Goniometric Light Reconstruction for Augmented Reality Image Synthesis

Thorsten Grosch

Institut für Computervisualistik
Universität Koblenz-Landau
Universitätsstraße 1
D-56070 Koblenz
grosch@uni-koblenz.de

Stefan Müller

Institut für Computervisualistik
Universität Koblenz-Landau
Universitätsstraße 1
D-56070 Koblenz
stefanm@uni-koblenz.de

Wolfram Kresse

Fraunhofer IGD
Fraunhoferstraße 5
D-64283 Darmstadt
wkresse@igd.fhg.de

Abstract. Reconstruction of light and material properties from digital photographs is a challenging problem. We describe a new image-based approach for the reconstruction of light and material properties from a set of photographs. Using these values in a lighting simulation makes it possible to augment the photograph with correct lighting conditions in all regions of the photograph. Since light sources are the most important component for the correct appearance of virtual objects, special emphasis is put on the reconstruction of a light goniometric as a realistic light source description. Furthermore we describe a new data structure to rapidly detect changing regions in the photograph for a fast update after moving objects. For high quality images we show how to use anti-aliasing techniques to improve the display quality especially at the edges between real and synthetic objects.

1 Introduction

Inserting virtual objects in photographs with correct appearance is an active area of research in the last few years. The first attempt was made by Nakamae et al.[Na86], where virtual buildings were placed in a photograph. The base for most of the following work was done by Alain Fournier et al.[FGR93]. They introduced the differential rendering technique for displaying virtual objects with correct lighting and shadows in a photograph based on a global illumination simulation. Drettakis et al.[DRB97] improved this method using the line-space hierarchy[DS97] for interactive update rates when moving a virtual object. Loscos et al.[LDR00][Lo99] achieved better display quality by separating direct and indirect light and recomputing direct light for each pixel. A system working in real-time by using the graphics hardware was introduced by Gibson et al.[GM00] who use shadow maps and a multi-pass rendering algorithm for virtual object shadows. Before virtual objects can be inserted in a photograph, illumination characteristics must be reconstructed. Many publications exist for the reconstruction of material properties from photographs. Fournier et al.[FGR93] and Drettakis et al.[DRB97] are using heuristics for approximate reflectance values. Loscos et al.[Lo99] use an inverse radiosity approach with approximate indirect light. Yu et al.[Yu99] developed an inverse radiosity algorithm for the reconstruction of diffuse and specular material properties from a set of high-dynamic-range photographs[DM97]. The algorithm by Boivin and Gagalovicz[BG01] uses different reflection models, starting with diffuse. Most of these publications use the assumption of known light sources. There are only a few attemps for the reconstruction of light sources. Fournier[FGR93] uses a best-fit approach for the emittance values of area light sources with known positions. Debevec[De98] introduced the omni-directional image as a description of incoming light and Sato et al.[SSI99a] use a fish-eye view photograph as incoming light description. In a following paper, Sato et al.[SSI99b] calculate a radiosity distribution on a hemisphere from several pixel brightnesses in shadow regions. Parallel to our work, Gibson et al. [GHH01] developed an image-based approach to reconstruct a set of virtual light sources which mimic the behaviour of direct and indirect light. Gibson et al.[GM00] use the image of a light probe to reconstruct a set of directional light sources which best approximate the shadow of a known object. The drawback of the light-based methods is that they are only valid for a small part of the scene, the light sources are expected to be far away. Recently, Unger et al. [Un03] described how to calculate spatially-varying light in a small region and Gibson et al. [Gi03] improved their algorithm for the current graphics hardware. All these methods do not calculate the real properties of the light sources and sampling from a whole hemisphere is usually very time consuming. We therefore reconstructed the real characteristics of the light source to get a light description that is valid in all regions in the photograph.

2 Algorithm Description

This chapter describes the reconstruction of light and material properties from a set of high-dynamic-range photographs[DM97].

2.1 Light Goniometrics

We use light goniometrics [Sc00] for the description of lights since this is one of the most realistic approximations of a real light source. Recently, Goesele[Go03] presented an even better way to capture and render real light sources. A light goniometric describes the luminous intensity of a light source for any outgoing direction (Θ, φ). Fig. 1 shows a light goniometric and the information captured from a light probe for comparison. The luminous intensity of only one direction of the spotlight is captured in the light probe image, if we move the virtual object away from the light probe position, the change in light intensity is lost.

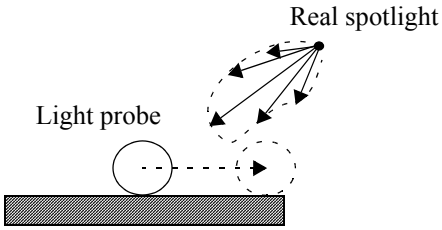

Fig. 1. Light goniometric: The light probe captures the correct light intensity only for one position, if we move away from this position, the intensity is no longer correct.

Each realistic light source has its own intensity distribution, using point or area lights will always lead to inaccurate results. The radiosity equation for a patch E with light goniometrics as light sources is

$$B_E = \rho_E \cdot \left(\sum_i \frac{I_i(\Theta, \varphi) \cdot \cos\Theta_{Ei}}{d_i^2} + \sum_S B_S F_{ES} \right)$$

The equation is split into a direct part describing the incoming irradiance from the lights and an indirect part with the illumination from the surrounding patches. I_i is the luminous intensity of a light source in direction (Θ, φ).(Fig. 2)

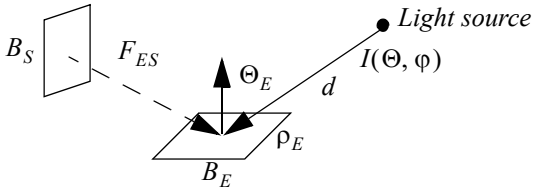

Fig. 2. Patch E receives direct light from a light source and indirect light from patch S.

2.2 Inverse Radiosity

Inverse radiosity describes the recalculation of light and material properties from a converged radiosity solution. This is a diffucult problem if neither the light intensity nor the

reflectance values are known. Look at the right side of Fig. 7: Both simulations look similar, but the left picture was created with a white light source whereas the box on the right picture has a bright green light source and low reflectance values. Without any additional information it is often impossible to reconstruct the correct combination of light intensity and surface reflectance, there are many combinations leading to the same result. [GHH01] and [SSI99b] developed iterative algorithms which usually converge towards the correct solution but they depend on good starting values and it is uncertain if the correct combination was found. We therefore decided to overcome this ambiguity with one user-specified reflectance value, for example the known material of an object in the photograph. With a known reflectance value we can solve the radiosity equation for the unknown luminous intensity in this direction (if there is only one light):

$$I(\Theta, \varphi) = \left(\frac{B_E}{\rho_E} - \sum_S B_S F_{ES} \right) \cdot \frac{d^2}{\cos\Theta_E} \qquad \text{(EQ1)}$$

On the other hand, if we have information about the light sources, we can solve the equation for reflectance:

$$\rho_E = \frac{B_E}{\sum_i \frac{I_i(\Theta, \varphi) \cdot \cos\Theta_{Ei}}{d_i^2} + \sum_S B_S F_{ES}} \qquad \text{(EQ2)}$$

We will use this fact in the algorithm described in the next section.

2.3 Algorithm Description

The following algorithm reconstructs the light goniometric and the material properties with one user-defined reflectance value ρ_{Ref}. We assume that this diffuse reflectance value is constant over the whole polygon, so we first calculate the luminous intensity for all directions leading to patches with material ρ_{Ref} using EQ1. For each patch we take the average high-dynamic-range value of all pixels inside the patch as the patch radiosity B_E. Since every real object has a textured surface, we calculate the average surface brightness instead of using every single pixel value which might confuse the reconstruction algorithm. Then we search for a small angle α_{Min} between the direction of a known luminous intensity I_1 and an unknown luminous intensity I_2 (Fig. 3). Since the angle is small, we can assume that the intensities will be similar, so we set $I_2 = I_1$ and calculate the unknown reflectance value ρ using EQ2. With the known reflectance value we can reconstruct the luminous intensity for all directions leading to patches with this material. These two steps are repeated until all patches visible from the light source are found. Fig. 7 shows the reconstruction of a light goniometric and the radiosity simulation with the reconstructed light and reflectance values. An augmented photograph based on this simulation is shown on Fig. 8.

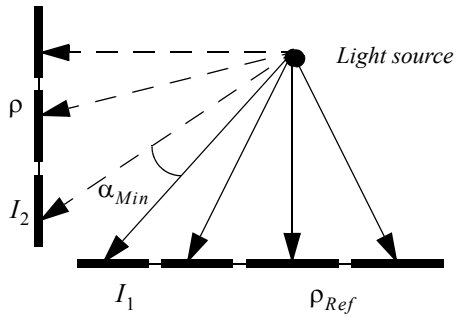

Fig. 3. Reconstruction of light and material : The reflectance value of the lower polygon is known, so we can determine the luminous intensity in this direction. The reflectance value of the left polygon is unknown, but since the intensities I_1 and I_2 are similar, we can approximate this reflectance value.

2.4 Light Source Position

If the light source is not visible on one of the photographs, we can estimate the light source position from the shadows. A corner point of a shadow V_S and the corresponding geometry vertex V_G define a so-called *shadow-line* with the light source position somewhere on this line. If we can find two or more of these shadow-lines in the photograph, the light source position must be the common intersection point. If we only see the shadow edge e_S and the corresponding geometry edge e_G we can conclude that the light source must be in the *shadow-plane* defined by the two edges.

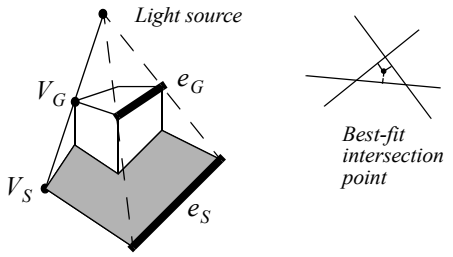

Fig. 4. The unknown light source position is detected as the best-fit intersection point from shadow-lines (defined by V_S and V_G) and shadow-planes (defined by e_S and e_G)

Trying to calculate a single common intersection point of all shadow-lines and shadow-planes does not work in pratice, because of small errors in the user-defined positions. Therefore we calculate a best-fit intersection point with a minimal sum of squared distances to all shadow-lines and shadow-planes. A detailed description of this process can be found in [Gr01].

3 User Interaction

Since our algorithm runs only in software at the moment, we developed some data structures to decrease the update time after moving the virtual object.

3.1 Extended Pixel Data Array

The main problem with many realistic lighting models is that no user interaction is possible because the lighting simulation takes a lot of time. Therefore we introduce a new data structure to locate changing regions in the photograph due to virtual objects. After the first radiosity simulation we store the following information for each pixel:

```
Vec pos (3D position)
Patch *patch (patch belonging to pixel)
Color B_Dir (direct radiosity for the pixel)
Color B_Ind (interpolated indirect radiosity)
float FF (unoccluded formfactor)
float visib (visibility to light source)
Shaft pixelShaft (volume between pixel and light)
```

Our data structure is based on the Pixel Data Array introduced by Loscos et al.[Lo99]. Loscos et al. used the shafts [HW91] belonging to the links created in the hierarchical radiosity simulation to decide whether the direct light for a pixel has changed or not. For large patches this means a certain over-estimation, therefore we create a so-called *pixel shaft* describing the volume between pixel position and light source (Fig. 5). The direct light for a pixel is only recomputed if the virtual object intersects the pixel shaft, otherwise we use the value from the first simulation stored in the Pixel Data Array. If the virtual object overlaps the pixel shaft, visibility is recomputed and the unoccluded formfactor is weighted with the new visibility. For visibility computation, we use ray-casting with a few rays placed around the light source position in a jittered grid. The shape of the grid is provided by the user and approximates the extent of the real light source. This grid is recursively subdivided along the longest axis. For each resulting small box, several random points are precomputed and stored. One point of each small box is selected for visibility computation. Because of the recursive subdivision, we can create a jittered grid with different resolutions. An oracle function can be adapted to decide per pixel how many shadow samples are necessary for a given image quality.

3.2 Pixel Blocks

Comparing the 3D position of one pixel with its neighbour pixels, there are often similar points. Fig. 5 shows four points with their positions. If we build the bounding box of these points and create a shaft to the light source we can exclude direct light calculation for the whole pixel block if the virtual object is outside this shaft. Connecting adjacent pixel blocks leads to a hierarchy of pixel blocks with their pixel block shafts. The whole

picture is the root of this *pixel shaft hierarchy*. The right image in Fig. 6 visualizes the recursive traversal of the hierarchy, each colored block describes a region without changes in direct light detected with one shaft test. With this hierarchy we do not have to check each pixel and the direct light recomputation is much faster, especially for small virtual objects. If a shaft is affected by the virtual object, we recursively test the children shafts for intersection, otherwise we ignore the pixel block associated with this shaft.

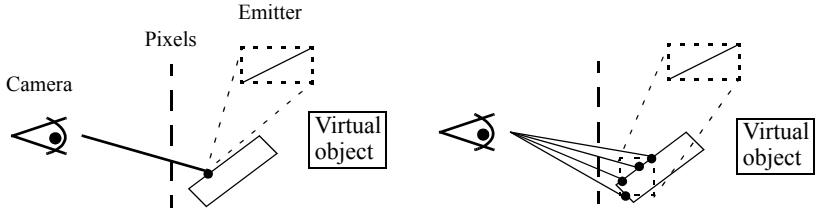

Fig. 5. Pixel shaft and pixel block shaft. Left image: Before recalculating the direct light for a pixel, we test if the virtual object intersects the shaft between the pixel and the light source. Right image: The virtual object is outside the shaft of a pixel block, the direct light for the pixel block remains unchanged.

4 Display Quality

For a realistic integration of virtual objects in photographs we use the differential rendering described in [FGR93]. Like [Lo99], we use final gathering to improve the image quality. Indirect light is taken from the radiosity simulation and interpolated between the vertex radiosities. Direct light is recomputed for each pixel and updated after object movement if necessary. For high quality images, we calculate the differential rendering on a higher resolution than the image resolution. The reason for this is shown in Fig. 6: The chair does not look like a real object because of the aliasing effects along the border between real and virtual pixels. Taking a picture with a camera, the resulting pixel color is always an average value of the colors seen through this pixel window. To mimic this behaviour, we subdivide pixels into subpixels and search pixels with at least one virtual subpixel. The differential rendering is done for all subpixels and as a result, the average value is displayed. We found that it is not necessary to calculate the whole picture with a higher resolution, for shadows and changes in indirect light the normal resolution is sufficient. Moreover, we do not calculate the indirect light for each subpixel and use the pixel midpoint value stored in the Pixel Data Array instead without visible changes.

5 Conclusions and Future Work

We presented a new approach for the reconstruction of a realistic light source description and diffuse reflectance values from a set of photographs. The pixel shaft hierarchy was introduced as a data structure to find regions in the photograph with changes in direct

light when inserting and moving a virtual object. The main drawback of our method is the limitation to one light source. The problem with the reconstruction of many light goniometrics are the ambiguities: The brightness of a surface lit by two different light sources can result from an infinite number of two different light gonimetrics producing the same light distribution on the surface. A first step towards reconstructing multiple light sources could be the high-dynamic-range photograph of a light probe which describes a first approximation of all light sources. Then, the user sets the general shape of each light source, like a specular lobe for a spotlight. The reconstruction algorithm then tries to reconstruct light goniometrics which are correct for the light probe position and as close as possible to the user-defined shape. Moreover, our algorithm is limited to diffuse environments at the moment, the extension to specular reflection is a task we left as future work.

Statistics

The statistics were made with an SGI Infinite Reality (R10000, 194 MHz). The photograph resolution is 341 x 512 pixels (Fig. 6). The radiosity simulations took 3.28 seconds, creating the Pixel Data Array took 44.16 seconds. The update time after moving the chair depends mainly on the number of shadow rays used for direct light computation. A high quality picture takes 21.6 seconds, a fast update with one shadow ray takes 3.3 seconds. Anti-aliasing took 101.2 seconds with a 3 x 3 resolution for the sub-pixels. The memory requirements were about 36 MB (mainly for the Pixel Data Array). The light goniometric shown on Fig. 7 was reconstructed from some photographs with a resolution of 768 x 512 pixels. The reconstructed 3D geometry consists of 27 polygons with 282 patches. The reconstruction time was 97.9 seconds, 48.7 seconds for initalizing the patch radiosities and 49.2 seconds for the reconstruction of light and materials.

Fig. 6. The left image was generated with standard differential rendering. The middle image was calculated on a higher resolution, the virtual object looks better integrated in the photograph. Right image: Changes in direct light are rapidly identified with the pixel shaft hierarchy. Each colored block represents a region in the photograph without changes in direct light, identified with only one shaft test

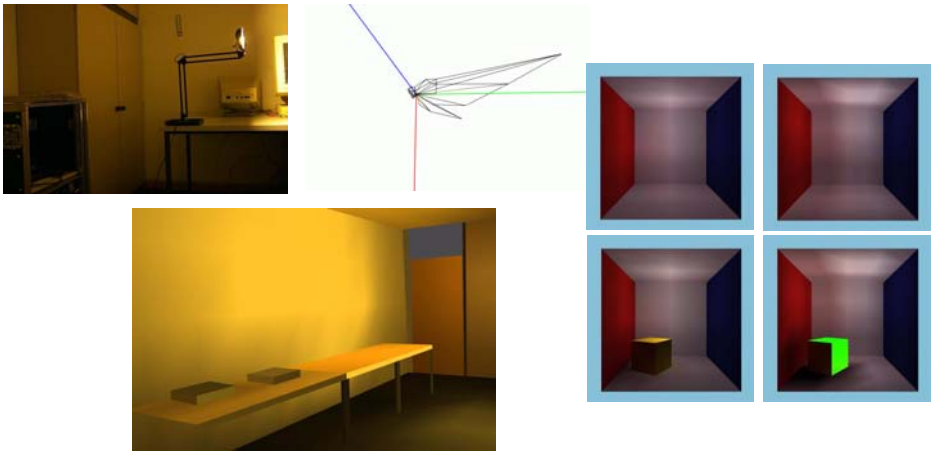

Fig. 7. The upper left image shows the light source that was used on the photograph (Fig. 8). The light goniometric for this light was reconstructed from some photographs and is shown in the middle image. The inaccuracies are mainly because of the incomplete 3D reconstruction of the geometry. The user-defined reflection value was the left wall with (0.7, 0.7, 0.7). The bottom left image is a radiosity solution based on the reconstructed light and materials. Note that the simulation values are similar to the photograph in Fig. 8. The 3D model was reconstructed using computer vision techniques described in [Gu00]. Right side: The top row shows a radiosity simulation with a white light (left box) and a green light (right box) with different reflection values. Both simulations are similar, a yellow box is inserted in the scene (bottom row), showing the difference.

Fig. 8. Some virtual objects were added to the image. Object appearance and shadow brightness are consistent with the photograph.

Acknowledgements

We would like to thank Dirk Reiners for the photographs of his living room and his desktop lamp. We also thank Frank Schoeffel for the help with Frame Maker and Didier Stricker for the support with the 3D reconstruction.

References

[BG01] Boivin, S., Gagalowicz, A.: *Image-Based Rendering of Diffuse, Specular and Glossy Surfaces from a Single Image, SIGGRAPH 2001*

[De98] Debevec, P.: *Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography.* Proceedings of SIGGRAPH 98 (July 1998), 189-198

[DM97] Debevec, P.E., and Malik, J.: *Recovering high dynamic range radiance maps from photographs.* Proceedings of SIGGRAPH 97 (August 1997), 369-378

[DRB97] Drettakis, G.; Robert, L.; Bougnoux, S.: *Interactive Common Illumination for Computer Augmented Reality*, Eighth Eurographics Workshop on Rendering (1997). In: Dorsey, J.; Slusallek, P. (Hrsg.): Rendering Techniques '97. Springer Verlag, 1997, pp. 45–56.

[DS97] Drettakis, G.; Sillion, F.: *Interactive Update of Global Illumination Using a Line-Space Hierarchy*, ACM Computer Graphics (SIGGRAPH '97 Proceedings), 31(3):57–64, 1997.

[FGR93] Fournier, A., Gunavan, A.S., and Romanzin, C.: *Common Illumination between real and computer generated scenes.* In Proceedings of Graphics Interface '93 (San Francisco, CA, May 1993), Morgan Kaufmann, pp. 254-262

[GHH01] Gibson, S., Howard, T.J., Hubbold, R.J.: *Flexible Image-Based Photometric Reconstruction using Virtual Light Sources*, Eurographics 2001,Manchester, UK, September 2001

[GM00] Gibson, S.; Murta, A.: *Interactive Rendering with Real-World Illumination*, Eleventh Eurographics Workshop on Rendering. In: Péroche, B.; Rushmeier, H. (Hrsg.): Rendering Techniques 2000. Springer Verlag, 2000, pp. 365–376.

[Gi03] Gibson, S.;Cook, J.;Howard, T.;Hubbold, R.:*Rapid Shadow Generation in Real-World Lighting Environments*, Eurographics Symposium on Rendering 2003

[Go03] Goesele, M.; Granier, X.;, Heidrich, W.; Seidel, H.P.:*Accurate Light Source Acquisition and Rendering.* ACM Transactions of Graphics (Proceedings of ACM SIGGRAPH 2003) July 2003, Vol. 22, Number 3, pp 621-630

[Gr01] Grosch, T.: *Interaktive Erweiterung und Beleuchtungsrekonstruktion von Fotos mit Radiosity*, Diploma Thesis, Technische Universität Darmstadt 2001

[Gu00] Gulliou, E., Menevaux, D., Maisel, E., Bouatouch, K.:*Using vanishing points for camera calibration and coarse 3D reconstruction from a single image*. In Visual Computer 16 (2000) 7, 396-410

[HW91] Haines, E.A.; Wallace, J.R.: *Shaft Culling for Efficient Ray-Traced Radiosity*, Proceedings of the Second Eurographics Workshop on Rendering, Barcelona, Spanien, 1991.

[LDR00] Loscos, C.; Drettakis, G.; Robert, L.: *Interactive Virtual Relighting of Real Scenes*, IEEE Transactions on Visualization and Computer Graphics, 6(4):289–305, 2000.

[Lo99] Loscos, C.; Frasson, M.-C.; Drettakis, G.; Walter, B.; Granier, X.; Poulin, P.: *Interactive Virtual Relighting and Remodeling of Real Scenes,* Tenth Eurographics Workshop on Rendering. In: Lischinski, D.; Ward-Larson, G. (Hrsg.): Rendering Techniques '99. Springer Verlag, 1999, pp. 329–340.

[Na86] Nakamae, E., Harada, K., Ishizaki, T., and Nishita, T.: *A montage method: The overlaying of the computer generated images onto a background photograph.* Computer Graphics (Proceedings of SIGGRAPH 86) 20, 4 (August 1986), 207-214

[Sc00] Schmidt-Clausen, H.J.:*Grundlagen der Lichttechnik,* Unterlagen zur Vorlesung 2000, TU Darmstadt

[SSI99a] Sato, I., Sato, Y., and Ikeuchi, K.: *Acquiring a radiance distribution to superimpose virtual objects onto a real scene*. IEE Transactions on Visualization and Computer Graphics 5, 1 (January - March 1999), 1 - 12

[SSI99b] Sato, I., Sato, Y., and Ikeuchi, K.: *Illumination distribution from brightness in shadows: Adaptive extimation of illumination distribution with unknown reflectance properties in shadow regions*. In Proceedings of IEEE ICCC'99, pages 875-882, September 1999

[Un03] Unger, J.; Wenger, T.; Hawkins, T.; Gardner, A.; Debevec, P.:*Capturing and Rendering With Incident Light Fields*, Eurographics Symposium on Rendering 2003, pp 1-10

[Yu99] Yu, Y., Debevec, P., Malik, J., and Hawkins, T.: *Inverse global illumination : Recovering reflectance models of real scenes from photographs*. In A.Rockford, editor, Computer Graphics (Proceedings of SIGGRAPH 99), volume 19, pages 215-224. Addison Wesley Longman, August 1999