

Poisson-Based Continuous Surface Generation for Goal-Based Caustics

Yonghao Yue

Columbia University

Kei Iwasaki

Wakayama University and UEI Research

Bing-Yu Chen

National Taiwan University and UEI Research

Yoshinori Dobashi

Hokkaido University and UEI Research

and

Tomoyuki Nishita

UEI Research and Hiroshima Shudo University

We present a technique for computing the shape of a transparent object that can generate user-defined caustic patterns. The surface of the object generated using our method is smooth. Thanks to this property, the resulting caustic pattern is smooth, natural, and highly detailed compared to the results obtained using previous methods. Our method consists of two processes. First, we use a differential geometry approach to compute a smooth mapping between the distributions of the incident light and the light reaching the screen. Second, we utilize this mapping to compute the surface of the object. We solve Poisson's equation to compute both the mapping and the surface of the object.

Categories and Subject Descriptors: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling; I.3.8 [Computer Graphics]: Applications

General Terms: Algorithms, Design

Additional Key Words and Phrases: Caustics, surface generation, continuous, Poisson equation

ACM Reference Format:

Yonghao Yue, Kei Iwasaki, Bing-Yu Chen, Yoshinori Dobashi, Tomoyuki Nishita. 2014. Poisson-based continuous surface generation for goal-based caustics. *ACM Trans. Graph.* 33, 3, Article 31 (May 2014), 7 pages. DOI: <http://dx.doi.org/10.1145/2580946>

1. INTRODUCTION

In computer graphics, many methods have been proposed to compute the shape or structure of an object to obtain a desirable illumination effect, such as reflectance fields [Fuchs et al. 2008], surface reflections [Weyrich et al. 2009], shadows [Mitra and Pauly 2009; Baran et al. 2012; Bermano et al. 2012], subsurface scattering [Dong et al. 2010; Hařan et al. 2010], shading [Alexa and Matusik 2010], holography [Regg et al. 2010], and refractive steganography [Papas et al. 2012].

Among these methods, caustics provide us with some of the most impressive illumination effects. Our aim is to compute the shape of the transparent object that produces the caustic pattern specified by the user (Figure 1). The setting is shown in Figure 2. We assume that the light incident on the object (the corresponding surface is

called the *incident surface*) is nearly parallel, and more or less normal to the incident surface, so that the light is in essence not refracted at the incident surface. On the opposite side of the object, there is a surface called the *refractive surface*, where the light can be refracted. When we place this object between the light source and a planar screen, we want the refracted light reaching the screen to form the desired caustic pattern. The problem dealt with in this article is the computation of the shape of this refractive surface. In Figure 1, we show a real example of an object constructed in this way and the resulting caustic pattern.

Although previous methods [Finckh et al. 2010; Papas et al. 2011; Yue et al. 2012] have been proposed for realizing the required caustics, they are somewhat restricted. Finckh et al.'s method [2010] can generate only simple caustics. Papas et al.'s method [2011] and Yue et al.'s method [2012] can generate more complex caustics; however, these methods generate non smooth surfaces, affecting the quality of the caustics. In this article, we propose a new method with which both continuous surfaces and caustics can be generated, with the aim of producing high-quality caustic patterns.

For the production of high-quality caustic patterns, we want to consider not only the case where the object is placed in focus, but also the case where the object is placed out of focus. To produce high-quality caustics in both cases, five criteria need to be met: (i) the need to produce complex patterns such as general images, (ii) the need to produce high-resolution images, and (iii) the need to have a continuous dynamic range (with high quality even in regions of low intensity). In addition, (iv) the focus range must be wide, and (v) the caustic pattern produced must be continuous even when the object is not in focus. The criteria (iv) and (v) are needed because we do not want the quality of the generated caustic pattern to be sensitive to the locations and orientations of the light source, object, and screen (i.e., caustics can be produced even when the placement is not perfect).

Our aim is to generate object surfaces that match all of the criteria listed before, whereas previous methods have at least one criterion which cannot be satisfied (see Section 2).




Fig. 1. The fabricated object (right) and its caustics (left) on the screen. Please see Figure 7 for the shape of the refractive surface of the object.




Fig. 2. An illustration of the proposed display system.

2. PREVIOUS APPROACHES

The approaches taken in the previous methods [Finckh et al. 2010; Papas et al. 2011; Yue et al. 2012] can be classified into the following two types. The first approach, taken by Finckh et al. [2010], randomly perturbs the refractive surface to find the solution. The second approach, taken by Papas et al. [2011] and Yue et al. [2012], first determines the relationship between the light on the incident surface and that on the screen, and then computes the shape of the refractive surface. In the first approach, the solution space is complex, and there can be many solutions that are optimal locally. Papas et al. [2011] pointed out that the result generated using this approach can converge to a solution that does not span the full contrast range of a complex target image (criterion (i) in Section 1).

In the second approach, the key element is the way in which the relationship between the light on the incident surface and that on the screen is found. With previous methods only a discrete relationship could be found. In those previous methods, the incident surface (or the refractive surface) and the screen are subdivided into sets of small regions, and the relationships between these small regions are established. Papas et al.’s method [2011] subdivides the incident surface into a grid and the caustics into a set of Gaussian kernels, and uses simulated annealing to find the relationship. In Yue et al.’s method [2012], the object can be rearranged using a set of discrete sticks.

In these methods [Papas et al. 2011; Yue et al. 2012], the light incident on neighboring regions generally arrives at distant regions on the screen. Due to this discontinuity, the following problems oc-




Fig. 3. An illustration of the computation of the relationship.

cur. First, each small region usually needs to be of a certain size, which restricts the resolution of the caustic pattern (criterion (ii) in Section 1). Second, the dynamic range of the caustics is somewhat restricted (criterion (iii) in Section 1). This is because the intensity of the caustics is basically determined by the number of overlapping regions on the screen, and since each small region on the incident surface is of a similar size, the intensity of the caustics becomes discrete. Furthermore, the range over which the caustics are in focus is short (criterion (iv) in Section 1), and if the object is placed out of focus, the resulting caustic pattern would be a set of discrete spots (criterion (v) in Section 1). Our goal is to eliminate discretization of the resulting image and to enhance the stability of the projected image to perturbations in the physical display setting.

In this article, we propose a new method for computing the shape of the object. Our method can be classified as the second approach described previously. The distinct aspect of our method is that we compute a continuous relationship. Owing to this property, the refractive surface becomes continuous, and the distribution of the light leaving the refractive surface is also continuous. The problems described earlier can be resolved as a result of this.

3. OUR METHOD

Our method consists of two steps. In the first step, we compute the relationship between the light reaching the incident surface and that reaching the screen. In the second step, we compute the shape of the refractive surface from the relationship.

3.1 Computing the Relationship

As shown in Figure 3, the z -axis is aligned with the direction of the incident parallel light (the light is coming from $+z$ toward $-z$), and the incident surface is a plane perpendicular to the z -axis. Points on the incident surface can be represented using x and y coordinates. As the light is perpendicular to the incident surface, it enters the object without refraction and reaches the refractive surface. The z coordinate of any point on the refractive surface can be represented as a single-valued function of x and y . The light refracts at the refractive surface and then reaches the screen. To distinguish the coordinates on the screen from those on the incident surface, we use u and v to describe the x and y coordinates on the screen. For simplicity, we use (x, y) or (u, v) and drop the z coordinate to describe the points on the incident surface and the screen.

The light path can be described by the relationship between a point $\mathbf{p}(x, y)$ on the incident surface and a point $\mathbf{q}(u, v)$ on the screen. In our method, we consider the case when this relationship is one-to-one, and represent the caustics by changing the density of the light reaching the screen. Thus, we need to find from which regions of the incident surface we need to gather light in order to form the intensity distribution of the desired caustic pattern. This problem can be formulated as the problem of finding a mapping from $\mathbf{q}(u, v)$ to $\mathbf{p}(x, y)$ satisfying the following two conditions.




Fig. 4. (a) The difference D (red and blue regions indicate positive and negative portions); (b) pressure field ϕ ; (c) vector field $\nabla\phi$.

The first condition is that the mapping is continuous over the entire domain. If this mapping is continuous, we can obtain a continuous refractive surface. The second condition is that this mapping preserves the light energy. Imagine the light beam shown in Figure 3. If the light incident on an infinitesimal region dS around point $\mathbf{p}(x, y)$ on the incident surface has intensity $L(x, y)$ (let us use L for simplicity), and the intensity of this light on an infinitesimal region dS' around point $\mathbf{q}(u, v)$ on the screen is $L'(u, v)$ (let us use L' for simplicity), then $LdS = L'dS'$. The aim of our method is to find this mapping. Since it is difficult to analytically compute this mapping, we use a geometric flow approach, as described next.

The mapping from $\mathbf{q}(u, v)$ to $\mathbf{p}(x, y)$ can be regarded as parameterization of x and y using u and v . That is, we start from an initial parameterization $x = u$ and $y = v$, and then maintain the continuity of the parameterization while modifying it so that it increasingly satisfies the conservation of light energy.

Under the parameterization of u and v , the ratio dS/dS' is equivalent to the determinant J of the Jacobian given by

$$J = \begin{vmatrix} \frac{\partial x}{\partial u} & \frac{\partial y}{\partial u} \\ \frac{\partial x}{\partial v} & \frac{\partial y}{\partial v} \end{vmatrix}. \quad (1)$$

Hence, $L' = LJ$. If the parameterization, or the mapping, satisfies $LJ = C$ at an arbitrary point, where C is the intensity of the caustics at the corresponding point on the screen, then that mapping is the solution to the problem.

During computation, $LJ = C$ is not always satisfied, and there is a difference. Let this difference $D(x, y)$ (again, we use D for simplicity) be

$$D = LJ - C. \quad (2)$$

We reduce this difference (see Figure 4(a)) by updating the parameterization to continuously modify the point $\mathbf{p}(x, y)$ on the incident surface, while fixing the corresponding point $\mathbf{q}(u, v)$ on the screen. Although we could consider the gradient ∇D (where $\nabla = (\frac{\partial}{\partial x}, \frac{\partial}{\partial y})$) and modify the points as $\frac{\partial \mathbf{p}}{\partial t} = \nabla D$ (where t is a virtual time to represent the computation process), ∇D is usually not continuous and may take an infinite value. Thus it is difficult to stably perform the computation using ∇D .

Instead, inspired by a concept in computational fluid dynamics, we modify the parameterization as follows (our method is related to one that generates area-preserving parameterization [Zou et al. 2011]). First, there are regions where the difference D is positive or negative. Analogous to those in computational fluid dynamics, we regard the positive and negative regions as source and sink, respectively. To relax the source and sink, we solve Poisson's equation

$$\nabla^2 \phi = -D, \quad (3)$$

in order to compute the pressure field ϕ (see Figure 4(b)), and update the points $\mathbf{p}(x, y)$ along the gradient of ϕ . That is, we let

$$\frac{\partial \mathbf{p}}{\partial t} = \nabla \phi. \quad (4)$$

Unlike ∇D , $\nabla \phi$ (see Figure 4(c)) is a much smoother vector field, thus we can stably update the points $\mathbf{p}(x, y)$. In our method, we repeatedly update the parameterization by computing the difference using Eq. (2), solving the Poisson's equation (3), and updating \mathbf{p} according to Eq. (4).

A physical interpretation of the preceding technique is as follows. If we want a region on the screen to be brighter/darker, the corresponding light beam needs to collect light from a wider/narrower region (i.e., dS should be larger/smaller). This implies that we have to locally “expand” or “shrink” dS . By solving Eq. (3) and computing $\nabla \phi$, we get a “compressible” flow field that realizes this expansion or shrinkage. By iteratively updating the parameterization, the contrast of the resulting caustics increases and approaches the contrast of the input pattern.

In our implementation, we use a triangular mesh to represent the parameterization. To compute the light energy, instead of computing the Jacobian, we compute a median dual mesh (where we have a face for each vertex in the triangular mesh, each of which is generated by connecting the midpoints of the edges and the centroids of the triangles adjacent to the corresponding vertex in the triangular mesh). We then integrate the intensities in each face in the median dual mesh to compute the light energy. Poisson's equations (3) and (7) are solved using a first-order Galerkin finite-element method. To update the vertices of the triangular mesh, we discretize Eq. (4) using the equation $\mathbf{p}(t_{n+1}) = \mathbf{p}(t_n) + \nabla \phi \Delta t$, where n indicates the n -th computation step, and Δt is determined so that there will be no flipped triangles. We obtain Δt by first regarding $\mathbf{p}(t_{n+1})$ as a function of Δt . Next, for each triangle, we compute the maximum Δt_i that the (i -th) triangle can take such that the area assigned to it remains nonnegative. Then, we compute the minimum value Δt_m among Δt_i (i.e., $\Delta t_m = \min_i \Delta t_i$). Finally, we set Δt to $0.5\Delta t_m$.

3.2 Computing the Refractive Surface

As we have assumed the refractive surface to be a single-valued function, its z coordinate can be represented as $z = h(x, y)$. A normal vector \mathbf{N} to the refractive surface can be represented by

$$\mathbf{N} = \left(\frac{\partial h}{\partial x}, \frac{\partial h}{\partial y}, -1 \right), \quad (5)$$

where \mathbf{N} is normalized so that its z coordinate is -1 . Considering only the x and y components, we have

$$\mathbf{N}_{xy} = \nabla h, \quad (6)$$

where $\mathbf{N}_{xy} = (N_x, N_y)$ is a vector composed of the x and y components of \mathbf{N} . Taking the divergence of both sides of Eq. (6), we obtain another Poisson's equation (similar to Yu et al. [2004])

$$\nabla^2 h = \nabla \cdot \mathbf{N}_{xy}. \quad (7)$$

Thus, if the normal vector at an arbitrary point on the refractive surface is known, we can use Eq. (7) to compute the shape of the refractive surface. To compute the normal vectors, we use Snell's law based on the parameterization obtained from Section 3.1, and get

$$\mathbf{N}_{xy} = \frac{1}{k}(\mathbf{q} - \mathbf{p}), \quad (8)$$




Fig. 5. Simulated caustic patterns. (a) The input pattern (©Playboy Magazine Nov. 1972); (b) the result using Finckh et al.'s method [2010]; (c) the result using Papas et al.'s method [2011]; (d) the result using our method. The images of (b) and (c) are reprinted from Papas et al. [2011].




Fig. 6. Computation of the relationship. (a) to (e): Results after 1, 3, 5, 7 and 25 steps, respectively. The top row shows the isoparametric lines for displaying the parameterization (please enlarge the figure to see the details). The bottom row shows the difference D (red and blue portions indicate positive and negative regions, and the difference is normalized such that the average intensity of the desired caustics corresponds to 100%).

where $(\mathbf{q} - \mathbf{p})$ means $(u - x, v - y)$,

$$k = \eta \sqrt{\|\mathbf{q} - \mathbf{p}\|^2 + (H - h)^2} - (H - h). \quad (9)$$

η is the refractive index (1.49 for an acrylate basin), and H is the distance between the screen and the transparent object.

Since the computation of the right-hand side of Eq. (7) needs h , we use an iteration method as follows. First, we initialize the refractive surface as a plane. Then, we iterate the following two steps. In the first step, we compute Eqs. (8) and (9) using the current refractive surface. In the second step, we solve Eq. (7) and update the refractive surface. Usually, only after a few iterations, we obtain convergence.

3.3 A Remark on the Uniqueness of the Solution

Let us briefly consider the relationship between a point $\mathbf{p}(x, y)$ on the incident surface and a point $\mathbf{q}(u, v)$ on the screen, obtained using our method. When we take the difference $(u - x, v - y)$ between the coordinates for all the points, this difference can be regarded as a vector field. According to the Helmholtz decomposition, a vector field can generally be decomposed into an incompressible field

(divergence-free field) and a compressible field (curl-free field). In our method, the initial parameterization is curl free, and since $\nabla \phi$ always satisfies $\nabla \times \nabla \phi = 0$, only curl free components are added into the parameterization. Thus, the parameterization obtained using our method contains no divergence-free components. On the other hand, any divergence-free flow would preserve the area. Thus, (when the incident light distribution is uniform) any divergence-free mapping, which preserves the area and hence preserves the light energy, can be combined with the mapping obtained using our method to produce the same desired caustic pattern (i.e., take a divergence-free mapping which maps (x', y') to (x, y) and define a composite mapping from (x', y') to (u, v)). In the sense that our mapping does not contain divergence-free components, our mapping can be regarded as the “minimum” mapping.

4. RESULTS

We show the results when using Figure 5(a) as the input. The resolution of the input image is 512×512 . The progress of the computation of the relationship is shown in Figure 6, which shows that it almost converges after seven steps. Based on the computed rela-




Fig. 7. The refractive surface computed using our method.



Fig. 8. Simulated results with different distances between the object and the screen. Left to right: the distances are $0.5\times$, $1.0\times$, $1.5\times$, and $2.0\times$ the focusing length.




Fig. 9. Other examples. From left to right: input patterns, isoparametric lines of the mapping, computed surfaces, and simulated caustic patterns.

tionship, we calculated the refractive surface as shown in Figure 7 (the computation described in Section 3.2 took three iterations to converge). When computing the surface, we assumed the dimensions of the object to be $100\text{mm} \times 100\text{mm}$ in the x - and y - axes, and the focal length to be 100mm . The maximum difference in the depth (z -axis) of the object was 14.4mm . In Figure 5(d), we show the simulated caustics using an inverse ray-tracing method. We can see that the texture of the hat and the details of the hair are nicely reproduced. Comparison to the simulated results obtained using previous methods [Finckh et al. 2010; Papas et al. 2011] is shown in Figure 5. Compared to the previous methods, our method is able to reproduce the detail.




Fig. 10. Real caustic patterns produced by fabricated objects using Papas et al. [2011] (left) and our method (right). The left image is reprinted from Papas et al. [2011]. Note that the photography environments are different, altering the impression of the caustic patterns.



Fig. 11. From left to right: photographs of the fabricated objects and their caustics, the simulated caustic patterns, and the caustic patterns of the fabricated objects.



Fig. 12. From left to right: the parameterization computed using remeshing, the refractive surface, the simulated caustics without remeshing, the simulated caustics with remeshing, and the caustics produced by the fabricated object.

In Figure 8, we show the simulated results when the distance between the refractive surface and the screen is changed. When we set the distance to half or 1.5 times the focal length, the caustic pattern is slightly distorted; nevertheless, we can still recognize the pattern. Moreover, the caustic pattern changes smoothly with the distance.

Some additional results are shown in Figure 9. Using our method, we can also generate vector-graphics-like caustic patterns as shown in Figure 9(a). Figure 9(e) shows an example containing Chinese characters.

In Figure 1, we show the fabricated object for Figure 5(a) and its caustic pattern. This object was fabricated using a CNC machine (Okuma FMR-40) with a ball end mill. We applied three passes to mill the surface: two rough milling passes and a finishing milling pass. In the first rough milling pass, we used a ball end mill of 3mm radius and the spacing (pitch) between neighboring milling paths was 0.5mm . In the second rough milling pass, the radius

was 0.5mm and the pitch was 0.05mm. The purpose of these rough milling passes is to progressively mill the surface and to reduce the mechanical load during the finishing pass. In the finishing milling pass, the radius was 0.5mm and the pitch was 0.02mm. The fabrication quality is essentially dependent on the setting of the finishing milling pass. The total fabrication time for these passes was approximately six days. We can vaguely perceive the vertical stripes, which are due to the quality of the path of the end mill (see Section 5 for a detailed discussion). In Figure 10, we compare the real caustic patterns produced by the fabricated object using our method against one fabricated using that of Papas et al. [2011]. Since with our computation, the relationship between the light reaching the incident surface and that reaching the screen is continuous, we can obtain caustic patterns with continuous tone representation also for dark regions, whereas the dark regions in Papas et al.’s caustic pattern tend to be too dark.

In Figure 11, we show the fabricated objects for Figures 9(a) and (e). As shown in Figure 11(a), we can see that the sharp features, such as the thin dark regions and the crisp corners, are also reproduced in the caustics generated by the fabricated object. We were also able to represent the characters in the caustics of the fabricated object, as shown in Figure 11(b).

The settings used in fabricating the object in Figure 11(a) are the same as those used for the Lena example. For the fabrication of the object in Figure 11(b), we used an NC machine (Roland MDX-40A) with a ball end mill. We applied three passes to mill the surface as well, and the settings of the radius and the pitch are the same as those used for the Lena example.

5. LIMITATIONS AND DISCUSSIONS

Since we are considering fully continuous parameterization, we cannot handle regions with zero intensity. To handle zero intensity, we need to extend our method to incorporate singular points or lines in the parameterization, in order to enable the splitting of light beams and to prevent light reaching such regions.

A limitation of the current implementation is that the mesh could get stuck (i.e., we cannot advect the mesh vertices). This could happen when the input pattern contains “localized high-contrast regions”, that is, when there is a very dark region adjacent to a bright region. In such a case, the triangles (in the mesh) corresponding to the bright/dark regions tend to expand/shrink greatly. Hence, triangles between these regions tend to be very thin and to impose severe restriction on Δt (see Section 3.1). It is interesting to note that when the mesh gets stuck, if we use the corresponding parameterization to compute the object surface, the resulting caustic pattern is still recognizable. The drawback is that the contrast is lower than required. The examples in Figures 9(c) and (f) are such cases.

Introducing a remeshing technique could prevent the mesh getting stuck. We implemented a simple remeshing by first detecting triangles which admit small Δt , then removing the vertices of the triangles, and inserting random vertices using dart throwing. In this way, thin triangles will be automatically removed, and we will be able to advect the mesh with large Δt . As shown in Figure 12, the contrast of the resulting caustics increases and reaches the same level as that of the input pattern. (The results shown in Section 4 are all computed without remeshing.) One drawback of this simple remeshing technique is that the isoparametric lines could get “dirty”, and the edges appearing in the caustics could become wavy. We think we could further introduce a mesh smoothing technique to keep the mesh “clean”. We leave this improvement for future work.

Under the current fabrication settings, we obtain slightly blurred caustics and striped artifacts. A possible reason for the blurriness is




Fig. 13. Left: an illustration of the milling path. Right: a top view. Paths shown in black color are for milling the surface and those shown in gray color are “subsidiary” paths, each of which is used to position the end mill to the next section without touching the surface (the end mill goes up, then moves to the next starting position straightly and horizontally, and then goes down).




Fig. 14. Left: an exaggerated illustration of a cross-section through the refractive surface (black line) and the fabricated surface (red line) due to the finite radius of the end mill (the materials milled out by the end mill are shown as dashed red lines). Right: an exaggerated illustration of the light refraction with respect to the fabricated surface.



Fig. 15. Simulated caustics taking into account bumps in the refractive surface. The radius of the end mill was 0.5mm in each case. The pitch of the milling path was set to: (a) 0.02mm, (b) 0.005mm, and (c) 0.02mm. In (c), deviations in the paths in the z direction are taken into account and are modeled using white noise.

as follows. For milling, the paths of the ball end mill are taken as illustrated in Figure 13 (for simplicity, we only consider the finishing milling pass). Since the pitch between consecutive paths is finite (0.02mm) and the ball end mill has a finite radius (0.5mm), the milled surface will have tiny bumps on it (see the exaggerated illustration in Figure 14). These bumps will cause some of the refracted light to deviate slightly, creating blurred caustics. In Figure 15(a), we show our simulation of the caustics taking these bumps into account. We find that if we set the pitch to 0.005mm and use the same ball end mill, we can significantly reduce this blurriness as simulated in Figure 15(b). Note that using a smaller pitch means that the total length of the milling path will be longer, hence the end mill will gall more during the milling process, which could be another cause of the artifact. A highly strengthened end mill would be needed for this fine setting.

A possible reason for the striped artifact is as follows. With the current setting, the fabrication proceeds by repeatedly milling the surface from one side to the other (as in Figure 13), and each time

the ball end mill enters the material, a physical collision takes place between the end mill and the surface. The impact of this sometimes causes the end mill to slightly deviate from its intended coordinates, and this deviation will be carried over the remaining milling process. In Figure 15(c), we model this deviation as white noise (in the z direction), and obtain similar striped artifacts in the simulated caustics (in this example, the pitch is set to the same as our fabrication setting, 0.02mm). Hence, we believe these stripe artifacts can be reduced if we are able to lessen this physical impact (for example, by using a single smoothly connected milling path, which traces mostly along the contour lines of the surface, in order to keep the end mill moving within the surface as much as possible and to reduce the events where the end mill leaves and reenters the surface).

6. CONCLUSIONS AND FUTURE WORK

In this article, we have presented a new method for computing the smooth refractive surfaces that generate user-specified caustics. Using our method, the relationship between the light reaching the incident surface and that reaching the screen is computed as a continuous function. As a result, the computed refractive surface is smooth. We have demonstrated that we can generate continuous and high-quality caustics through the simulated results, and have shown the effectiveness of our method. For future work, we would like to develop a clean remeshing technique which enables better edge representation for high-contrast input patterns.

Acknowledgments

This work is supported by JSPS Grant-in-Aid for Research Activity Start-up (23800011). The stay of Y. Yue at Columbia University is funded by JSPS Postdoctoral Fellowships for Research Abroad. We are deeply grateful to MTM Corporation (Japan) for fabricating the transparent objects and for useful discussions.

REFERENCES

- ALEXA, M. AND MATUSIK, W. 2010. Reliefs as images. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2010)* 29, 4, 60:1–60:7.
- BARAN, I., KELLER, P., BRADLEY, D., COROS, S., JAROSZ, W., NOWROUZEZAHRAI, D., AND GROSS, M. 2012. Manufacturing layered attenuators for multiple prescribed shadow images. *Computer Graphics Forum (Proc. of EUROGRAPHICS 2012)* 31, 2, 603–610.
- BERMANO, A., BARAN, I., ALEXA, M., AND MATUSIK, W. 2012. Shadowpix: Multiple images from self shadowing. *Computer Graphics Forum (Proc. of EUROGRAPHICS 2012)* 31, 2, 593–602.
- DONG, Y., WANG, J., PELLACINI, F., TONG, X., AND GUO, B. 2010. Fabricating spatially-varying subsurface scattering. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2010)* 29, 4, 62:1–62:10.
- FINCKH, M., DAMMERTZ, H., AND LENSCH, H. P. A. 2010. Geometry construction from caustic images. In *Proceedings of the 11th European Conference on Computer Vision*. 464–477.
- FUCHS, M., RASKER, R., SEIDEL, H.-P., AND LENSCH, H. 2008. Toward passive 6D reflectance field displays. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2008)* 27, 3, 58:1–58:8.
- HAŠAN, M., FUCHS, M., MATUSIK, W., PFISTER, H., AND RUSINKIEWICZ, S. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2010)* 29, 4, 61:1–61:10.
- MITRA, N. J. AND PAULY, M. 2009. Shadow art. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia 2009)* 28, 5, 156:1–156:7.
- PAPAS, M., HOUIT, T., NOWROUZEZAHRAI, D., GROSS, M., AND JAROSZ, W. 2012. The magic lens: refractive steganography. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia 2012)* 31, 6, 186:1–186:10.
- PAPAS, M., JAROSZ, W., JAKOB, W., RUSINKIEWICZ, S., MATUSIK, W., AND WEYRICH, T. 2011. Goal-based caustics. *Computer Graphics Forum (Proc. of Eurographics 2011)* 30, 2, 503–511.
- REGG, C., RUSINKIEWICZ, S., MATUSIK, W., AND GROSS, M. 2010. Computational highlight holography. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia 2010)* 29, 6, 170:1–170:12.
- WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2009)* 28, 3, 32:1–32:6.
- YU, Y., ZHOU, K., XU, D., SHI, X., BAO, H., GUO, B., AND SHUM, H.-Y. 2004. Mesh editing with poisson-based gradient field manipulation. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2004)* 23, 3, 644–651.
- YUE, Y., IWASAKI, K., CHEN, B.-Y., DOBASHI, Y., AND NISHITA, T. 2012. Pixel art with refracted light by rearrangeable sticks. *Computer Graphics Forum (Proc. of EUROGRAPHICS 2012)* 31, 2, 575–582.
- ZOU, G., HU, J., GU, X., AND HUA, J. 2011. Authalic parameterization of general surfaces using lie advection. *IEEE Transactions on Visualization and Computer Graphics* 17, 12, 2005–2014.

Received February 2013; revised December 2013; accepted January 2014