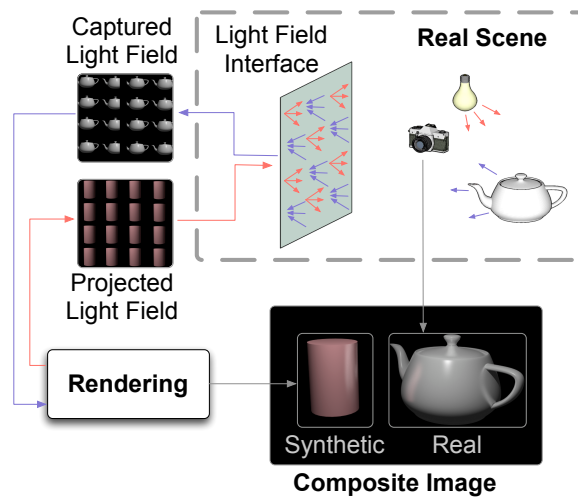


Light Field Transfer: Global Illumination Between Real and Synthetic Objects


Oliver Cossairt

Shree Nayar
Columbia University *

Ravi Ramamoorthi



(a) Dataflow for light field transfer



(b) Global illumination rendered using light field transfer

Figure 1: (a) A conceptual diagram of the proposed light field transfer method. Indirect lighting is transferred between real and virtual scenes via a light field interface. (b) A composite image rendered using this algorithm without any knowledge of the geometry or material properties of real objects. Note that a variety of interreflections are visible though all objects on the left are synthetic and all objects on the right are real. Dynamic objects such as the real hand in (b) can be included because the proposed light field transfer is done in near real-time.

Abstract

We present a novel image-based method for compositing real and synthetic objects in the same scene with a high degree of visual realism. Ours is the first technique to allow global illumination and near-field lighting effects between both real and synthetic objects at interactive rates, without needing a geometric and material model of the real scene. We achieve this by using a light field interface between real and synthetic components—thus, indirect illumination can be simulated using only two 4D light fields, one captured from and one projected onto the real scene. Multiple bounces of interreflections are obtained simply by iterating this approach. The interactivity of our technique enables its use with time-varying scenes, including dynamic objects. This is in sharp contrast to the alternative approach of using 6D or 8D light transport functions of real objects, which are very expensive in terms of acquisition and storage and hence not suitable for real-time applications. In our method, 4D radiance fields are simultaneously captured and projected by using a lens array, video camera, and digital projector. The method supports full global illumination with restricted object placement, and accommodates moderately specular materials. We implement a complete system and show several example scene compositions that demonstrate global illumination effects between dynamic real

and synthetic objects. Our implementation requires a single point light source and dark background.

CR Categories: I.3.7 [COMPUTER GRAPHICS]: Three-Dimensional Graphics and Realism—Raytracing; I.4.1 [COMPUTER GRAPHICS]: Three-Dimensional Graphics and Realism—Virtual Reality;

Keywords: augmented reality, global illumination, light field, image-based relighting

1 Introduction

Compositing real and synthetic objects in the same scene is important for many computer graphics applications such as visual effects, augmented reality, and architectural rendering. The makers of many of today's feature films need to insert digital actors into a scene containing real actors and props. To make composite scenes such as these compelling, lighting from local and distant sources must be consistent between real and synthetic components. Moreover, when synthetic and real objects are in close proximity to one another, indirect lighting plays an important role in creating realistic renderings. Near-field lighting models have been introduced for synthetic and real objects independently, but a Hybrid Global Illumination (HGI) method that incorporates both real and synthetic objects and accounts for the interplay between them has yet to be proposed. Most existing methods only account for distant lighting, and do not consider interreflections between real and synthetic scenes. In addition, many existing techniques are labor intensive.

Hybrid scenes would benefit from a HGI method that incorporates near-field illumination as well as indirect light transfer between real and virtual domains. In this paper, we propose just such a method, the first of its kind. The method is efficient because it does not require an estimate of geometry and material properties. At the same time, it also does not require the capture and simulation of large image-based datasets. Owing to this efficiency, we are able to show the first videos of time-varying composite scenes including dynamic objects, with both direct and indirect lighting.

*e-mail: {ollie,nayar,ravir}@cs.columbia.edu

ACM Reference Format
Cossairt, O., Nayar, S., Ramamoorthi, R. 2008. Light Field Transfer: Global Illumination Between Real and Synthetic Objects. *ACM Trans. Graph.* 27, 3, Article 57 (August 2008), 6 pages. DOI = 10.1145/1360612.1360656 <http://doi.acm.org/10.1145/1360612.1360656>.

Copyright Notice
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, fax +1 (212) 869-0481, or permissions@acm.org.
© 2008 ACM 0730-0301/2008/03-ART57 \$5.00 DOI 10.1145/1360612.1360656
<http://doi.acm.org/10.1145/1360612.1360656>

To implement our system, we require an interface to measure light rays from, and emit light rays into the real scene. As shown in Fig. 1(a), this interface consists of a light field projector and a light field camera. The camera captures a light field of the real scene, which is sent to the rendering engine to relight the synthetic objects. The rendering engine generates a light field of the synthetic scene, which is sent to the projector to relight the real objects. This exchange of light fields between real and virtual worlds can be done iteratively so as to allow multiple bounces of indirect lighting. The book-keeping cost is only two 4D datasets, which can be processed in real-time with current computing power.

Our light field transfer method has several advantages over previous work. Since it is purely image-based, no model of the geometry and reflectance of the real scene is needed. Conventional image-based relighting methods could be employed, but they require a full 8D dataset for complete near-field and indirect lighting effects (both incoming and outgoing). In contrast, we project/acquire only a single 4D light field on/from the real object for each bounce, corresponding to the near-field illumination actually present in the scene and reflected from virtual objects—we do not need to acquire the appearance of the real scene under all possible lighting conditions as in conventional image-based relighting. Moreover, we achieve interactive frame rates for both acquisition and display, that are extremely difficult with either traditional computer graphics models and global illumination rendering, or image-based relighting of high-dimensional datasets. This enables HGI with time-varying scenes, where the real and synthetic objects can be moving or even changing in geometry and reflectance.


Though our method reduces dataset complexity, sampling is still a bottleneck. We restrict the shininess of specular objects in our scenes to compensate for coarse sampling granularity at our light field interface. Our method adds a new geometric constraint: all rays transferred between real and virtual domains must intersect the light field interface. The planar interface used in our implementation requires that real and synthetic objects do not occlude one another. Our examples use a simplified, single source direct lighting model, and require a dark background. In the examples, we assume a linear mapping between camera and projector intensities, and we ignore cross-talk between captured and projected color channels. Possible improvements that allow greater flexibility in scene configurations are straightforward and discussed in Section 6. Despite the aforementioned constraints, we show that our method elegantly supports a variety of scene configurations.

Our technique can be used to create composite images such as the one shown in Fig. 1(b). In this example, all the objects on the left are synthetic and rendered (a hand holding a photo frame above a wooden table), and all the objects on the right are real and captured with a camera. Although the rendered and captured images are generated independently, objects in each include indirect lighting from one another. In the rendered image, the synthetic photo frame shows glossy reflections of objects in the real scene. A reflection of the synthetic photograph is visible on the real metal bowl. When the rendered and captured images are composited, the final result seamlessly combines real and synthetic objects in the same image with consistent direct and global illumination, that also changes correctly in real-time as the real and synthetic objects are moved. Other examples show diffuse-diffuse interactions and soft shadows (Fig. 5). We can interactively control the geometric and material properties of our computer-generated models, and composite them in real-time with captured video footage of real scenes.


2 Related Work

A number of techniques have been introduced that allow indirect lighting for hybrid scenes. Most require either geometric or material properties of the real scene to be known. Others relax this constraint, but require exorbitantly large datasets.

Parallel Rendering The idea of introducing an interface to transfer indirect illumination was first presented by Arnaldi et al. [1991]



(a) Illustration of the hardware layout.



(b) Photograph of the actual hardware used.

Figure 2: The hardware used for light field transfer. The system consists of a light field capture and projection unit, a view camera, and direct lighting. A camera and projector share a common lens array that multiplexes 2D radiance patterns into 4D patterns.

for parallel rendering. In this method, virtual walls subdivide a complex synthetic scene and global illumination for each subdivision is computed in parallel. Light field transfer can be understood as an extension of the Virtual Walls method, suitable for HGI.

Inverse Rendering: Yu et al. [1999] introduce an inverse global illumination method to estimate the material properties of real objects in a scene based on known object geometry and lighting. Ramamoorthi and Hanrahan [2001] extend this work by showing that lighting can also be recovered under complex illumination if several photographs are used. Once both geometry and material properties of real objects are known, HGI can be achieved by conventional raytracing techniques. However, applying these approaches in general scenarios is often difficult, the parametric reflectance models used are not adequate for many materials, and real-time acquisition and inverse rendering for dynamic scenes is not possible.

Augmented Reality: Jacobs and Loscos [2006] provide a good overview of a variety of HGI techniques, that are usually referred to in the field as “mutual illumination” or “common illumination.” Some notable works are [Fournier et al. 1993; Debevec 1998; Gibson and Murta 2000], which introduce methods for achieving HGI by using an *a-priori* model of real object geometry and depth, respectively. Sato et al. [1999] achieve similar results by calculating real object geometry using stereo matching methods. Whether implicitly or explicitly, all these methods require shape and material properties of real objects in order to achieve a HGI solution.

Data Driven Rendering: Unger et al. [2003] capture 4D light fields from real-world objects and use them to relight synthetic objects. Masselus et al. [2003] and Sen et al. [2005] introduced image-based methods for relighting real-world scenes by capturing 4D and 6D datasets, respectively. While these methods enable global illumination effects for hybrid scenes, they only allow uni-directional radiance transfer, and hence are inadequate for a complete HGI solution. In principle, methods in Garg et al. [2006] can be used to achieve HGI because they capture the full 8D reflectance field of real objects. However, large datasets and long capture times prohibit this technique from being practical for real-time systems.

End-To-End Light Field Systems: Several systems have been developed in an attempt to provide end-to-end video conferencing that gives users the ability to choose their own viewpoint [Matusik and Pfister 2004; Yang et al. 2008]. These systems capture light fields of a scene, then transmit them to a light field projection system which is viewed by a user at a distant location. These systems are used for communication purposes and not for HGI rendering.

Adaptive Systems: A handful of systems have been proposed which utilize projectors to modify the appearance of real objects. Fujii et al. [2005] and Raskar et al. [2001] use a projector-camera system for photometric calibration and material editing, respectively, but neither approach can be used directly for HGI rendering.

3 Light Field Transfer

We now give a brief overview of our system. The next section discusses implementation of the various components in detail.

Physical and Virtual Layout: Figure 2 shows a diagram of the hardware layout. On the left are the real objects, light source and view camera. The view camera captures the images in which real and synthetic scenes are composited. A lens array and a camera/projector on the right provide the light field interface, that captures and projects light fields from/onto the real scene.

The renderer contains identical components in a virtual scene. It uses a set of projective lights and virtual cameras to project/capture light fields to/from the light field interface. The direct lighting is a point source, and a virtual camera captures an image for the final composite. Once a global coordinate system is chosen, the position and orientation of the direct lighting, view camera, and light field interface are calibrated to the virtual scene components.

Iterative Light Transfer: One full iteration of light transfer consists of the following steps. These can be repeated for multiple bounces of global illumination. **Step1:** The lens array from Fig. 2 multiplexes a 4D radiance field of the real scene into a 2D image that is captured by the light field camera. Lens arrays have been widely used in a similar fashion in previous light field systems. **Step2:** The captured light field is fed to the renderer, which converts the light field from a 2D image to a set of projective lights that illuminate the synthetic object. **Step3:** A set of virtual cameras are used to render a light field of the synthetic scene, illuminated by the captured light field. **Step4:** The light field from the synthetic scene is passed as a large 2D image to the projector (tiling images from each of the virtual cameras). It is converted to a 4D radiance pattern by the lens array, after which it illuminates the real object.

4 Implementation

We now discuss the various components of our system, in particular the hardware, calibration for acquisition, and rendering.

4.1 Hardware

Lens Array: A lens array is used to multiplex 4D light fields as a 2D image for capture and projection. The lens array consists of a hexagonal grid of Fresnel lenses and was purchased from Fresnel Tech (part #310). This lens array is a 8" x 10.5" sheet consisting of 10 x 12 lenses of approximately 1" diameter and focal length.

Capture: Two cameras are needed for light field transfer—a light field camera, and a view camera. Light field capture is done with a Lumenera Lw570 CCD camera, which is capable of capturing 2592x1944 images at 7 fps. In practice, 2048x1536 images are captured at slightly higher frame rate. View capture is performed with a Point Grey Firefly camera, capable of capturing 1024x768 images at 60Hz. An example of a 4D light field captured by this camera via the lens array is shown in Fig. 3(a).

Projection: Light field projection is accomplished with a 1024x768 pixel Epson digital projector. Each lens in the array covers an 80x80 pixel region, which results in greater angular resolution than spatial resolution. An example of a 4D light field projected by this system is shown in Fig. 3(b).




Figure 3: Light fields used in Fig. 5. (a) The 4D light field of the real scene that is used to relight the synthetic scene. The light field is undistorted and coerced onto a regular grid (as shown) before being input to the renderer. (b) The resulting 4D light field of the synthetic scene is projected onto the real scene. The light field is pre-distorted and coerced onto a hexagonal grid before projection.

4.2 Calibration

There are three aspects of calibration necessary for our system. The light field camera must be aligned to the lens array, the light field camera and projector must be aligned, and the real and virtual scenes must be aligned. All these calibrations are done only once, before the light field transfer system is used.

Lens Array Calibration: To calibrate the captured light field, a mapping of pixels in camera space into rays in a global coordinate system must be found. For this, we use a variation of the technique used in [Yang et al. 2008]. First, a homography between lens array coordinates and image coordinates is found. We assume each lenslet has identical intrinsic parameters, which allows us to calibrate each lenslet once we find its optical center. Since our lenslets have short focal length, and exhibit significant radial distortion, we use the calibration method in [Zhang 2000] to find a precise pixel-to-ray mapping. To register the lenslet optical centers, the light field of an on-axis light source located at infinity is captured. The bright regions in the resulting image correspond to on-axis rays captured by each lens in the array. We then use the Hough transform, to find the pixel locations for these on-axis rays.

Projector Calibration: To calibrate the light field projector and camera, a homography between the image space coordinates of each device is estimated. This is done by projecting a rectangle onto the lens array surface, capturing its image, and finding a homography between the captured and projected rectangles. Using this homography, light field images are transformed from projector space to camera space. Once in camera space, a mapping to global coordinates is known, described in the previous paragraph.

Global Coordinate Calibration: To align real and virtual scenes, we measure the location and orientation of a projected light field relative to our hardware. To achieve this, a synthetic light field of a checkered plane is rendered and projected onto the real scene. A diffuse surface is placed at the location of the projected light field, and its image is captured by the view camera. The extrinsic parameters of the checker pattern are found and used to place the synthetic objects in the renderer. Since direct lighting for our scenes was limited to a single point source, the relative position of this light was measured by hand and used to position a virtual light.

4.3 Rendering

For rendering, we must first develop a technique for image synthesis with a full light field as incident illumination. We describe a novel solution using projective light sources in graphics hardware. Beyond this, there are two rendering paths in our implementation—light field rendering and view rendering. Both rendering paths are necessary to achieve a final HGI rendering, but the image from the latter is the one used in the final composite. The result of light field rendering is used to relight the real scene.

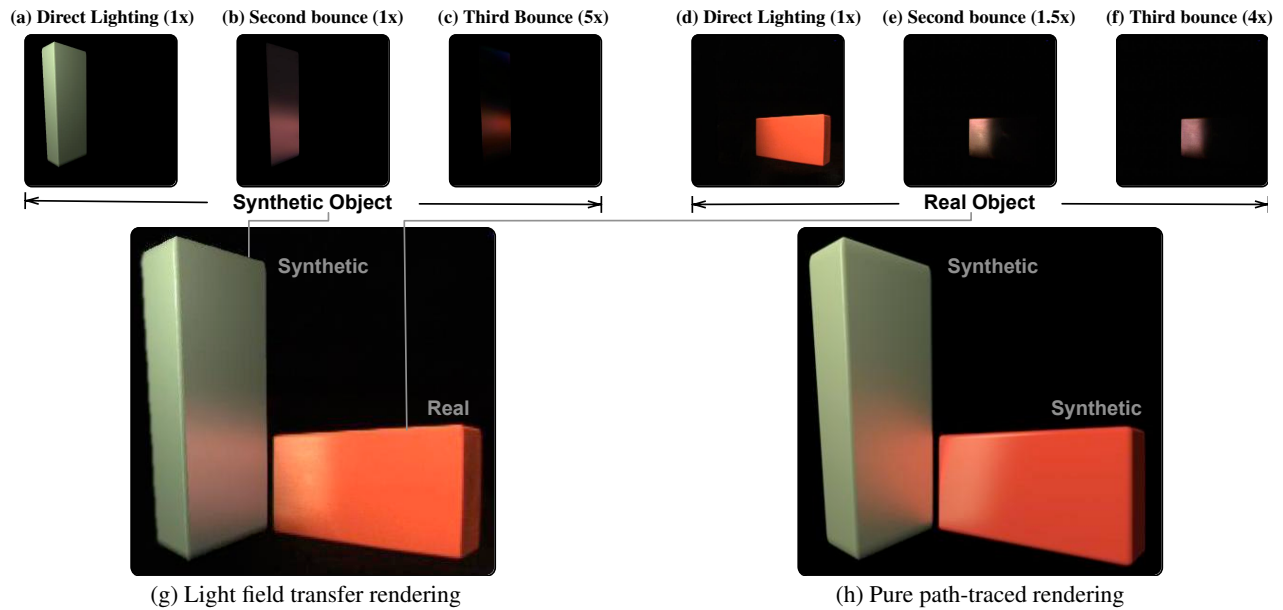


Figure 4: The light field transfer technique is used to decompose the bounces of light transport for a hybrid scene. (a) and (d) show direct illumination. (b) and (e) show contribution from only the first iteration of transfer. (c) and (f) show contribution from only the second iteration. (g) shows the composite image with the sum of contributions from the three stages of illumination. (h) shows a very similar computer generated scene rendered with full path tracing (all bounces included).

Rendering with Incident Light Fields: A significant challenge is how to illuminate the virtual scene with a full light field. Prior methods have used brute force illumination computations that take hours to days [Unger et al. 2003]. Instead, we create a set of projective light sources, with spatial locations and a field of view that matches the lens array. Relighting with light fields can then be implemented in graphics hardware. We use OpenGL, with GPU shader programs that take dynamic textures of captured light fields as input. The shaders calibrate the captured light fields, crop 2D slices for each source location, and project these slices onto the synthetic scene. The scene can then be rendered in the standard way, with illumination-BRDF computations in the pixel shader. The reflectances of the objects can be changed interactively. While our renderer does not compute full global illumination, it correctly illuminates convex objects. If necessary, greater flexibility in geometry can be attained using alternative real-time rendering algorithms, such as precomputed radiance transfer.

Rendering Output Light Fields: Light field rendering is accomplished as described in Sec. 3. The desired positions of synthetic objects in global coordinates are chosen. The relative position and orientation of direct lighting is registered based on the calibration. In addition, incident light field illumination is used for lighting the scene as described in the previous paragraph. A set of virtual cameras render different perspectives of the synthetic scene, and these images are tiled into a 2D image. The cameras have fields of view and spatial locations that are matched to the lens array. The light field is generated by a multiple pass OpenGL renderer that sequentially fills regions of the output frame-buffer. The output image is sent to the light field projector and transferred to the real scene.

View Rendering: A virtual camera is placed in the synthetic scene with the same location and orientation in global coordinates as the view camera. The relit synthetic objects are rendered from this perspective for use in the final compositing stage.

Final Composite: Compositing final HGI images was achieved by rendering an alpha layer in the previous stage. The images captured by the view camera and generated by view rendering are then composited using this alpha matte. This method assumes that real objects do not occlude synthetic objects. Clearly, if depth information from the real scene is available, depth comparisons will allow compositing scenes with arbitrary occlusion relationships.

Performance: Rendering is performed on an NVIDIA 8800GTX GPU. Video from the light field camera is streamed to the GPU for use as a texture. Despite the fact that high resolution data is transferred to the GPU, data transfer is not the bottleneck in rendering performance. Light field generation is the most expensive stage in the rendering pipeline, averaging 1-2 fps for most of the example scenes presented in this paper. Light field relighting follows with an average of 1-4fps, and video streaming at 7fps.

5 Results

We show a simple example, verifying that our system captures global illumination effects between real and synthetic objects. Then, we show several examples highlighting specific effects such as glossy reflections, soft shadows, and diffuse interreflections.

Multiple Bounce Global Illumination: Light field transfer enables multiple bounces of global illumination between real and synthetic scenes. Because our algorithm iteratively transfers indirect lighting, we can pause computation to inspect the contributions from individual iterations, and measure the contribution of each bounce, separately. By inspecting each bounce independently, we gain insights into the global effects in the scene and also determine when our algorithm converges. This has interesting (though loose) analogies with the direct-indirect separation of [Nayar et al. 2006], but we can now see each individual bounce separately.

Figure 4 shows the decomposition of lighting for a simple scene consisting of a synthetic green cube and a real red block. The figure shows the contribution to each object from direct, first bounce indirect, and second bounce indirect illumination. Notice that both real and synthetic objects mutually reflect onto each other. The final result, including a sum of direct and indirect contributions is also shown, along with an all-synthetic scene having similar material properties that is rendered with path tracing for comparison and verification. As expected, the contribution of the second bounce for indirect illumination is very small (see scale factors in the figure captions). For most scenes, a high quality result is obtained from just one or two iterations of our algorithm.

Soft Shadows and Diffuse-Diffuse Interactions: The example scene in Fig. 5 depicts shadowing and diffuse-diffuse interactions. Here, the virtual scene consists of a sun and an orbiting planet. The real scene consists of a mannequin head. There is no direct lighting

in this scene, and the only illumination of the mannequin comes from the synthetic sun. The synthetic planet casts a shadow on the mannequin, and some light reflected off the mannequin illuminates the shadowed side of the planet. As the planet moves closer to the mannequin in (b), the soft shadow becomes more distinct, and illumination from the mannequin on the planet increases.

Glossy Reflections: In Fig. 6, a hybrid scene consists of a virtual scene with a synthetic photograph and frame held by a hand above a table. The real scene contains a metallic bowl and statue placed on a checkered surface. A glossy reflection of the synthetic photograph is visible in the real bowl. The sequence shows more details in the reflected image as the frame moves closer to the bowl. This scene is also shown with a slightly different configuration in Fig. 1(b), where we also see glossy reflections of the real hand, orange, and bust in the synthetic photo frame.

6 Discussion on Limitations

Several straightforward improvements of our demonstration system could lend greater utility for practical applications. Restrictions on occlusion relationships can be relaxed with non-planar light field interfaces. Advances in imaging and projector technology will allow finer light field sampling and the use of shinier materials. By placing a light probe in the scene, environment mapping can be used to incorporate more complex direct lighting. Effective color calibration faces two significant challenges. First, the mapping of pixel intensities for camera and projector are usually non-linear and different from one another. Second, The RGB color filters for the camera and projector may have different pass-bands, introducing the possibility of cross-talk between projected and captured color channels. Calibration techniques described in [Grossberg et al. 2004] can be used to address both of these issues. Though these improvements would introduce greater rendering precision and flexibility in scene configurations, the following fundamental limitations may still preclude wider adoption:

- Synthetic objects cannot shadow illumination of real objects by real sources.
- Complete flexibility in material selection is complicated by the large number of samples required for mirror reflections.
- The method is *online*, and therefore requires synthetic models to be finished at the time of rendering.
- High dynamic range light field transfer is currently limited by digital projectors, which are only capable of 24-bit color.

7 Conclusions and Future Work

We have shown a method to achieve global illumination for hybrid scenes, that does not require any geometric or material knowledge of real objects. Despite the serious limitations on geometric configuration and illumination, our method is the first to provide consistent lighting between real and synthetic objects in near-real time. Our light field transfer algorithm also allows individual bounces of indirect lighting to be inspected and used to help minimize computation. These developments introduce new opportunities for visual effects and augmented reality applications.

In the future, it would be interesting to build a system for HGI rendering with complete human bodies. Such a system would require brighter projectors and custom lens arrays. It would provide more realistic lighting for scenes that include both real and synthetic actors. Another interesting configuration might be to create a smaller, portable system consisting of a cube, where each surface is capable of light field transfer. Synthetic objects would remain inside the cube as it moved about during a performance. Such a system would overcome some limitations in object placement, and could be realized using micro-projector technology.


There are a number of applications where consistent illumination is desired for real and virtual objects, that also interact and interreflect with each other. We have developed a light field transfer method as an important step towards this goal.

Acknowledgments


This work was supported in part by the NSF (CCF-04-46916, CCF-03-05322, CCF-07-01775, CCF-05-41259), a Sloan Research Fellowship, an ONR Young Investigator award N00014-07-1-0900, and a grant from the Columbia University Initiative in Science and Technology.

References


- ARNALDI, B., PUEYO, X., AND VILAPLANA, J. 1991. On the division of environments by virtual walls for radiosity computation. In *EGWR '91*, 198–205.
- DEBEVEC, P. 1998. Rendering synthetic objects into real scenes. In *SIGGRAPH '98*, 189–198.
- FOURNIER, A., GUNAWAN, A., AND ROMANZIN, C. 1993. Common illumination between real and computer generated scenes. In *Proc. of Graph. Interface*, 254–262.
- FUJII, K., GROSSBERG, M., AND NAYAR, S. 2005. A projector-camera system with real-time photometric adaptation for dynamic environments. In *CVPR '05*, 814–821.
- GARG, G., TALVALA, E., LEVOY, M., AND LENSCH, H. 2006. Symmetric photography: Exploiting data-sparseness in reflectance fields. In *EGSR '06*, 251–262.
- GIBSON, S., AND MURTA, A. 2000. Interactive rendering with real-world illumination. In *EGWR '00*, 365–376.
- GROSSBERG, M., PERI, H., NAYAR, S., AND BELHUMEUR, P. 2004. Making one object look like another: controlling appearance using a projector-camera system. In *CVPR '04*, 452–459.
- JACOBS, K., AND LOSCOS, C. 2006. Classification of illumination methods for mixed reality. *Comp. Graph. Forum* 25, 1, 29–51.
- MASSELUS, V., PEERS, P., DUTRÉ, P., AND WILLEMS, Y. 2003. Relighting with 4D incident light fields. In *SIGGRAPH '03*, 613–620.
- MATUSIK, W., AND PFISTER, H. 2004. 3D TV: a scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes. In *SIGGRAPH '04*, 814–824.
- NAYAR, S., KRISHNAN, G., GROSSBERG, M., AND RASKAR, R. 2006. Fast separation of direct and global components of a scene using high frequency illumination. In *SIGGRAPH '06*, 935–944.
- RAMAMOORTHY, R., AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *SIGGRAPH '01*, 117–128.
- RASKAR, R., WELCH, G., LOW, K.-L., AND BANDYOPADHYAY, D. 2001. Shader Lamps: Animating real objects with image based illumination. In *EGWR '01*, 89–101.
- SATO, I., SATO, Y., AND IKEUCHI, K. 1999. Acquiring a radiance distribution to superimpose virtual objects onto a real scene. *IEEE Trans. on Vis. and Comp. Graph.* 5, 1, 1–12.
- SEN, P., CHEN, B., GARG, G., MARSCHNER, S., HOROWITZ, M., LEVOY, M., AND LENSCH, H. 2005. Dual photography. In *SIGGRAPH '05*, 745–755.
- UNGER, J., WENGER, A., HAWKINS, T., GARDNER, A., AND DEBEVEC, P. 2003. Capturing and rendering with incident light fields. In *EGSR '03*, 141–149.
- YANG, R., HUANG, X., LI, S., AND JAYNES, C. 2008. Toward the light field display: Autostereoscopic rendering via a cluster of projectors. *IEEE Trans. on Vis. and Comp. Graph.* 14, 1, 84–96.
- YU, Y., DEBEVEC, P., MALIK, J., AND HAWKINS, T. 1999. Inverse global illumination: recovering reflectance models of real scenes from photographs. In *SIGGRAPH '99*, 215–224.
- ZHANG, Z. 2000. A flexible new technique for camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell.* 22, 11, 1330–1334.



(a) The planet is far from the mannequin.




(b) The planet moves closer.




(c) The planet is in front of the nose.


Figure 5: In these figures, a real mannequin head is illuminated by a synthetic sun, and a synthetic planet casts a shadow on the mannequin. As the planet approaches the mannequin in (b), the shadow becomes more distinct and the backside of the planet becomes illuminated by the mannequin. (Please see submitted video.)



(a) The photo frame is far from the shiny bowl.



(b) The photo frame moves closer.



(c) The photo frame is very close to the shiny bowl.

Figure 6: Examples showing glossy reflections of a synthetic photograph on a shiny bowl in the real scene. As the frame moves closer, its reflection enlarges and more detail is visible. (Please see submitted video.)