

A Reflective Light Stage

Pieter Peers Tim Hawkins Paul Debevec

USC Institute for Creative Technologies
Technical Report ICT-TR-04.2006

December 23, 2006

Abstract

We present a novel acquisition device to capture high resolution 4D reflectance fields of real scenes. The device consists of a concave hemispherical surface coated with a rough specular paint and a digital video projector with a fish-eye lens positioned near the center of the hemisphere. The scene is placed near the projector, also near the center, and photographed from a fixed vantage point. The projector projects a high-resolution image of incident illumination which is reflected by the rough hemispherical surface to become the illumination on the scene. We demonstrate the utility of this device by capturing a high resolution hemispherical reflectance field of a specular object which would be difficult to capture using previous acquisition techniques.

1 Introduction

Image-based relighting has been a popular research topic in recent years (e.g., [3, 15, 21, 25, 33]). A key element in the image-based relighting process is the notion of a reflectance field [3]. This reflectance field is an $8D$ function, describing the relation between incident and exitant illumination on a bounding volume surrounding the scene.

One of the most popular techniques for acquiring such a reflectance field is by capturing the scene under different illumination conditions. Due to the linearity of light transport, a relit image can be created by combining the captured photographs weighted proportionally to the contribution of the illumination condition in the desired incident light field.

Capturing this $8D$ function completely is currently still impractical. Therefore, a number of simplifications are commonly applied to the $8D$ reflectance field. First, a fixed vantage point is assumed. This reduces the dimensionality of the exitant illumination to a $2D$ field (i.e., a photograph of the scene) instead of the $4D$ original exitant light field. Secondly, we assume that incident illumination originates from a distant source, effectively reducing the incident illumination to directional illumination, a $2D$ field [3]. Recently, this second restriction has been relaxed [19] at the cost of an increase in acquisition complexity.

Current reflectance field acquisition devices either sample the incident light field (e.g., the Light Stage [3]), or use a dense sampling device, such as a CRT monitor (e.g., [35]), that only covers a small portion of the entire incident domain.

In this report we present a novel device for capturing $4D$ reflectance fields of real scenes, which consists of a rough specular hemispherical mirror, that reflects controlled illumination emitted from a projector equipped with a fish-eye lens onto the scene. This device has the following properties with respect to the captured reflectance fields:

- **No moving elements:** an acquisition device without moving elements has a number of advantages over devices which require mechanical movement. For one, the acquisition can potentially be completed in less time, since no time is spent moving parts around. Also, a device without moving parts is in general easier to design, construct, and calibrate.
- **Hemispherical incident illumination:** a $4D$ reflectance field is characterized by a direction (i.e., spherical) incident light field. The presented de-

vice can sample half of the sphere of incident directions without mechanical moving parts, and could be easily extended to capture the whole sphere by rotating the camera and scene by 180 degrees and acquiring a second dataset.

- **Continuous, high-resolution reflectance fields:** are needed to capture scenes containing glossy and specular materials. Methods relying on sampling are limited, due to Nyquist's theorem, in the kinds of materials they can accurately capture (e.g., materials exhibiting BRDFs with high frequency components are notoriously difficult to sample without aliasing).

The remainder of this report is organized as follows. First previous presented acquisition devices are briefly investigated in section 2. Next, the new acquisition device is described in detail (section 3). Calibration issues and some results are discussed in section 4. Finally, section 5 concludes this report.

2 Previous Work

Two principal approaches have been used in image-based relighting to acquire reflectance fields.

The first one, and probably the most well-known, are the techniques based on sampling. These methods include the different Light Stage devices [3, 4, 6, 9, 12, 13, 14, 16, 17, 18, 19, 20, 23, 31], that basically sample the incident light field from a spherical bounding volume, by either moving a light source over this bound of the volume [3], or by switching discretely placed light sources on the surrounding sphere on or off [4]. The former is also used in Masselus *et al.* [19] to capture $6D$ reflectance fields. In this case, however, the single light source is replaced by a projector, enabling angularly varying control of the incident illumination control for each lighting direction.

Sample-based techniques generally suffer from aliasing when capturing high frequency reflectance functions. Mohan *et al.* [22] and Fuchs *et al.* [7] solve this problem by pre-filtering the reflectance functions. Pre-filtering is achieved by using (overlapping) extended area light sources. The extended area light source is obtained by aiming a light source at a (sufficiently large) diffuse surface, which in turn reflects the illumination on the scene. Both Mohan *et al.* and Fuchs *et al.* implement this by using the surrounding bounding volume (i.e., room) as a diffuse reflector, and place (and move) the light source inside the volume. Pre-filtering can reduce aliasing effects in the captured reflectance field, at the cost of smoothing out high frequency details in the reflectance functions.

An interesting variant on the sampling approach, are the dual methods [10, 29]. These methods build on the generalized reciprocity principle to swap the position of the camera and the light source. Sen *et al.* [29] use this principle to parallelize the acquisition of $6D$ reflectance fields by positioning multiple cameras around the scene and illuminate it from a single projector at a fixed position (i.e., the dual camera), and consequently suffer from the same aliasing problems as other sample-based methods. Hawkins *et al.* [10], however, use this principle to acquire high resolution, continuous reflectance fields. They use a laser and a galvanometer instead of a projector, but instead of capturing the exitant light field directly as Sen *et al.*, they first reflect it off a diffuse sphere surrounding the scene, before capturing it with a single camera equipped with a fish-eye lens. Two dimensions of the incident light field are integrated out due to the diffuse reflection. The resolution of the two remaining dimensions are only limited by the camera resolution. A disadvantage of this setup is the need to compensate for the significant amount

of light interreflected within the sphere, the loss of the polarization of the incident illumination (shown to be useful for reflectometry in [3]), and the fact that only a small percentage of the exitant illumination is reflected towards the camera.

A second approach for capturing reflectance functions uses a high resolution, semi-continuous, controllable light source such as a CRT or an LCD projector. Because of the large number of (very small) light sources it is impractical to turn on each light source individually. Various techniques have been developed to overcome this problem [1, 21, 26, 27, 32, 34, 35]. A disadvantage of using a CRT monitor as controllable light source is the relatively small fraction of the sphere of directions that is easily covered (for example, one sixth). Furthermore, the size of the scene is also limited by the size of the emitter.

Other notable acquisition devices in different domains in computer graphics are the mirrored half-dome used by Ward [30] to capture BRDFs (Bidirectional Reflectance Distribution Functions), and the concave parabolic mirror used by Dana [2] to capture BTFs (Bidirectional Texture Functions). Ward projects a moving collimated light source onto a flat material sample sitting near the center of a half-silvered hemisphere. The full hemisphere of radiant light reflected from the sample is reflected by the half-silvered mirror back toward the center where it is imaged using a camera with a fish-eye lens, yielding a 2D slice of the BRDF. They note that the device is not accurate enough to capture highly specular surfaces due to the limited extent of the collimated beam, and due to the inaccuracies in shape of the mirror. Dana uses a parabolic concave mirror to view and illuminate a single surface point of a textured sample. Different surface points are sampled by translating either the mirror or the sample. This design requires that the incident and exitant illumination directions follow the same optical path which is achieved using a beam splitter. However, this device does not cover the complete hemisphere of viewing and lighting directions.

3 Device Description

We set two goals when designing the apparatus:

1. The device should be able to generate a high resolution continuous incident light field on a scene of reasonable size.
2. The device should be able to generate incident illumination from as much of the sphere of directions as possible.

Ideally we would like to have a spherical emitter that can send different intensities from a specific direction. This is the approach followed by the Light Stage devices [3]. However, constructing a continuous incident field of illumination in this manner is difficult. As shown in the previous section, the most successful continuous acquisition devices are based on display technologies. It is thus logical to base our novel device on such a technology.

We assume that the incident illumination is only directional and not spatially varying across the scene, as is the case when a small scene is lit by a distant environment. Our approach, as a general overview, uses a concave hemi-ellipsoidal mirrored surface to reflect light emitted from a fisheye video projector at one of its foci back in toward the scene (such as a small object or a human face) placed at its other focus. In this arrangement, any illumination emitted from the projector will reflect back toward the scene due to the geometry of the ellipse.

Because spherical displays are not readily available, we resort to a similar setup as Ward et al. [30]. We use a concave spheroid mirror-like surface (e.g., an ellipsoid, with two identical axes, and a third longer axis), and place a projector equipped with a fish-eye lens at one focus. The scene is placed at the other focus (see Figure 1).

A video projector with a fisheye lens behaves similarly to a point light source (though its illumination generally varies with angle). With a perfectly specular ideal ellipsoid mirror, light emanating from the projector would focus to a small point at the other focus, which would not be enough spatial extent to illuminate a small scene. To broaden the area that is illuminated by the projector, we use a rough specular paint to coat the inside of the spheroid mirror (also seen in Figure 1). This rough specular surface will ensure that incident illumination from the projector is projected in a small solid angle in the direction of the focus point, but without reflecting too much illumination towards other points on the mirror (which would cause undesirable interreflection). A diffuse coating, such as used in the Dual Light Stage [10], would exhibit a larger degree of interreflection relative to the albedo of




Figure 1: A diagram showing the path of light in the reflective light stage.

the material. Furthermore, the rough specular surface maintains additional properties of the incident illumination such as polarization, which is potentially useful for performing reflectance measurements on specular and diffuse reflectance independently as in [3]). The rough specularity also reduces focusing problems that arise from a mirror surface that deviates from a true spheroid shape. The spheroidal surface in this work was formed using an acrylic material with a bumpy mesostructure to add additional roughness to the specular reflection.

The size of the major axis of the spheroid can be easily computer given de desired distance of the foci, and the length of the minor axis:

$$a^2 = c^2 + b^2,$$

where a is the length of the major axis, b is the length of the smaller axis, and c is the distance from a focus point to the center. For example, when a minor axis of 1 meter is desired, and a focal distance of 0.25 meter, then the major axis has a length of 1.03 meter.

It is somewhat easier to build a hemispherical surface than an hemiellipsoidal surface, in that acrylic blow-molding companies frequently have circular molds

and apertures available. If we look at the ellipsoid case above, we see that the major and minor axes only differ 3% in length. Given the fact that we also use a rough specular mirror, some deviation from the *perfect* spheroid can be tolerated without introducing significant artifacts. Since the difference between the axes for this typical case is very small, we opted for using a hemispherical mirror.

Figure 2 shows our acquisition device of approximately one meter in diameter. In order to photograph the scene, a hole is drilled through the top of the hemisphere. The hemisphere is constructed by inflating heated acrylic into a hemispherical shape. The acrylic used contains a *bumpy* texture, which helps to *roughen* the specular reflections. The hemisphere is coated by a thin layer of silver car-paint, itself a rough specular material with high reflectivity, without the usual gloss layer to avoid mirror-like reflections. A sample of this material is shown in figure 3.




Figure 2: Photographs of our setup. Left: the camera, scene and projector are marked. Right: only hemisphere is shown.




Figure 3: The inside of the hemisphere has a *bumpy* texture, coated with a thin layer of silver car-paint, a sample of which is shown above.

4 Results and Discussion

We demonstrate the usefulness of the reflective light stage by capturing a reflectance field of a shiny object. The reflectance field is captured by subdividing the incident hemisphere of illumination into a 20×15 set of discrete solid angles that completely cover the hemisphere, but do not overlap. An HDR photograph [5] is captured while illuminating each of the solid angles individually. We create this set of solid angles by regularly discretizing the projector image plane into square regions.

Due to the fisheye lens, the exact mapping between projector pixels and incident directions does not fit a common projection model. To calibrate for this mapping, we emit horizontal and vertical stripe patterns on a mirrored sphere positioned at the location of the sphere. From the observed reflections of the stripe patterns, it is straightforward to determine the mapping between the projector pixels and incident illumination directions. In Figure 4 the resulting calibration is shown for the setup depicted in Figure 2. The black spot (Figure 4(a)) in the center is the camera (a hole in the spherical reflector). Because our projector has a 4 : 3 aspect ratio, the bottom part of the hemisphere is not covered. The direction in the mapping corresponding to the projector is also marked (Figure 4(b)).




Figure 4: The mapping between projector pixels and incident illumination directions illustrated in a latitude-longitude parameterization; red and green intensity indicate the vertical and horizontal index of the lighting condition corresponding to each direction. (a) A hole in the spherical mirror allows the camera to photograph the scene. (b) The direction of the projector lens.




Figure 5: The test object used to demonstrate our acquisition device.

Once the mapping between projector pixels and incident illumination directions is known, it is straightforward to create a relit image in a similar manner to [3]. To illustrate our technique, we capture a shiny teapot wind-up toy (Figure 5). This reflectance of this scene would be difficult to capture well using any of the previous sampling methods, and it would not be possible to capture such a large portion of the reflectance field with any of the other (CRT based) acquisition methods. The results in Figure 6 show that we are able to obtain relit results with a continuous specular reflection even with the moderate sampling resolution of 20×15 .


As long as the incident illumination resolution is limited (as it is in this report), the acquisition duration is still reasonable. However, to capture very high resolution reflectance fields, more sophisticated methods will be required such as using wavelet noise [27]. Other illumination patterns such as (hemi)spherical harmonics and multiplexed illumination [28] merit further investigation for use with this device.



(a)



(b)



(c)

Figure 6: Relit images of the scene computed from a reflectance field captured with the reflective light stage. Three different illumination conditions are used to generate the relit images: (a) uniform white illumination, (b) the Grace Cathedral light probe image, and (c) the Galileo's Tomb light probe image.

5 Conclusions and Future Work

We have presented a new device to acquire high resolution, alias-free reflectance fields for image-based relighting. The device consists of a rough specular hemispherical mirror, a projector equipped with a fish-eye lens, and a digital camera. The device has no moving parts and is able to acquire the reflectance response of a scene from the full hemisphere of incident illumination.

The effectiveness of the device was demonstrated on a small scene containing a shiny teapot for which a hemispherical reflectance field was captured. This example also shows the need to develop new sampling methods that capture the response of the object to higher resolution incident illumination conditions and that are able to capture the reflectance field with sub-linear acquisition complexity.

Currently, our setup acquires only 4D reflectance fields. Extending the system to the capture of 8D reflectance fields could be done by using a sharp (as opposed to rough) concave mirror surface and aiming a video projector and camera through a lenslet array [11, 24] placed at one of the foci as in Figure 7. In this system, the projector and camera receive the same view of the lenslet array using a half-silvered mirror in a manner similar to [8]. In this arrangement, the lenslet array transforms the 2D image from the projector and camera into 4D excitant and incident light fields with both spatial and angular variation. The hemispherical surface further reflects the emitted field of illumination toward the subject as a hemispherical incident light field, and then returns the reflected light back toward the lenslet array so that its angular and spatial variation can be imaged by the camera. Since both projectors and cameras already exist at resolutions exceeding 8 million pixels, useful sampling resolution in all eight dimensions could be attainable, allowing for spatially- and angularly-varying illumination to be simulated on the scene for any virtual viewing position. Clearly, a sublinear acquisition complexity method as in [29] or [27] would be required to capture such an 8D dataset in a reasonable amount of time. This system presents an alternative approach to the kaleidoscopic 8D reflectance field acquisition system proposed in [8].

Acknowledgements

The authors gratefully acknowledge the support of Tomas Pereira and Brian Miller in acquiring and painting the mirrored hemisphere. This work was sponsored by the USC Office of the Provost and USA RDECOM. The content of the information does not necessarily reflect the position or the policy of the US Government, and no official endorsement should be inferred.




Figure 7: An 8D reflectance field could be captured with a reflective light stage by placing a lenslet array at one focus and using a sharp (as opposed to rough) mirrored surface for the reflection.

References

- [1] Yung-Yu Chuang, Douglas E. Zongker, Joel Hindorff, Brian Curless, David H. Salesin, and Richard Szeliski. Environment matting extensions: Towards higher accuracy and real-time capture. In *SIGGRAPH 2000, Computer Graphics Proceedings*, Annual Conference Series, 2000.
- [2] Kirstin Dana. BRDF/BTF measurement device. In *ICCV*, pages 460–466, 2001.
- [3] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 145–156, 2000.
- [4] Paul Debevec, Andreas Wenger, Chris Tchou, Andrew Gardner, Jamie Waese, and Tim Hawkins. A lighting reproduction approach to live-action compositing. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 547–556, 2002.
- [5] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 369–378, 1997.
- [6] Per Einarsson, Charles-Felix Chabert, Andrew Jones, Wan-Chun Ma, Bruce Lamond, Tim Hawkins, Mark Bolas, Sebastian Sylwan, and Paul Debevec. Relighting human locomotion with flowed reflectance fields. In *EGSR'06: Proceedings of the 17th Eurographics Workshop on Rendering*, pages 183–194, 2006.
- [7] Martin Fuchs, Volker Blanz, and Hans-Peter Seidel. Bayesian relighting. In *EGSR '05: Proceedings of the 16th Eurographics workshop on Rendering*, pages 157–164, July 2005.
- [8] Jefferson Y. Han and Ken Perlin. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Transactions on Graphics*, 22(3):741–748, July 2003.
- [9] Tim Hawkins, Jonathan Cohen, and Paul Debevec. A photometric approach to digitizing cultural artifacts. In *VAST '01: Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*, pages 333–342, 2001.

- [10] Tim Hawkins, Per Einarsson, and Paul E. Debevec. A dual light stage. In *EGSR '05: Proceedings of the 16th Eurographics Workshop on Rendering*, pages 91–98, 2005.
- [11] Aaron Isaksen, Leonard McMillan, and Steven J. Gortler. Dynamically reparameterized light fields. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 297–306, July 2000.
- [12] Andrew Jones, Andrew Gardner, Mark Bolas, Ian McDowall, and Paul Debevec. Performance geometry capture for spatially varying relighting, SIGGRAPH 2005, Sktech.
- [13] Melisa L. Koudelka, Sebastian Magda, Peter N. Belhumeur, and David Kriegman. Image-based modeling and rendering of surfaces with arbitrary BRDFs. In *2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2001), Vol II*, pages 568–575, 2001.
- [14] Zhoulin Lin, Tien-Tsin Wong, and Heung-Yeung Shum. Relighting with the reflected irradiance field: Representation, sampling and reconstruction. In *2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2001), Vol I*, pages 561–567, December 2001.
- [15] Céline Loscos, Marie-Claude Frasson, George Drettakis, Bruce Walter, Xavier Granier, and Pierre Poulin. Interactive virtual relighting and remodeling of real scenes. In *EGWR '99 :Proceedings of the 10th Eurographics workshop on Rendering*, volume 10, pages 235–246, Jun 1999.
- [16] Tom Malzbender, Dan Gelb, and Hans Wolters. Polynomial texture maps. In *SIGGRAPH 2001, Computer Graphics Proceedings*, Annual Conference Series, pages 519–528, 2001.
- [17] Wai man Pang, Tien-Tsin Wong, and Pheng-Ann Heng. Estimating light vectors in real time. *IEEE Computer Graphics and Applications*, 24(3):36–43, 2004.
- [18] Vincent Masselus, Philip Dutré, and Frederik Anrys. The free-form light stage. In *EGWR '02: Proceedings of the 13th Eurographics workshop on Rendering*, pages 247–256, 2002.
- [19] Vincent Masselus, Pieter Peers, Philip Dutré, and Yves D. Willems. Relighting with 4D incident light fields. *ACM Transactions on Graphics*, 22(3):613–620, 2003.

- [20] Vincent Masselus, Pieter Peers, Philip Dutré, and Yves D. Willems. Smooth reconstruction and compact representation of reflectance functions for image-based relighting. In *EGSR '04: Proceedings of the 15th Eurographics workshop on Rendering*, pages 287–298, 2004.
- [21] Wojciech Matusik, Matthew Loper, and Hanspeter Pfister. Progressively-refined reflectance functions from natural illumination. In *EGSR '04: Proceedings of the 15th Eurographics Workshop on Rendering*, pages 299–308, 2004.
- [22] Ankit Mohan, Jack Tumblin, Bobby Bodenheimer, Cindy Grimm, and Reynold J. Bailey. Table-top computed lighting for practical digital photography. In *EGSR '05: Proceedings of the 16th Eurographics workshop on Rendering*, pages 165–172, 2005.
- [23] Francesc Moreno-Noguer, Shree K. Nayar, and Peter N. Belhumeur. Optimal illumination for image and video relighting. In *Proceeding of IEEE European Conference on Visual Media Production (CVMP)*, 2005.
- [24] Ren Ng, Marc Levoy, Mathieu Bredif, Gene Duval, Mark Horowitz, and Pat Hanrahan. Light field photography with a hand-held plenoptic camera. Technical Report CSTR 2005-02, Stanford University Computer Science, April 2005.
- [25] Jeffrey S. Nimeroff, Eero Simoncelli, and Julie Dorsey. Efficient Re-rendering of Naturally Illuminated Environments. In *Fifth Eurographics Workshop on Rendering*, pages 359–373, 1994.
- [26] Pieter Peers and Philip Dutré. Wavelet environment matting. In *EGSR '03: Proceedings of the 14th Eurographics workshop on Rendering*, pages 157–166, 2003.
- [27] Pieter Peers and Philip Dutré. Inferring reflectance functions from wavelet noise. In *EGSR '05: Proceedings of the 16th Eurographics workshop on Rendering*, pages 173–182, 2005.
- [28] Yoav Y. Schechner, Shree K. Nayar, and Peter N. Belhumeur. A theory of multiplexed illumination. In *ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, page 808, 2003.
- [29] Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R. Marschner, Mark Horowitz, Marc Levoy, and Hendrik P. A. Lensch. Dual photography. *ACM Transactions on Graphics*, 24(3):745–755, 2005.

- [30] Gregory J. Ward. Measuring and modeling anisotropic reflection. *Journal of Computer Graphics*, 26(2):265–272, 1992.
- [31] Andreas Wenger, Andrew Gardner, Chris Tchou, Jonas Unger, Tim Hawkins, and Paul Debevec. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics*, 24(3):756–764, 2005.
- [32] Yonatna Wexler, Andrew W. Fitzgibbon, and Andrew Zisserman. Image-based environment matting. In *EGWR '02: Proceedings of the 13th Eurographics workshop on Rendering*, 2002.
- [33] Tien-Tsin Wong, Pheng-Ann Heng, Siu-Hang Or, and Wai-Yin Ng. Image-based rendering with controllable illumination. In *EGWR '97: Proceedings of the 8th Eurographics workshop on Rendering*, pages 13–22, June 1997.
- [34] Jiayuan Zhu and Yee-Hong Yang. Frequency-based environment matting. In *12th Pacific Conference on Computer Graphics and Applications (PG 2004)*, pages 402–410, 2004.
- [35] Douglas E. Zongker, Dawn M. Werner, Brian Curless, and David H. Salesin. Environment matting and compositing. In *SIGGRAPH 1999, Computer Graphics Proceedings*, Annual Conference Series, pages 205–214, 1999.