layers. Then we swap the peeling depth buffer and do the next peeling pass. In theory, the peeling operation should be continued until no pixels on the screen are updated. In our current implementation, the number of the peeling passes for each scene is specified by the user.

Determining mesostructure shadow silhouettes is done in the same way except we render the scene for the lighting direction. Instead of shading the surface pixels  $p_f$  with BTF, we render their depths to the shadow map buffer.

#### 4.2 Shadow casting on BTF-mapped surfaces

To render the BTF-mapped surface realistically, we also need to consider the shadows cast on the BTF-mapped surface. A problem with existing rendering techniques is that the mesostructure has no effect near shadow boundaries. Figure 8a shows an example of shadow testing with existing BTF rendering algorithms. The two reference surface points  $p_1$  and  $p_2$  have the same viewing and lighting directions  $v$  and  $l$ . Without taking the mesostructure into




Fig. 8. **a** Without MDF, both  $p_1$  and  $p_2$  are mistakenly considered lit as indicated by red arrows. **b** With MDF,  $p_1$  is lit whereas  $p_2$  is in shadow




Fig. 9. Shadow computation geometry

consideration, we would mistakenly regard both  $p_1$  and  $p_2$  as lit.

To generate realistic shadows, we need to combine the effects of mesoscale geometry with shadow maps cast by macroscale objects, which can be easily done with MDF. For simplicity, we assume that we have a point light source casting hard shadows.

As illustrated in Fig. 9, for each surface point  $p$  with texture coordinate  $(x, y)$ , we find the intersection point  $p'$  of the viewing ray and the mesostructure surface  $p' = p + s D(x, y, v_p) v_p$ , where  $s$  is the texture scale at  $p$  and  $v_p$  is the viewing direction defined in the local coordinate frame. Using the shadow map we test whether  $p'$  is in the macroscale shadow. Since the self-shadowing effect of mesostructure has been included in the BTF, we generate the shadow map for a BTF-mapped surface blocker such that the depth of each point is measured from the “base” of the mesostructure. As shown in Fig. 9, at point  $p''$  the shadow map depth is  $d(p'') + \frac{d_{max}}{\cos \theta_l}$ , where  $d(p'')$  is the depth from the light source calculated using the macroscale geometry and  $d_{max}$  is the maximum depth of the mesostructure.

This algorithm can be easily integrated with the visibility computation algorithm described in the pseudocode. Specifically, to generate the shadow map, we write the  $z(p_f) + \frac{d_{max}}{\cos \theta_l}$  into the shadow buffer, where  $\theta_l$  is the light direction computed in the local coordinate frame at  $p_f$ . To render the final result, we compute the mesostructure point with depth as  $z(p_f) + d(p_f)$  for  $p_f$  and test it against the shadow map. Depending on whether a pixel is in shadow, we shade point  $p$  with either an ambient term  $L_a(x, y, v_p)$  or the BTF value  $T(x, y, v_p, l_p)$ , where  $L_a(x, y, v_p)$  is obtained by integrating the BTF value for all lighting directions  $l = (\theta_l, \phi_l)$  on the hemisphere  $\Omega$ :

$$L_a(p, v_p) = \int_{\Omega} T(x, y, v_p, l) d\omega_l.$$

## 5 Experimental results

*Implementation details.* We have implemented our new rendering algorithm on a PC with a PIV 2.8-GHz CPU and Geforce 6800GT graphics card. In our implementation,

we first render the shadow map for the lighting direction. Then we render the final result with the BTF. To fetch the BTF and MDF data in rendering, we first compute the texture scale  $s$ , viewing direction  $v$ , lighting direction  $l$ , texture coordinate  $(x, y)$ , and texture coordinate frame for each vertex. After rasterization, these vertex attributes are interpolated to pixels. To render the BTF with graphics hardware, we decompose the 6D BTF into several lower-dimensional maps using singular value decomposition (SVD) [13]. Specifically, we reorganized the 6D BTF data into a 2D matrix  $A$  where the rows are indexed by  $x, y, \theta_l$ , and  $\phi_l$  and the columns are indexed by viewing angle  $\theta_v, \phi_v$ . Applying SVD to  $A$  produces a set of 4D eigenmaps  $E_i(x, y, \theta_l, \phi_l)$  and corresponding 2D weight maps  $W_i(\theta_v, \phi_v)$ . By keeping a small number of eigenmaps (four for all BTF data used in this paper), the BTF data are compressed and then can be reconstructed as

$$BTF(x, y, \theta_v, \phi_v, \theta_l, \phi_l) = \sum_i W_i(\theta_v, \phi_v) E_i(x, y, \theta_l, \phi_l).$$

In our current implementation, the 4D eigenmap is reorganized into a 3D texture for rendering, by the means of packing lighting directions to the third dimension.

To render the 4D MDF data, we directly reorganize the data into 3D texture by packing viewing directions to the third dimension. The trilinear interpolation is used for texture filtering.

*Experimental results.* Table 1 lists all MDF and compressed BTF data are used in this paper, where the peanut and the weave data are captured from real materials with our BTF-MDF capturing device. The other data are rendered from synthetic mesostructures by ray tracing. For BTFs, we store each captured 24-bit RGB image. MDF depth data are quantized to 8 bits per pixel and stored in MDF for rendering. Note that the 4D MDF data are small enough, which can be easily loaded in graphics hardware without compression.

Figure 12 demonstrates a scene rendered with BTFs and MDFs measured from real materials. Both the coarse weave BTF/MDF data on the tori and the peanut BTF/MDF data on the floor are captured from real materi-

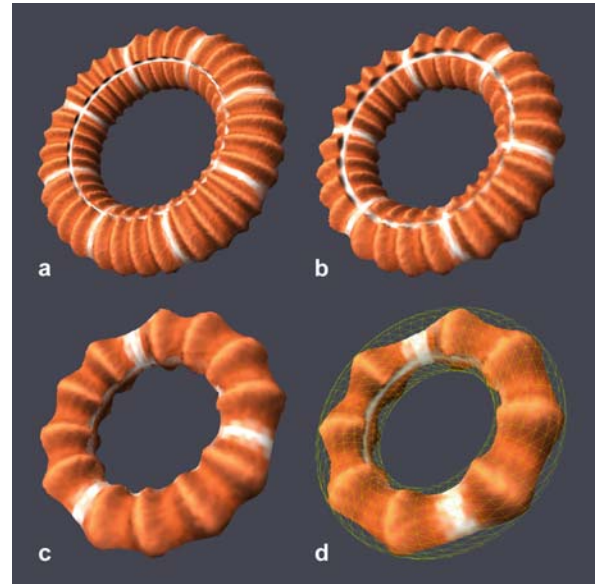


**Fig. 10a–c.** Comparison of results rendered **a** by BTF only, **b** by BTF and MDF, and **c** ground truth

als using our device. Two peeling passes are used to render mesostructure silhouettes and mesostructure shadow silhouettes. Note the mesostructure silhouette of the tori, the detailed shadow boundary on the floor, and other visual effects rendered by both the BTF and the MDF. Also, the detailed shadow boundaries cast on the floor are consistently rendered with the implicit geometry rendered with the BTF on the floor.

Figure 13 demonstrates a scene decorated with synthetic grass and brick BTF/MDFs, where three peeling passes are used in rendering. The mesostructure silhouette on the torus knot is convincingly rendered. Also observe the mesostructure shadow silhouette on the floor. Another rendering result of BTF/MDF data is exhibited in Fig. 14, where two peeling passes are used to generate the mesostructure silhouette effects.


Table 2 lists the rendering performance of our algorithm for different scenes shown in the paper, where the output image resolution is  $800 \times 600$ .




**Fig. 11a–d.** Rendering results of the same torus mapped with brick BTF/MDF data in different texture scales **a** with  $8 \times 3$  tiles, **b** with  $6 \times 2$  tiles, **c** with  $3 \times 1$  tiles, and **d** with  $2 \times 1$  tiles. The base triangle mesh is shown as the *yellow wireframes* in **d**

**Table 1.** Sampling resolution and size of BTF/MDF data


Sample	Image res.	Light & view resolution	BTF data (MB)	MDF data (MB)
Blue pipe	$64 \times 64$	$12 \times 5 \times 12 \times 5$	1.3	0.9
Brick	$64 \times 64$	$12 \times 5 \times 12 \times 5$	1.3	0.9
Peanut	$128 \times 128$	$12 \times 8 \times 12 \times 8$	7.5	6.3
Weave	$128 \times 128$	$12 \times 8 \times 12 \times 8$	7.5	6.3
Grass	$64 \times 64$	$12 \times 8 \times 12 \times 8$	5.3	3.9
Block	$64 \times 64$	$12 \times 5 \times 12 \times 5$	1.3	0.9



**Fig. 12.** Rendering result of BTF-mapped surfaces with real-world BTF/MDF samples



**Fig. 13.** BTF-mapped surface with synthetic BTF and MDF data



**Fig. 14a–c.** Three examples of BTF-mapped surfaces with shadow and silhouette effects **a** with weave BTF/MDF, **b** with block BTF/MDF, and **c** with brick BTF/MDF

*Limitations.* One issue is that the BTF is measured under spatially uniform lighting. Introducing macroscale blockers causes the illumination to vary spatially. This has a subtle and complicated influence on mesostructure interreflections that our method does not account for. Qual-

itatively, our method generates overly sharp shadows that would be softened if interreflections were simulated accurately. Note that this is a problem for any method based on BTFs. Another issue is that the BTF and MDF data are measured from a plane and then mapped onto

**Table 2.** FPS of BTF/MDF rendering for different scenes

Scene	Triangle number	BTF only	With Object silhouette	With object & shadow silhouette
Fig. 1	11,952	186	75	47
Fig. 12	3,200	110	42	27
Fig. 14	11,952	115	30	18
Fig. 14a	4,276	186	75	49
Fig. 14b	4,276	110	40	25
Fig. 14c	4,276	115	85	44

curved surfaces. Our peeling-based, depth-interval test is only an approximation of the true deformation of the mesostructure geometry onto a curved surface. As shown in Fig. 11, the method we propose yields consistent results at different surface curvatures. Even in the extreme situation, the bottom-right torus mapped in the tile-2 scale, the mesostructure and silhouette boundary appear matched with each other and consistent as the curvature varies. Figure 10 compares the rendering results of a cylinder mapped with mesostructure rendered with BTF, BTF/MDF, and displaced geometry. The displaced geometry mapped onto the cylinder is rendered by ray tracing.

Note that our approach provides convincing visual effects that are consistent with the appearance rendered with the surface BTF. In particular, the shadow generated by our algorithm is very similar to that generated by detailed geometry. The silhouette generated by our method also provides a good approximation of the ground truth.

## 6 Conclusion


We have presented an algorithm that augments BTF rendering to include previously neglected, detailed silhouette effects using the MDF. The MDF can be measured from real materials without undue additional work in BTF acquisition. Experimental results demonstrate that our new algorithm can render realistic mesostructure silhouettes and shadows, including shadows on BTF-mapped surfaces, in real time, and at low storage cost. For future work we are interested in automatically obtaining geometry information solely from the BTF.

**Acknowledgement** We are grateful to the anonymous reviewers, whose comments were very valuable in refining our paper. Many thanks to Yasuyuki Matsushita for helping us set up capture devices.

## References

- Blinn, J.F.: Simulation of wrinkled surfaces. In: Computer Graphics (SIGGRAPH '78 Proceedings), **12**(3), 286–292 (1978)
- Bouguet, J.Y., Perona, P.: 3d photography on your desk. In: IEEE International Conference on Computer Vision (ICCV98), pp. 43–50 (1998)
- Cook, R.L.: Shade trees. In: Computer Graphics (SIGGRAPH '84 Proceedings), **18**(3), 223–231 (1984)
- Curless, B., Levoy, M.: Better optical triangulation through spacetime analysis. In: IEEE International Conference on Computer Vision (ICCV95), pp. 987–994 (1995)
- Dana, K.J.: Brdf/btf measurement device. In: Proceedings of the 8th IEEE International Conference on Computer Vision (ICCV), **2**, 460–466 (2001)
- Dana, K.J., van Ginneken, B., Nayar, S.K., Koenderink, J.J.: Reflectance and texture of real-world surfaces. *ACM Trans. Graph.* **18**(1), 1–34 (1999)
- Diefenbach, P.J.: Pipeline rendering: interaction and realism through hardware-based multi-pass rendering. Ph.D. thesis, Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104 US (1996)
- Furukawa, R., Kawasaki, H., Ikeuchi, K., Sakauchi, M.: Appearance based object modeling using texture database: acquisition, compression and rendering. In: Eurographics Workshop on Rendering, pp. 257–266 (2002)
- Han, J.Y., Perlin, K.: Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Trans. Graph.* **22**(3), 741–748 (2003)
- Heidrich, W., Daubert, K., Kautz, J., Seidel, H.P.: Illuminating micro geometry based on precomputed visibility. In: Computer Graphics (Proceedings of SIGGRAPH 2000), pp. 455–464 (2000)
- Hirche, J., Ehlert, A., Guthe, S., Doggett, M.: Hardware accelerated per-pixel displacement mapping. In: GI '04: Proceedings of the 2004 Conference on Graphics Interface, pp. 153–158. Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, ON, Canada (2004)
- Kaneko, T., Takahehi, T., Inami, M., Kawakami, N., Yanagida, Y., Maeda, T., Tachi, S.: Detailed shape representation with parallax mapping. In: ICAT2001(11th International Conference on Artificial Reality and Tele-Existence), pp. 205–208 (2001)
- Kautz, J., McCool, M.D.: Interactive rendering with arbitrary BRDFs using separable approximations. In: Lischinski, D., Larson, G.W. (eds.) *Rendering Techniques '99*, Eurographics, pp. 247–260. Springer, Berlin Heidelberg New York (1999)
- Koenderink, J.J., Doorn, A.J.V.: Illuminance texture due to surface mesostructure. *J. Opt. Soc. Am.* **13**(3), 452–463 (1996)
- Liu, X., Hu, Y., Zhang, J., Tong, X., Guo, B., Shum, H.Y.: Synthesis and rendering of bidirectional texture functions on arbitrary surfaces. *IEEE Trans. Visual. Comput. Graph.* **10**(3), 278–289 (2004)
- Liu, X., Yu, Y., Shum, H.Y.: Synthesizing bidirectional texture functions for real-world surfaces. In: Proceedings of SIGGRAPH, pp. 97–106 (2001)
- Malzbender, T., Gelb, D., Wolters, H.: Polynomial texture maps. In: Proceedings of SIGGRAPH, pp. 519–528 (2001)
- Mammen, A.: Transparency and antialiasing algorithms implemented with the virtual pixel maps technique. *IEEE Comput. Graph. Appl.* **9**(4), 43–55 (1989)
- Max, N.: Horizon mapping: shadows for bumped mapped surfaces. *Visual Comput.* **4**(2), 109–117 (1988)
- Mueller, G., Meseth, J., Klein, R.: Compression and real-time rendering of measured BTFs using local PCA. In: Proceedings of Vision, Modeling and Visualisation (2003)
- Policarpo, F., Oliveira, M.M., Comba, J.L.D.: Real-time relief mapping on arbitrary polygonal surfaces. In: ACM SIGGRAPH 2005 Symposium on

- Interactive 3D Graphics and Games (I3D 2005) (2005)
22. Rushmeier, H., Balmelli, L., Bernardini, F.: Horizon map capture. In: Proceedings of Eurographics (2001)
  23. Sattler, M., Sarlette, R., Klein, R.: Efficient and realistic visualization of cloth. In: Eurographics Symposium on Rendering, pp. 167–178 (2003)
  24. Sloan, P.P., Cohen, M.F.: Interactive horizon mapping. In: Eurographics Workshop on Rendering, pp. 281–286 (2000)
  25. Sloan, P.P., Liu, X., Shum, H.Y., Snyder, J.: Bi-scale radiance transfer. *ACM Trans. Graph.* **22**(3), 370–375 (2003)
  26. Suykens, F., vom Berge, K., Lagae, A., Dutré, P.: Interactive rendering with bidirectional texture functions. In: Proceedings of Eurographics (2003)
  27. Tong, X., Zhang, J., Liu, L., Wang, X., Guo, B., Shum, H.Y.: Synthesis of bidirectional texture functions on arbitrary surfaces. *ACM Trans. Graph.* **21**(3), 665–672 (2002)
  28. Vasilescu, M.A.O., Terzopoulos, D.: TensorTextures: multilinear image-based rendering. *ACM Trans. Graph.* **23**(3), 336–342 (2004)
  29. Wang, L., Wang, X., Tong, X., Lin, S., Hu, S., Guo, B., Shum, H.Y.: View-dependent displacement mapping. *ACM Trans. Graph.* **22**(3), 334–339 (2003)
  30. Wang, X., Tong, X., Lin, S., Hu, S., Guo, B., Shum, H.Y.: Generalized displacement maps. In: Eurographics Symposium on Rendering (2004)
  31. Yamauchi, Y., Sekine, M., Yanagawa, S.: Bidirectional texture mapping for realistic cloth rendering. In: SIGGRAPH 2003 Sketches and Applications (2003)
  32. Zhang, Z.: Flexible camera calibration by viewing a plane from unknown orientations. In: IEEE International Conference on Computer Vision (ICCV99), pp. 666–673 (1999)



JIAPING WANG is a first-year Ph.D. student at the Institute of Computing Technology, Chinese Academy of Sciences in the graduate school of the Chinese Academy of Sciences. Before that, he received his B.S. from Ningbo University in 2002. He has been on an internship at Microsoft Research Asia (MSRA) since 2003. His research interests include appearance modeling and capturing, hardware accelerated rendering, high-dimensional texture analysis, and compression. He is also interested in early vision, perception, and machine learning.

DR. XIN TONG is a researcher/project lead of the Internet Graphics Group, Microsoft Research Asia. After receiving his Ph.D. in computer graphics from Tsinghua University, Beijing, in July 1999, he joined Microsoft Research China as an associate researcher. Prior to that, he received his B.S. and Master's degree in computer science from Zhejiang University in 1993 and 1996, respectively. His research interests include appearance modeling and rendering, image-based rendering, tex-

ture synthesis, and natural phenomena simulation.

DR. JOHN SNYDER is a senior researcher at Microsoft Research. He received his B.S. from Clarkson University, Postdam, NY, in 1984 and his Ph.D. from the California Institute of Technology, Pasadena, CA, in 1991 and has been at Microsoft Research since 1994. His research interests include geometry representation and processing and real-time algorithms for global illumination.

DR. YANYUN CHEN is an associate researcher at Microsoft Research Asia. His current research interests include photorealistic and nonphotorealistic rendering and image-based rendering. Dr. Chen received his Ph.D. from the Institute of Software, Chinese Academy of Sciences, in 2000 and his M.S. and B.S. from the Chinese University of Mining and Technology in 1996 and 1993, respectively.

DR. BAINING GUO is the research manager of the Internet Graphics Group at Microsoft Re-

search Asia. Before joining Microsoft, Baining was a senior staff researcher at Microcomputer Research Labs at Intel Corp. in Santa Clara, CA, where he worked on graphics architecture. Baining received his Ph.D. and M.S. from Cornell University and his B.S. from Beijing University. Baining is an associate editor of IEEE Transactions on Visualization and Computer Graphics. He holds over 30 granted and pending US patents.

DR. HEUNG-YEUNG SHUM is currently the managing director of Microsoft Research Asia. Dr. Shum joined Microsoft Research in 1996 after receiving his Ph.D. in robotics from the School of Computer Science at Carnegie Mellon University. He has authored and coauthored over 100 papers in computer vision, computer graphics, and robotics and received more than 20 US patents. He is on the editorial boards for IEEE Transactions on Circuit System Video Technology (CSVT), IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), the International Journal of Computer Vision, and Graphical Models.