

Consistent Interactive Augmentation of Live Camera Images with Correct Near-field Illumination

Thorsten Grosch*
MPI Informatik
Computer Graphics Group
Saarbruecken, Germany

Tobias Eble†
Institute for Computational Visualistics
University of Koblenz-Landau, Germany

Stefan Mueller‡
Institute for Computational Visualistics
University of Koblenz-Landau, Germany

Abstract

Inserting virtual objects in real camera images with correct lighting is an active area of research. Current methods use a high dynamic range camera with a fish-eye lens to capture the incoming illumination. The main problem with this approach is the limitation to distant illumination. Therefore, the focus of our work is a real-time description of both near - and far-field illumination for interactive movement of virtual objects in the camera image of a real room. The daylight, which is coming in through the windows, produces a spatially varying distribution of indirect light in the room; therefore a near-field description of incoming light is necessary. Our approach is to measure the daylight from outside and to simulate the resulting indirect light in the room. To accomplish this, we develop a special dynamic form of the irradiance volume for real-time updates of indirect light in the room and combine this with importance sampling and shadow maps for light from outside. This separation allows object movements with interactive frame rates (10 - 17 fps). To verify the correctness of our approach, we compare images of synthetic objects with real objects.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual Reality; Radiosity; Color, Shading, Shadowing and Texture

Keywords: Augmented Image Synthesis, Global Illumination

1 Introduction

Augmenting real camera images with virtual objects has many applications in the movie industry, cultural heritage and architecture, like augmented building sites or virtual furniture in real rooms. Recent hardware developments allow real-time capturing of high dynamic range (HDR) images of the current environmental light which is necessary for a consistent illumination of the virtual objects. The assumption often used for natural illumination is distant lighting, which is acceptable for outdoor applications but fails for an indoor scenario, where a correct near-field description of the

light is necessary because of the spatially varying indirect light. Our work is therefore focussed on interactive augmentation of camera images, showing a real room under time-varying daylight. We use a HDR video camera (IMS Chips) with a fish-eye lens which captures the daylight outside the room and simulate the resulting indirect light in the room on-the-fly. The display of the virtual object is a combination of graphics hardware for direct daylight and a time-varying irradiance volume for color bleeding effects.



Figure 1: Various virtual objects are placed in a real Cornell Box with arbitrary, time-varying illumination. Both direct and spatially varying indirect light can be displayed at interactive frame rates. The bottom row shows direct lighting, indirect lighting and the combination of both.

Our main contributions are:

- We develop a system for interactive illumination of virtual objects in real camera images which enables the display of objects with an arbitrary far-field and a diffuse near-field illumination.

*e-mail: grosch@uni-koblenz.de

†e-mail: tobias.eble@uni-koblenz.de

‡e-mail: stefanm@uni-koblenz.de

- We present a dynamic version of the irradiance volume which automatically adapts to temporally varying incident light
- For direct light, we show how to draw samples in the environment map which contain only the necessary illumination that is falling through the real windows onto the virtual object.

The rest of the paper is organized as follows: In Section 2 we review existing work and in Section 3 we explain our approach. The necessary preparations are described in Section 4. Section 5 explains the precomputations and Section 6 shows how to use them for interactive illumination. The sampling for direct illumination is described in Section 7 before we demonstrate our results in Section 8 and conclude in Section 9.

2 Previous Work

Inserting virtual objects with consistent illumination in real photographs was first shown by [Nakamae et al. 1986] in the form of virtual buildings. [Fournier et al. 1993] invented the differential rendering technique to display a virtual object in a real photograph based on two lighting simulations. This method was extended by [Debevec 1998] for natural illumination captured in the HDR image of a light probe. [Sato et al. 1999] developed a similar technique based on a fish-eye image. Several improvements exist for moving objects [Drettakis et al. 1997], final gathering [Loscos et al. 1999], inclusion of light goniometrics [Grosch et al. 2003] and differential photon mapping [Grosch 2005a]. A real-time augmentation of photographs, implemented on graphics hardware, was invented by [Gibson and Murta 2000][Gibson et al. 2003], and extended to panoramic images by [Grosch 2005b].

Several techniques for the sampling of environment maps exist for fast rendering of objects in natural illumination (see [Reinhard et al. 2005] for an overview). Another way for real-time illumination in a natural environment is based on environment map compression [Sloan et al. 2002], using spherical harmonics [Ramamoorthi and Hanrahan 2001] or wavelets [Ng et al. 2003]. Both approaches typically assume distant lighting and an extension to spatially varying light is difficult: Localized low-frequency lighting is possible by using spherical harmonic gradients [Annen et al. 2004]. A few methods include indirect lighting from point and spot lights [Hasan et al. 2006][Kristensen et al. 2005][Kontkanen et al. 2006].

We present a system which is able to display virtual objects under temporally and spatially varying light at interactive frame rates. Our system is based on the work of [Havran et al. 2005], who used an HDR video camera with a fish-eye lens permanently capturing the illumination. Virtual objects with correct illumination and shadows can be displayed at interactive rates using a multi-pass shadow mapping [Williams 1978] method in combination with temporal filtering of the sampling points. The main drawback is the limitation to distant light, which was improved by [Korn et al. 2006] with an HDR stereo camera. Here, the virtual object was displayed in a real photograph.

Beside the illumination information, information about the camera pose must be recovered when integrating a virtual object in a real image (tracking). Only a few approaches try to insert virtual objects with correct illumination in moving images. The first work was presented by [Kanbara and Yokoya 2002] by tracking a light probe. [Agusanto et al. 2003] used the image of a light probe and assumed static lighting. [Haller et al. 2003] showed shadows of virtual objects on real objects by using shadow volumes of manually placed point lights.

The tracking in our work is based on the marker tracking of AR-Toolkit [Kato and Billinghurst 1999]. We use two HDR cameras:

One camera is looking at the scene while the other camera is recording the environment illumination.

3 Our Approach

Like in [Havran et al. 2005], we create a set of sampling points for the current image of the fish-eye camera and illuminate the virtual object with a collection of shadow maps. Using this method inside a room is difficult because of the spatially varying illumination and the appearance of a virtual object will depend on the fish-eye camera position instead of its own object position. For example, a virtual object will be bright if the fish-eye camera is placed in a region where the sunlight is coming through a window and it will be dark if the camera is located in a shadow region. To avoid placing many cameras at different locations in the room, we developed a dynamic version of the irradiance volume [Greger et al. 1998] which adapts to the daylight coming in through the windows. The irradiance volume contains only indirect light which is reflected from the diffuse walls and objects in the room. The direct daylight is measured with a HDR fish-eye camera, located *outside* the room and processed separately with the graphics hardware in multiple render passes with a set of shadow maps. The interpolation error introduced by the discretization of the irradiance volume is therefore acceptable, because only slowly varying reflected indirect light is used and this can be well approximated with a coarse grid. Shadow leaking artifacts, as described in [Kontkanen and Laine 2006] and [Grosch 2005b], are mainly avoided due to the separation of direct and indirect light. To reduce the large memory requirements of the irradiance volume, spherical harmonics with nine coefficients can be used to describe the incoming indirect illumination at each grid point, as noticed by various authors [Mantiuk et al. 2002][Nijasure et al. 2005][Gibson et al. 2003]. Our idea is visualized in Figure 2: An HDR camera is recording the daylight from outside (we assume distant lighting here) and the indirect lighting in the room is computed from a dynamic irradiance volume, adjusted to the direct lighting. Although an analytic description of sun and sky is possible [Perez et al. 1993], we cannot simulate the current turbidity of the sky, especially the moving clouds. We therefore decided to use a fish-eye camera to observe the whole daylight from outside, including the light which is reflected from the environment.




Figure 2: Basic idea: The illumination of a virtual object is separated into two components: Direct light coming from outside and indirect light within the room. An HDR video camera is recording the daylight from outside and the resulting indirect light in the room is simulated.

4 Preparations

We first generate a 3D model of our scene manually with computer vision techniques from images [Hartley and Zisserman 2004]. All diffuse materials are reconstructed from the HDR images of two cameras, as described in [Ritschel and Grosch 2006]. The dis-

tortion of the fish-eye lens is computed from a checkerboard pattern. For photometric calibration, we recorded a MacBeth Color Checker with different illumination, measured the radiance of all color patches and fitted a logarithmic curve for the mapping between radiance and pixel color, as described in [Hoefflinger 2007]. The transfer function, required for spherical harmonics illumination, is computed and stored for all vertices of the virtual object.

5 Precomputation of Basis Illuminations

Our approach is based on the *superposition* principle of light. In [Nimeroff et al. 1994], an image resulting from arbitrary daylight illumination is composed from a set of pre-computed basis images. We extend this idea to the creation of an arbitrary irradiance volume which is constructed from a set of pre-computed basis irradiance volumes.

5.1 Basis Irradiance Volumes

The daylight is represented by a hemisphere around the scene which is subdivided in a user-defined number of N distinct regions. For each region, we compute a lighting simulation with an area light emitting 1 cd/m^2 , placed at this region. By limiting our approach to diffuse materials, a radiosity simulation [Sillion and Puech 1994] can be performed for this task. Next, we compute a *basis irradiance volume* inside the room for the resulting radiance distribution. The emitter is excluded from this irradiance volume, only the *indirect* light in the room is collected. This process is repeated for all N regions of the hemisphere, resulting in N basis irradiance volumes. Figure 3 visualizes this process. The irradiance volume for a given illumination can be computed from these basis irradiance volumes, as described in the next section. Note that it is not necessary to store all radiosity simulations, only the basis irradiance volumes are required for later use.

5.2 Spherical Harmonics Compression

To reduce memory requirements, the information of each grid point of a (basis) irradiance volume is described in form of spherical harmonics. Therefore, the radiance distribution seen from a grid point is projected onto the spherical harmonics basis functions to obtain the coefficients for the incident radiance. Ramamoorthi and Hanrahan showed, that only the first nine coefficients are required for unshadowed diffuse illumination with negligible error. This means, that spherical harmonics are well suited to describe the diffuse reflected indirect light in a room at each spatial location with only small memory consumption. For the remainder, we still refer to this grid data structure as *irradiance volume*, although it contains the directional radiance information in spherical harmonics representation.

Using the notation of [Ramamoorthi and Hanrahan 2001], the incoming radiance L_{in} at position \mathbf{x} , viewing in direction ω , is described as a sum of basis functions:

$$L_{in}(\mathbf{x}, \omega) = \sum_{l=0}^{\infty} \sum_{m=-l}^l L_{lm}(\mathbf{x}) Y_{lm}(\omega) \quad (1)$$

with the coefficients

$$L_{lm}(\mathbf{x}) = \int_{\Omega} L_{in}(\mathbf{x}, \omega) Y_{lm}(\omega) d\omega \quad (2)$$




Figure 3: Precomputation of basis illumination: The (hemi)sphere around the scene is subdivided into a set of N regions (left) and a radiosity simulation is computed for each emitter (center). For each of these basis simulations, an irradiance volume is computed from the reflected light (right).

where $Y_{lm}(\omega)$ are the orthonormal spherical harmonics basis functions. We use the superscript (i) to describe quantities from the basis simulations, eg. $L_{lm}^{(i)}(\mathbf{x})$ are the coefficients at grid point \mathbf{x} for the basis illumination of emitter i .

6 Combination of the Basis Illuminations

For clarity, we first describe the combination of the patch radiances of the basis simulations (Section 6.1) for arbitrary environment light. This is not part of our algorithm, but it is the foundation for the combination of the basis irradiance volumes, which is described in Section 6.2.

6.1 Combining the Patch Radiances

Due to the *linearity* of light, we can combine the current illumination from the precomputed basis illuminations. If we replace the emitter radiance from 1 cd/m^2 to an arbitrary radiance L_i , all patch radiances from the basis simulations $L_{patch}^{(i)}$ will increase by the same scaling factor to $L_{patch}^{(i)} \cdot L_i$. When using multiple emitters, the resulting illumination is identical to the sum of the individual simulations, because light is *additive*. For N emitters with emitter radiances L_1, \dots, L_N , the resulting patch radiances are therefore a linear combination of the basis patch radiances

$$L_{patch} = \sum_{i=1}^N L_{patch}^{(i)} \cdot L_i \quad (3)$$

6.2 Combining the Irradiance Volumes

Applying Equation 2 for computing the spherical harmonics coefficients for this radiance distribution leads to:

$$\begin{aligned}
 L_{lm}(\mathbf{x}) &= \int_{\Omega} L_{in}(\mathbf{x}, \omega) Y_{lm}(\omega) d\omega \\
 &= \int_{\Omega} \left(\sum_{i=1}^N L_{in}^{(i)}(\mathbf{x}, \omega) \cdot L_i \right) Y_{lm}(\omega) d\omega \\
 &= \sum_{i=1}^N \int_{\Omega} L_{in}^{(i)}(\mathbf{x}, \omega) \cdot L_i \cdot Y_{lm}(\omega) d\omega \quad (4) \\
 &= \sum_{i=1}^N L_i \cdot \int_{\Omega} L_{in}^{(i)}(\mathbf{x}, \omega) \cdot Y_{lm}(\omega) d\omega \\
 &= \sum_{i=1}^N L_i \cdot L_{lm}^{(i)}(\mathbf{x})
 \end{aligned}$$

In this derivation, we describe the incoming radiance L_{in} as a linear combination of the basis radiances $L_{in}^{(i)}$ (step 1). The integral can now be written as a sum of integrals (step 2) and the constant factors L_i can be written in front of the integral (step 3). Now, the remaining integrals are identical to the original coefficients $L_{lm}^{(i)}$ of the basis irradiance volumes (step 4).

This means, that the spherical harmonics coefficients for the current illumination are a linear combination of the original coefficients of the basis irradiance volumes. The scaling factors L_i can be computed by averaging the pixels of region i in the fish-eye image (see Figure 4).

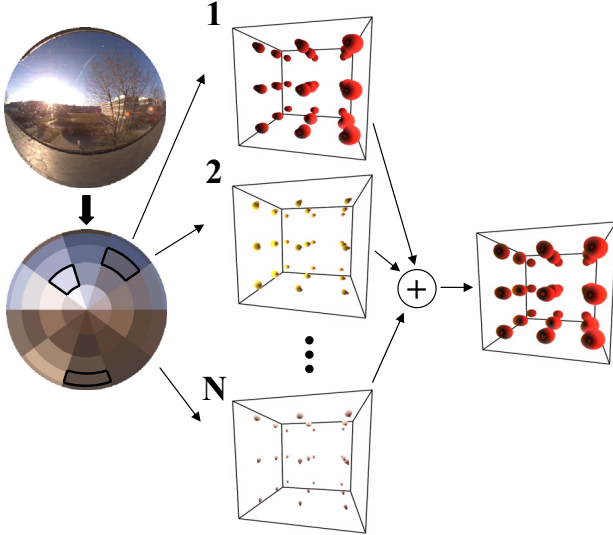


Figure 4: Combining basis illumination: For each region, we take the average of the pixels in the fish-eye image (left). The resulting average radiance values are the scaling factors for the basis irradiance volumes (center). The final irradiance volume is a linear combination of the basis irradiance volumes (right).

6.3 Displaying the Indirect Light

To obtain the indirect radiance of a vertex v at position \mathbf{x} , only a dot product of the coefficient vectors has to be computed [Sloan et al. 2002]:

$$L_{indirect}(\mathbf{x}) = \frac{\rho}{\pi} \sum_{l=0}^{\infty} \sum_{m=-l}^l L_{lm}(\mathbf{x}) C_{lm}^v \quad (5)$$

where C_{lm}^v are the coefficients of the transfer function at the vertex. The coefficients L_{lm} , representing the indirect light, are interpolated between grid vertices. Note that it is not necessary to compute Equation 4 for all grid vertices of the irradiance volume. Only the coefficients L_{lm} for the grid points inside the bounding box of the virtual object are computed, because they are required for interpolation.

7 Direct Daylight

Because direct daylight is more important than indirect lighting in the room, we do not use the irradiance volume for this type of light. Direct lighting causes sharp shadow boundaries and the interpolation errors would be too significant. To evaluate the direct light we have to compute

$$L_{direct}(\mathbf{x}) = \frac{\rho}{\pi} \int_{\Omega} L_{env}(\omega) V(\mathbf{x}, \omega) \cos(\theta) d\omega \quad (6)$$

where $L_{env}(\omega)$ is the direct light visible in the fish-eye image and $V(\mathbf{x}, \omega)$ is the binary visibility term which is zero, if the ray, starting from \mathbf{x} is blocked by a geometric object in direction ω , and one otherwise. To display accurate shadows, direct light is computed by taking M sampling points in the fish-eye image (importance sampling):

$$L_{direct}(\mathbf{x}) \approx \frac{\rho}{M\pi} \sum_{j=1}^M \frac{L_{env}(\omega_j) V(\mathbf{x}, \omega_j) \cos(\theta_j)}{p(\mathbf{x}, \omega_j)} \quad (7)$$

based on a density function $p(\mathbf{x}, \omega_j)$.

7.1 Sampling Strategy for Rooms with Windows

A simple choice to solve Equation 7 is to use a density $p \sim L_{env}$. While this works well if the virtual object is directly illuminated by the sun, the sampling can still be improved if the object is in shadow (see Figure 5). In this case, the window limits the solid angle of incoming direct light, and only daylight from this cone of directions is arriving at the virtual object, producing a soft shadow. To increase the image quality, sampling should be restricted to this region instead of wasting many samples for the sun and other invisible regions.

To draw samples with a density $p \sim L_{env} \cdot V$, we use four planes to describe the visible daylight region. For each of the four window edges, we compute a tangential plane at the bounding sphere of the virtual object, which is passing through the window edge. Two of these planes are visualized in Figure 5. Before drawing samples, we perform an *in-out* test for each pixel of the fish-eye image. If the corresponding point on the hemisphere is located outside, the pixel is set to black. Samples are drawn from the resulting image, which contains only the portion of daylight that illuminates the virtual object. Figure 6 shows a comparison of our sampling technique with sampling of the whole fish-eye image.




Figure 5: Depending on the position of the virtual object, only a fraction of the incoming illumination has to be considered for direct lighting.

In case of multiple windows, this process can be repeated for each window, resulting in a fish-eye image with multiple window regions. Because we use the bounding sphere of the virtual object, this is a conservative approximation of the visible region and no daylight is missing.

7.2 Multi-Pass Rendering

The calculation of direct daylight is based on the work of [Gibson and Murta 2000]: For each sampling point, illumination from one pointlight with one shadow map is rendered. The total illumination is computed by accumulation of the single images. To avoid darkening of existing shadows, two binary shadow images are generated for each point light: One with the virtual object and a second image with only the real geometry. Subtracting these two images results in a binary shadow image which contains only *new* shadows caused by the virtual object. The missing direct illumination is computed only in these regions and subtracted from the real image in each render pass. In addition, the planes described in Section 7.1 are used as *clipping planes* during rendering to avoid drawing any fragments outside the possible shadow region.

8 Results

As a representative test scene for a room, we built a small Cornell Box together with a moving light, representing time-varying daylight. An HDR video camera with a fish-eye lens is placed on top of the box, permanently recording the changing illumination. Markers are placed beside the box for tracking. In Figure 7 we show a moving virtual object which is placed at several locations inside the box. Note the consistent illumination including color bleeding from the walls which are not visible in the fish-eye image.

To verify the accuracy of our indirect light approximation, we placed virtual objects in a synthetic image (see Figure 8). In the left image, we used our technique to place three virtual teapots in the image of a room, rendered with *pbrt* [Pharr and Humphreys 2004]. In the right image, the whole scene was rendered with *pbrt*, showing only small differences.

A comparison with a real object is shown in Figure 9: Here, we placed a *real* teapot, generated with a 3D printer, into a Cornell Box. Although small differences in indirect lighting are visible, we believe that our approach is a good tradeoff between speed and quality. In an interactive application, it is important not to ignore the indirect light (see Figure 1). Small errors in direct light (e.g. wrong light direction and shadows) are more significant than hardly noticeable inaccuracies in indirect light.

Object movement is possible with interactive frame rates (10 - 17 fps), a detailed description is shown in Table 1. For still images, we achieve even real-time frame rates (22 - 116 fps). The additional time for moving images is caused by marker extraction, pose estimation and for uploading new camera images to the graphics card. For all measurements, we used 64 shadow maps for the direct light, stored in a single 512 x 512 texture, with 64 x 64 pixels for each shadow map. The image resolution is 512 x 496 (which is the maximum resolution for two HDR cameras).

We observed that a number of $N = 64$ basis emitters is sufficient for a good approximation of indirect light. The precomputation time for 64 basis radiosity simulations is about five minutes (2K patches), computing the coefficients for the basis irradiance volumes takes 2-3 seconds (for 8^3 grid vertices). The computation of 64 scaling values L_i requires 5 ms, the time for indirect light calculation of the virtual object is 0.13 ms (teapot) and 0.29 ms (bunny). Due to the spherical harmonics compression, the whole irradiance volume requires only 3.5 MB of memory. A visualization of the irradiance volume is shown in Figure 10 for different sun directions.

All measurements were made on an Intel Core 2 Duo, 2.13 GHz, 1GB RAM, equipped with an NVIDIA 8800 GTS graphics card.

Average Frame Rate	Standard Sampling	Window Sampling
Teapot (4K faces)	76..116 15..17	42..64 15..17
Bunny (15K faces)	22..40 10..12	22..35 10..12

Table 1: Frame rates for virtual objects of different size. In the first column, standard importance sampling is used; in the second column, window exclusion is activated. For still images, real-time frame rates are achieved (first values), in case of moving images, the frame rates are interactive (second values).




Figure 8: Indirect light approximation with a synthetic scene: The left image shows three virtual teapots, placed in a *pbrt* image of a room. In the right image, the whole scene was rendered with *pbrt*. Note the similar appearance of all teapots.

9 Conclusion and Future Work

In this paper, we presented a method for interactive illumination of virtual objects, which are placed in real camera images of rooms under time-varying lighting. For consistent illumination, both far- and diffuse near-field illumination is computed. For a correct near-field description we developed a dynamic version of the irradiance volume which automatically adapts to the incoming far-field illumination in real-time. The far-field illumination is measured with an HDR fish-eye video camera, placed outside the room. This distant




Figure 6: Illumination for a Cornell Box with a window. Top row: Original fish-eye image, excluded window region and final image used for sampling. Bottom row: Direct light computed with 64 samples from the whole fish-eye image (left), rejection sampling (center) and sampling only from the excluded region (right). Without the exclusion, many samples are used for the sun and only a few samples are useful. Rejection sampling introduces more flickering than sampling from the excluded window region, as can be seen in the accompanying video.




Figure 7: Application example: A user is walking with a camera in a real room with varying illumination. The camera images show augmentations with consistent illumination in all locations of the room. The room is represented by a small Cornell Box, position and orientation of the camera are obtained by marker tracking.




Figure 9: Comparison virtual - real: The left image shows a real Cornell Box, augmented with a virtual teapot. In the right image, the teapot is real (generated with a 3D printer).

illumination is displayed with the graphics hardware. A number of samples is taken from the fish-eye image and a multi-pass algorithm with shadow maps is used for display. To obtain a direct illumination with high quality, we developed a special sampling strategy to select only samples which correspond to incoming light directions that illuminate the virtual object through the windows of the room.

We obtain interactive frame rates for our testscenes and a similar appearance of virtual objects in comparison with real objects.

Currently, our system is restricted to diffuse materials and no indirect shadows can be displayed. Due to the limitation of the irradiance volume, no changes in indirect lighting, introduced by the virtual object, e.g. color bleeding from virtual to real objects, can be displayed.

Some images show colored noise (e.g. Figure 6 and 7), which becomes visible after a color saturation required for displaying images from the HDR camera. Alternatively, a low dynamic range (LDR) camera with lower noise and higher resolution could be used for viewing the virtual object. But in this case a color transformation between the camera images would be required. Different camera sensors with different spectral sensitivities, a time-varying white balance and saturated pixels must be considered here. We leave this topic as future work, for now we remove the camera noise by averaging a few successive frames (for still images). A higher resolution of the fish-eye camera would also be desirable, because the size of the sun is only a few pixels which introduces a slight color shift, because the color image is computed from a Bayer pattern.

Many research directions exist for future work: Direct illumination can be improved by using more sophisticated sampling, e.g. Quasi




Figure 10: Visualization of the irradiance volume for different light directions. Note that the irradiance volume adapts to the change in illumination, the update time for an 8^3 irradiance volume is 5 ms. For clarity, only 3^3 grid points are shown.

Monte Carlo methods [Pharr and Humphreys 2004] and filtering samples over time [Havran et al. 2005]. Moreover, we investigate if the inclusion of real light sources [Goesele et al. 2003] is possible in our simulation to remove the restriction to daylight. The fixed number N of constant basis functions we use for the creation of the basis simulations could be replaced by a hierarchical subdivision of the sphere in combination with an oracle function. One drawback of the irradiance volume is that the illumination of the scene remains unaffected by object motions. Therefore, an interesting topic is the adaptation of the irradiance volume to geometric changes, either virtual or real, as demonstrated in [Nijasure et al. 2005]. Virtual objects with glossy materials could be included with a final gathering approach, as shown in [Dmitriev et al. 2004].

The integration of our technique in a mobile AR system would also be interesting but there are several problems: First, the HDR cameras require special frame grabber hardware, a connection to a laptop is currently not possible. Secondly, displaying virtual objects in video-see-through glasses introduces the same problems like LDR cameras. Using optical see-through devices is even more difficult because an HDR display would be required for the virtual pixels.

We plan to test our augmentation method with images of real rooms. Our vision is, that a user with a head-mounted display can walk inside real rooms, and augmented images, indistinguishable from the real world, are shown in the glasses. For multiple rooms and windows, several HDR cameras could be placed at different windows, or alternatively a single camera on the roof of the building, directed upwards. For this single-camera solution, the missing illumination from the lower hemisphere (e.g. light reflected into the room from the ground outside) could be estimated from the known irradiance of the upper hemisphere.

References

AGUSANTO, K., LI, L., CHUANGUI, Z., AND SING, N. W. 2003.

$L_{in}(\mathbf{x}, \omega)$	Incoming radiance seen from point \mathbf{x} in direction ω
$L_{in}^{(i)}(\mathbf{x}, \omega)$	Incoming radiance seen from point \mathbf{x} in direction ω in basis simulation i
$Y_{lm}(\omega)$	Spherical harmonics basis function
$L_{lm}(\mathbf{x})$	Spherical harmonics coefficient for incoming light at gridpoint \mathbf{x}
$L_{lm}^{(i)}(\mathbf{x})$	Spherical harmonics coefficient for incoming light at gridpoint \mathbf{x} for basis simulation i
C_{lm}^v	Spherical harmonics coefficient for transfer function at vertex v
L_i	Radiance of emitter i
L_{patch}	Patch radiance value
$L_{patch}^{(i)}$	Patch radiance value in basis simulation i
$L_{env}(\omega)$	Radiance value in environment map
$L_{direct}(\mathbf{x})$	Outgoing radiance at point \mathbf{x} , originating from direct light
$L_{indirect}(\mathbf{x})$	Outgoing radiance at point \mathbf{x} , originating from indirect light

Figure 11: Notation

Photorealistic rendering for augmented reality using environment illumination. In *ISMAR*, 208–216.

ANNEN, T., KAUTZ, J., DURAND, F., AND SEIDEL, H.-P. 2004. Spherical Harmonic Gradients for Mid-Range Illumination. In *Proceedings of the Eurographics Symposium on Rendering*, 331–336.

DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proc. SIGGRAPH '98*, 189–198.

DMITRIEV, K., ANNEN, T., KRAWCZYK, G., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2004. A cave system for interactive modeling of global illumination in car interior. In *ACM Symposium on Virtual Reality Software and Technology (VRST 2004)*, ACM, Hong Kong, R. Lau and G. Baciuc, Eds., 137–145.

DRETTAKIS, G., ROBERT, L., AND BOUGNOUX, S. 1997. Interactive common illumination for computer augmented reality. In *Eighth Eurographics Workshop on Rendering*, 45–56.

FOURNIER, A., GUNAVAN, A., AND ROMANZIN, C. 1993. Common illumination between real and computer generated scenes. In *Proc. of Graphics Interface 1993*, 254–262.

GIBSON, S., AND HOWARD, T. 2000. Interactive reconstruction of virtual environments from photographs, with application to scene-of-crime analysis. In *Proc. of ACM Symposium on Virtual Reality Software and Technology 2000, Seoul, Korea, October 2000*, 41–48.

GIBSON, S., AND MURTA, A. 2000. Interactive rendering with real-world illumination. In *Eleventh Eurographics Workshop on Rendering*, 365–376.

GIBSON, S., COOK, J., HOWARD, T., AND HUBBOLD, R. 2003. Rapid shadow generation in real-world lighting environments. In *Eurographics Symposium on Rendering*, 219–229.

GOESELE, M., GRANIER, X., HEIDRICH, W., AND SEIDEL, H.-P. 2003. Accurate light source acquisition and rendering. In

- SIGGRAPH '03: *ACM SIGGRAPH 2003 Papers*, ACM Press, New York, NY, USA, 621–630.
- GREGER, G., SHIRLEY, P., HUBBARD, P. M., AND GREENBERG, D. P. 1998. The irradiance volume. *IEEE Computer Graphics and Applications* 18, 2, 32–43.
- GROSCHE, T., MUELLER, S., AND KRESSE, W. 2003. Goniometric light reconstruction for augmented reality image synthesis. In *Graphiktag im Rahmen der GI Jahrestagung, Frankfurt am Main*.
- GROSCHE, T. 2005. Differential photon mapping: Consistent augmentation of photographs with correction of all light paths. In *Eurographics 2005, Trinity College, Dublin, Ireland*.
- GROSCHE, T. 2005. PanoAR: Interactive augmentation of omnidirectional images with consistent lighting. In *Mirage 2005, Computer Vision / Computer Graphics Collaboration Techniques and Applications*, 25–34.
- HALLER, M., DRAB, S., AND HARTMANN, W. 2003. A real-time shadow approach for an augmented reality application using shadow volumes. In *VRST '03: Proceedings of the ACM symposium on Virtual reality software and technology*, ACM Press, New York, NY, USA, 56–65.
- HARTLEY, R., AND ZISSERMAN, A. 2004. *Multiple View Geometry in Computer Vision*. Cambridge University Press.
- HASAN, M., PELLACINI, F., AND BALA, K. 2006. Direct-to-Indirect Transfer for Cinematic Relighting. *ACM Trans. Graph.* 25, 3, 1089–1097.
- HAVRAN, V., SMYK, M., KRAWCZYK, G., MYSKOWSKI, K., AND SEIDEL, H.-P. 2005. Importance Sampling for Video Environment Maps. In *Eurographics Symposium on Rendering 2005*, ACM SIGGRAPH, Konstanz, Germany, K. Bala and P. Dutré, Eds., 31–42.
- HOEFFLINGER, B. 2007. *High-Dynamic-Range (HDR) Vision: Microelectronics, Image Processing, Computer Graphics (Springer Series in Advanced Microelectronics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- KANBARA, M., AND YOKOYA, N. 2002. Geometric and photometric registration for real-time augmented reality. In *IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR) (September 2002)*, 279.
- KANBARA, M., AND YOKOYA, N. 2004. Real-time estimation of light source environment for photorealistic augmented reality. In *ICPR (2)*, 911–914.
- KATO, H., AND BILLINGHURST, M. 1999. Marker tracking and HMD calibration for a video-based augmented reality conferencing system. In *Proceedings of the 2nd International Workshop on Augmented Reality (IWAR 99)*.
- KONTKANEN, J., AND LAINE, S. 2006. Sampling precomputed volumetric lighting. *Journal of Graphics Tools* 11, 3.
- KONTKANEN, J., TURQUIN, E., HOLZSCHUCH, N., AND SILLION, F. X. 2006. Wavelet Radiance Transport for Interactive Indirect Lighting. In *17th Eurographics Symposium on Rendering*, 161–172.
- KORN, M., STANGE, M., VON ARB, A., BLUM, L., KREIL, M., KUNZE, K., ANHENN, J., WALLRATH, T., AND GROSCHE, T. 2006. Interactive augmentation of live images using a HDR stereo camera. In *Dritter Workshop Virtuelle und Erweiterte Realität der GI-Fachgruppe VR/AR*, vol. 3, 107–118.
- KRAWCZYK, G., GOESELE, M., AND SEIDEL, H.-P. 2005. Photometric calibration of high dynamic range cameras. Research Report MPI-I-2005-4-005, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, April.
- KRISTENSEN, A. W., AKENINE-MILLER, T., AND JENSEN, H. W. 2005. Precomputed Local Radiance Transfer for Real-Time Lighting Design. *ACM Trans. Graph.* 24, 3, 1208–1215.
- LOSCOS, C., FRASSON, M., DRETTAKIS, G., WALTER, B., GRANIER, X., AND POULIN, P. 1999. Interactive virtual relighting and remodeling of real scenes. In *Tenth Eurographics Workshop on Rendering*, 329–340.
- MANTIUK, R., PATTANAIK, S., AND MYSKOWSKI, K. 2002. Cube-map data structure for interactive global illumination computation in dynamic diffuse environments. In *ICCVG*.
- NAKAMAE, E., HARADA, K., ISHIZAKI, T., AND NISHITA, T. 1986. A montage method: The overlaying of the computer generated images onto a background photograph. In *Computer Graphics (Proceedings of SIGGRAPH 86)*, 207–214.
- NG, R., RAMAMOORTHI, R., AND HANRAHAN, P. 2003. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph.* 22, 3, 376–381.
- NIJASURE, M., PATTANAIK, S. N., AND GOEL, V. 2005. Real-time global illumination on GPUs. *Journal of graphics tools* 10, 2, 55–71.
- NIMEROFF, J. S., SIMONCELLI, E., AND DORSEY, J. 1994. Efficient Re-rendering of Naturally Illuminated Environments. In *Fifth Eurographics Workshop on Rendering*, Springer-Verlag, Darmstadt, Germany, 359–373.
- PEREZ, R., SEALS, R., AND MICHALSKY, J. 1993. All-weather model for sky luminance distribution-preliminary configuration and validation. In *Solar Energy* 50, 3, 235–245.
- PHARR, M., AND HUMPHREYS, G. 2004. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann, August.
- RAMAMOORTHI, R., AND HANRAHAN, P. 2001. An efficient representation for irradiance environment maps. In *SIGGRAPH*, 497–500.
- REINHARD, E., WARD, G., PATTANAIK, S., AND DEBEVEC, P. 2005. *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting*. Morgan Kaufmann Publishers, Dec.
- RITSCHEL, T., AND GROSCHE, T. 2006. On-line estimation of diffuse materials. In *Dritter Workshop Virtuelle und Erweiterte Realität der GI-Fachgruppe VR/AR*, vol. 3, 95–106.
- SATO, I., SATO, Y., AND IKEUCHI, K. 1999. Acquiring a radiance distribution to superimpose virtual objects onto a real scene. *IEEE Trans. Vis. Comput. Graph.* 5, 1, 1–12.
- SILLION, F. X., AND PUECH, C. 1994. *Radiosity and Global Illumination*. Morgan Kaufmann Publishers, Inc., San Francisco, California.
- SLOAN, P.-P. J., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. Graph.* 21, 3, 527–536.
- WILLIAMS, L. 1978. Casting curved shadows on curved surfaces. In *SIGGRAPH '78: Proceedings of the 5th annual conference on Computer graphics and interactive techniques*, ACM Press, New York, NY, USA, 270–274.