Differential Photon Mapping - Consistent Augmentation of Photographs with Correction of all Light Paths

Thorsten Grosch

Institute for Computational Visualistics, University of Koblenz-Landau, Germany

Abstract

Augmenting images with consistent lighting is possible with differential rendering. This composition technique requires two lighting simulations, one simulation with only real geometry and another one with additional virtual objects. The difference of the two simulations can be used to modify the original pixel colors. The main drawback of differential rendering is that not all modified light paths can be displayed: The result of the lighting simulation is visible in a reflective object instead of the real environment, augmented with virtual objects. Moreover, many regions in the photograph remain unchanged and the same work is done twice without any visual effect. In this paper we present a new approach for augmenting a photograph with only a single photon mapping simulation. The changes in lighting introduced by a virtual object are directly simulated using a differential photon map. All light paths intersecting the virtual object are corrected. To demonstrate the correctness of our approach, we compare our simulation results with real photographs.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

1. Introduction and Previous Work

Adding virtual objects in photographs is a challenging task and was first mentioned by Nakamae et al. [NHIN86]. The differential rendering technique, an augmentation based on a lighting simulation, was introduced by Fournier et al. [FGR93]. Several improvements exist, like the fast update rates by Drettakis et al. [DRB97] and the final gathering approach by Loscos et al. [LFD*99]. Gibson et al. [GCHH03] developed a real-time differential rendering method for the graphics hardware, which was extended to panoramic images by Grosch [Gro05]. Augmented images with natural illumination were presented by Debevec [Deb98], using a high-dynamic-range [DM97] image of a light probe, a mirrored ball that captures the incoming illumination from all directions. All these approaches use differential rendering for the integration of virtual objects into the photograph. The idea is to reconstruct real geometry, camera parameters, materials and light sources from the photograph. This information can be then used in a lighting simulation which is similar to the photograph with a resulting radiance L_{Old}. After inserting virtual objects, a second lighting simulation is performed with the resulting radiance L_{New}. The virtual object is inserted in the photograph as it appears in this simulation. The difference $\Delta L = L_{New} - L_{Old}$ between the two simulations is added to the original photograph pixels to display the changes in illumination introduced by the virtual objects. Two main drawbacks exist with this composition technique: First, the simulation must be performed twice to calculate the difference ΔL which has to be added to the real pixels. As you can see in figure 1, large regions in the photograph remain unchanged because the two simulations are almost identical. Secondly, not all changed light paths can be simulated: Consider for example a virtual mirror object. The composition technique will display the result of the lighting simulation inside the mirror, but we expect to see the real environment, augmented with virtual objects. A very detailed reconstruction with textured polygons is required to overcome this problem. However, one of the strengths of differential rendering is that a coarse lighting simulation is sufficient to produce a convincing result. We therefore developed a new technique for the augmentation of photographs which directly simulates the changes in illumination if a virtual object is inserted. Our method is based on photon mapping [Jen95], which can simulate most of the visible lighting effects. The rest of the paper is organized as follows: Section 2 describes the reconstruction process, in section 3 we explain the photon distribution and the differential rendering. In section 4 we show our results before we conclude in section 5.

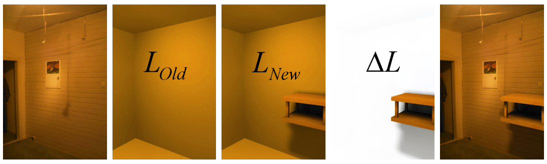

Figure 1: *Standard differential rendering: The difference between two simulations is used to modify the photograph.*

2. Reconstruction

Beside the normal photograph, we take a high dynamic range image of a light probe as input. Visible geometry on the photograph and camera parameters are reconstructed using object-based techniques described in [HZ04]. The light probe image serves as a description of incoming light with a parallel light source for each pixel. In addition, it is used to reconstruct diffuse reflectance values for all surfaces [Deb98] (currently, we assume a diffuse environment).

3. Augmenting the Photograph

3.1. Photon Distribution

To shoot photons towards virtual geometry, we calculate the bounding sphere with radius R of the virtual object. As shown in figure 2, each pixel in the environment map with radiance L_{pixel} carries the radiant flux $\phi_{pixel} = L_{pixel} \cdot \Omega_{pixel} \cdot \pi \cdot R^2$, where Ω_{pixel} is the solid angle for the current pixel. After summing up the flux of all pixels, we set the flux for one photon to $\phi_{photon} = \sum \phi_{pixel}/N$, with a used-defined number of photons N. This results in $\phi_{pixel}/\phi_{photon}$ photons for the current pixel.

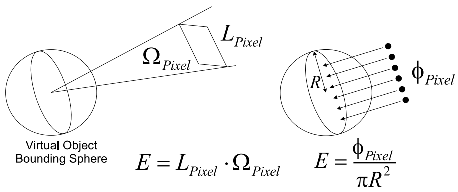

Figure 2: *Irradiance E resulting from one pixel of the environment map arriving at the virtual object.*

Each pixel in the environment map describes a parallel light source, as can be seen in figure 3. Therefore, the photons are first uniformly distributed on a disk with radius R perpendicular to the light direction and then shot towards the virtual object. If the photon hits the virtual object, we calculate the next intersection point with the real geometry and store a photon with the negative flux $-\phi_{photon}$. These shadow photons were introduced by [JC95], in our case they represent the missing light in the shadow region of the virtual object.

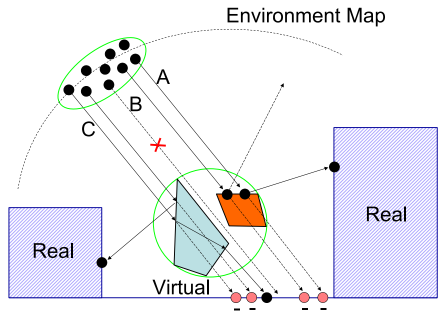

Figure 3: *Differential photon mapping: Photons are uniformly distributed on a circle perpendicular to the light direction. Shadow photons with negative flux are stored for photons which hit the virtual object. The left part of the virtual object is glass, the right part is diffuse.*

Otherwise, if the photon does not intersect the virtual object, it is completely ignored, because in this case the light path does not change (ray B). In comparison to standard differential rendering, these photons would be equal in both simulations and subtract to zero.

The following light paths for the photons which hit the virtual object are identical to normal photon mapping and shown in figure 3. For diffuse materials, the photon is stored and a randomly reflected direction is chosen (rays A). In case of reflective or refractive materials (rays C), we calculate a Fresnel reflectance value and select the reflected or refracted direction with Russian roulette [Jen01].

3.2. Displaying the Augmented Photograph

To display the augmented photograph, we send a ray through each pixel of our photograph and differentiate between real and virtual geometry, as shown in figure 4. In case of real (diffuse) geometry (pixel 1), we calculate the radiance estimate ΔL and add it to the pixel radiance: $L'_{pixel} = L_{pixel} + \Delta L$. Because we simulated a differential photon map, ΔL can be positive or negative. A positive radiance estimate indicates caustics or color bleeding, a negative radiance estimate occurs in shadow regions of the virtual object. If the primary ray hits a virtual object, the rendering process is different: Diffuse objects can be rendered directly using the radiance estimate at the intersection point (pixel 2). In this case the radiance estimate describes the total radiance because negative photons are only stored on real objects. For reflective or refractive virtual objects (pixel 3), we start a normal ray tracer for the reflected and refracted path until an intersection point with a diffuse material is found. If this point belongs to a virtual object, we return the radiance estimate. If the point belongs to real geometry, we first detect the corresponding pixel in the photograph by back-projecting the 3D point onto the photograph (pixel 4). This pixel contains the radiance L^*_{pixel} which is now visible through pixel 3. Then we calculate the radiance estimate ΔL and return the corrected pixel value $L^*_{pixel} + \Delta L$. This correction ensures that the augmented environment is visible in a reflective virtual object.

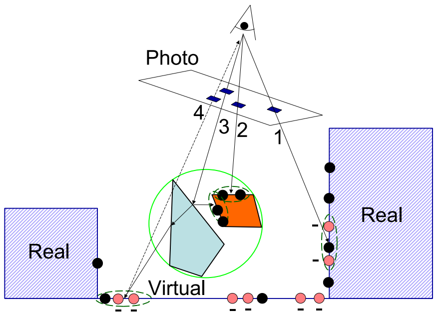

Figure 4: *Radiance estimate: The differential photon map contains the changes in radiance which must be added to the photograph.*

4. Results

In order to validate our approach, we compare our simulation results with real photographs, as shown in figure 5. Figure 6 demonstrates the improved image quality in comparison to standard differential rendering. Results for some more complex geometry are shown in figure 7.

5. Conclusions and Future Work

In this paper, we have presented a new approach for augmenting a photograph that is significantly faster than standard differential rendering. Instead of performing two complete lighting simulations, only the changes introduced by inserting a virtual object are directly simulated using a differential photon map. Moreover, our method improves the display quality, because all changed light paths can be calculated, especially reflective and refractive objects show the real, augmented environment. We are currently developing a final gathering version for high-quality rendering and a parallel version of our photon mapper. Moreover, we plan to integrate more complex reflection models for both real and virtual objects.

References

[Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proc. SIGGRAPH '98* (1998), pp. 189–198.

[DM97] DEBEVEC P., MALIK J.: Recovering high dynamic range radiance maps from photographs. In *Proc. SIGGRAPH '97* (1997), pp. 369–378.

[DRB97] DRETTAKIS G., ROBERT L., BOUGNOUX S.: Interactive common illumination for computer augmented reality. In *Eighth Eurographics Workshop on Rendering* (1997), pp. 45–56.

[FGR93] FOURNIER A., GUNAVAN A., ROMANZIN C.: Common illumination between real and computer generated scenes. In *Proc. of Graphics Interface 1993* (1993), pp. 254–262.

[GCHH03] GIBSON S., COOK J., HOWARD T., HUBBOLD R.: Rapid shadow generation in real-world lighting environments. In *Eurographics Symposium on Rendering* (2003), pp. 219–229.

[Gro05] GROSCH T.: PanoAR: Interactive augmentation of omni-directional images with consistent lighting. In *Mirage 2005, Computer Vision / Computer Graphics Collaboration Techniques and Applications* (2005), pp. 25–34.

[HZ04] HARTLEY R., ZISSERMAN A.: *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.

[JC95] JENSEN H. W., CHRISTENSEN N.: Efficiently rendering shadows using the photon map. In *Proceedings of Compugraphics 95* (1995), pp. 285–291.

[Jen95] JENSEN H. W.: Importance driven path tracing using the photon map. In *Rendering Techniques 95* (1995), pp. 326–335.

[Jen01] JENSEN H. W.: *Realistic Image Synthesis using Photon Mapping*. AKPeters, 2001.

[LFD*99] LOSCOS C., FRASSON M., DRETTAKIS G., WALTER B., GRANIER X., POULIN P.: Interactive virtual relighting and remodeling of real scenes. In *Tenth Eurographics Workshop on Rendering* (1999), pp. 329–340.

[NHIN86] NAKAMAE E., HARADA K., ISHIZAKI T., NISHITA T.: A montage method: The overlaying of the computer generated images onto a background photograph. In *Computer Graphics (Proceedings of SIGGRAPH 86)* (1986), pp. 207–214.

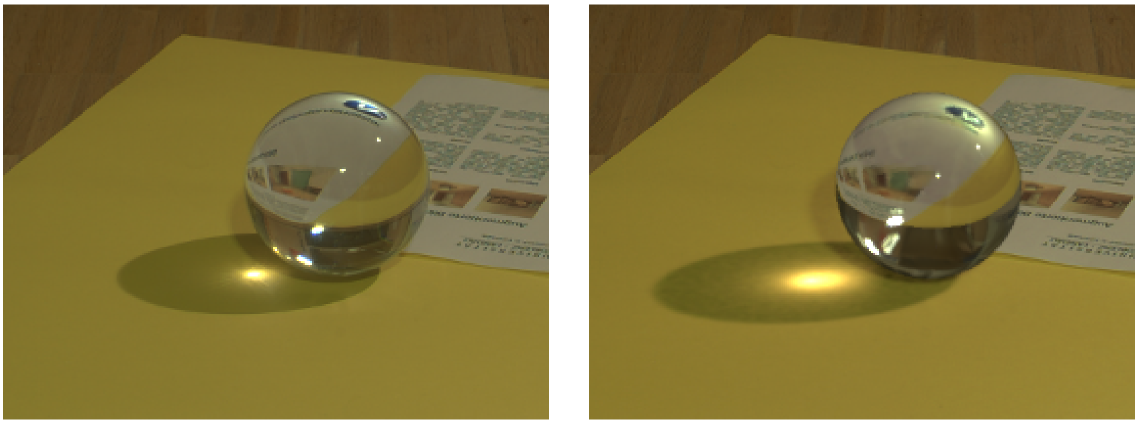

Figure 5: *Comparison for a glass sphere: The left image is the real photograph, the right image is our simulation with an estimated index of refraction for glass. The differential photon map makes it easy to simulate all these lighting effects which are difficult with the standard composition technique.*

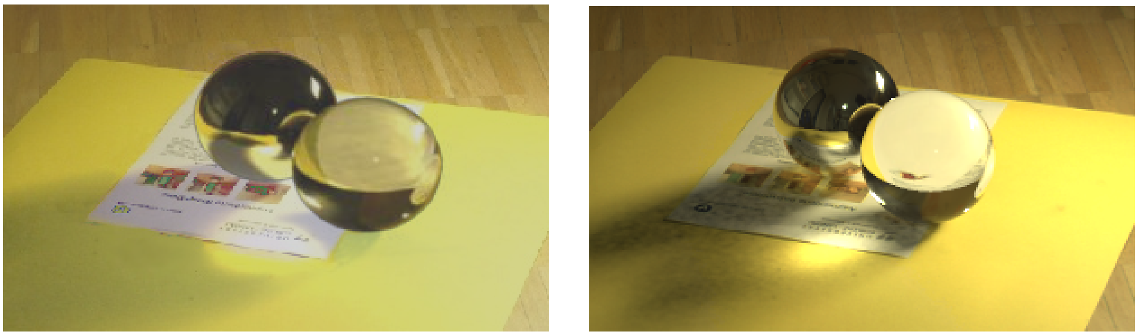

Figure 6: *The left image is standard differential rendering based on two final gathering simulations (rendered with Mental Ray©), which took about two hours. Inside the mirror sphere, the result of a coarse lighting simulation with untextured polygons is visible. The right image shows our simulation. Note the correct reflections: The real environment, augmented with virtual objects, is visible inside the spheres. The rendering time for the right image was only five minutes, because we exclusively simulate the changes in lighting.*

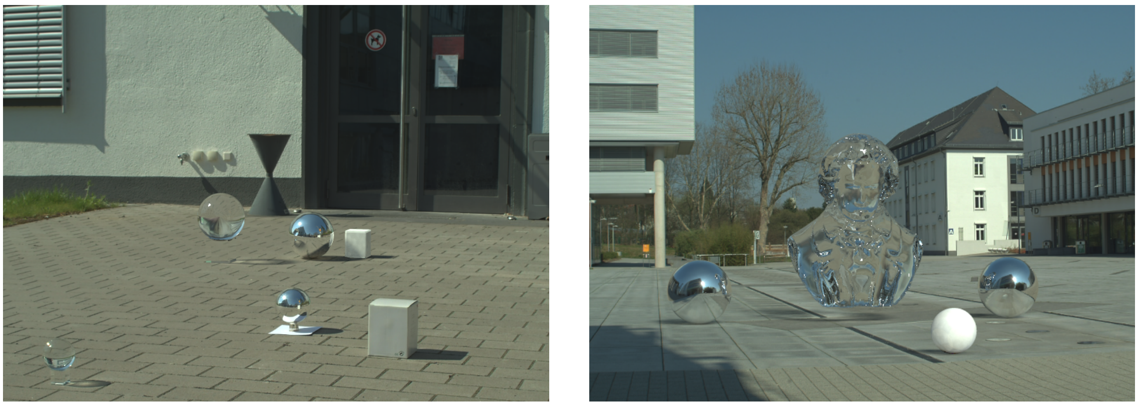

Figure 7: *The left image is a comparison between real objects in the foreground and virtual objects in the background. The simulation time was 8.5 minutes for 2.8 million photons. The right image shows some more complex geometry (10000 polygons). The simulation time was 1.8 hours for 28.000 photons. Both images were generated on a Dual Xeon PC, running at 2.8 GHz.*