Institut für Computergraphik und Algorithmen
Technische Universität Wien

Karlsplatz 13/186/2
A-1040 Wien
AUSTRIA

Tel: +43 (1) 58801-18688
Fax: +43 (1) 58801-18698

Institute of Computer Graphics and Algorithms
Vienna University of Technology

email:
technical-report@cg.tuwien.ac.at

other services:
http://www.cg.tuwien.ac.at
ftp://ftp.cg.tuwien.ac.at

Real-Time Rendering of Globally Illuminated Soft Glossy Scenes With Directional Light Maps

Heinrich Hey Werner Purgathofer

TR-186-2-02-05
March 2002

Abstract

Directional light maps allow to display globally illuminated static scenes with soft glossy surfaces at interactive frame-rates. They store the spatially and directionally varying incoming light at a surface in form of textures for a global set of incoming light directions. The directional light maps can be generated in a photon tracing preprocessing step. During a walkthrough they are used to render the globally illuminated surfaces with graphics hardware.

Real-Time Rendering of Globally Illuminated Soft Glossy Scenes With Directional Light Maps

Heinrich Hey, Werner Purgathofer

Vienna University of Technology
{hey,wp}@cg.tuwien.ac.at

Abstract. Directional light maps allow to display globally illuminated static scenes with soft glossy surfaces at interactive frame-rates. They store the spatially and directionally varying incoming light at a surface in form of textures for a global set of incoming light directions. The directional light maps can be generated in a photon tracing preprocessing step. During a walkthrough they are used to render the globally illuminated surfaces with graphics hardware.

1 Introduction

To display globally illuminated soft glossy surfaces not only the spatial distribution, but also the directional distribution of the illumination has to be known. Directional light maps support this by storing the incoming light at a surface for several incoming light directions. Each directional light map (DLM) is a texture that represents the spatially varying incoming light at a surface from one of these directions. The DLMs are generated in a global illumination simulation, eg. photon tracing, in a preprocessing step. Afterwards in an interactive walkthrough these DLMs are used for hardware accelerated rendering of the soft glossy surfaces including their view dependent global illumination.

By using a global set of incoming light directions for the DLMs of all surfaces in the scene, the hardware accelerated rendering can be efficiently done with only few state switches, which avoids expensive stalls of the hardware rendering pipeline.

In section 2 we give a short overview of related methods. Afterwards we describe the generation of the DLMs with photon tracing, and the hardware accelerated rendering in section 3. We finally present our results in section 4.

2 Related methods

The illumination on a diffuse surface can be represented with an illumination map [Arv86], which is a texture map that stores the spatially varying irradiance on the surface. This information, which is generated in a global illumination simulation, is enough to correctly display diffuse surfaces from arbitrary view points. If non-diffuse surfaces shall be displayed from arbitrary directions then information about the directionally variant incoming or outgoing light at the surface is required.

Light fields [LH96] and lumigraphs [GG+96] store the outgoing radiance of an object as 4-dimensional function on an image plane (light field), or on a cube that encloses the object (lumigraph). The outgoing radiance may also be stored directly at the surfaces of the object with surface light fields [MRP98],[WA+00], with wavelets [SS+00], or with eigen-textures [NSI99].

Graphics hardware light sources may be used to represent the outgoing radiance by fitting a small number of virtual light sources (usually 8 hardware light sources are available) for each object individually, so that the resulting phong lobes represent the glossy highlights on the object as best as possible [WA+97]. Virtual light sources may also be used to display a radiosity solution by placing them at the positions of the most contributing sending patches to illuminate a receiving glossy patch [SSS95]. Here the hardware light sources also have to be set for each glossy patch individually.

The incoming radiance from far away objects may be stored in an environment map [Hei99],[Hei01], which may be prefiltered for the rendering of reflections on glossy surfaces. Glossy reflections may also be rendered with an on-the-fly convolution of images of pure specular reflections [BH+99]. The incoming light may also be stored in a directional irradiance mesh [Stü98], or in a photon map which can be rendered at nearly interactive frame-rates by drawing splats of the photons on the surfaces using graphics hardware [SB97].

3 Directional light maps

A DLM $m_{s,\Psi}$ is a texture on a surface s that stores the spatially varying incoming light that s receives from an incoming light direction Ψ. The texels' values correspond to the irradiance of this light on a plane perpendicular to Ψ. For each soft glossy surface s in the scene with surface normal N_s, and for each direction $\Psi \in \Omega$ which is frontfacing to s ($0 < N_s \cdot \Psi$), a DLM $m_{s,\Psi}$ is stored. Ω is the predefined global set of light directions which is used for all surfaces. Note that for each surface DLMs are generated and processed during rendering for only 50% of the directions of Ω (the frontfacing ones).

For the efficiency of hardware accelerated rendering it is essential that all surfaces use the same set Ω of light directions, as explained in section 3.2. This global set of light directions also avoids illumination discontinuities at the borders of adjacent surfaces. Such discontinuities would arise if adjacent surfaces would be illuminated from different directions, as it would be the case if each surface would have its individual set of light directions.

3.1 Generation of directional light maps

The DLMs are generated in a preprocessing step. First of all, Ω has to be defined. This is done by selecting n_Ω uniformly distributed directions on the unit sphere [Shi92]. n_Ω is a user defined value, and determines the directional accuracy of the illumination resulting from the generated DLMs. A larger n_Ω allows directionally more precise illumination, but requires more texture memory and rendering time to store and render the larger number of DLMs.

Next, the DLMs are generated with photon tracing. A large number (typically several millions) of light paths are stochastically shot from the light sources into the scene. At each hit point x of a light path at a surface s, the light path's incoming power is splatted into that DLM of s which is directionally nearest to the light path's incoming direction Ψ_l. This is the DLM $m_{s,\Psi}$ where $\Psi_l \cdot \Psi$ is maximum. The splat is centered at that texel of $m_{s,\Psi}$ which maps to x. The splats' size and shape are user defined, and determine the resulting spatial blurring and noise in the DLMs. In our implementation we have used pyramidal splats which have been 1.5 texels wide. Alternatively, a more advanced photon density estimation, eg. local linear density estimation [WH+97], could be used for each DLM instead of the splatting. The texture resolution of $m_{s,\Psi}$ is selected proportionally to the area of the projection of s onto a plane perpendicular to Ψ.

Graphics hardware supports only texel values in the range [0,1], therefore the irradiance values have to be mapped to texel values. Let k be the directional hardware light source intensity that corresponds to the texel's irradiance on a plane perpendicular to Ψ, and let k_{max} be the directional hardware light source intensity that corresponds to the user defined maximum representable irradiance. The texel value is then

$$t = \min\left(\frac{k}{k_{max}}, 1\right).$$

The DLMs are packed together into large textures. Usually many (small) DLMs fit into one of these textures, thereby only needing few textures. DLMs with the same Ψ are preferably put into the same texture, because the DLMs are used in the order of their Ψ during rendering.

3.2 Interactive rendering

During an interactive walkthrough the DLMs of a surface s are used for hardware accelerated rendering of the view dependent global illumination on s. Each DLM $m_{s,\Psi}$ illuminates s by modulating a directional hardware light source which shines from direction Ψ with intensity k_{max}. The view dependent contributions of the DLMs on s are accumulated together in the image. The whole scene is rendered in the following steps:

- Clear the z-buffer and color buffer.
- Draw the scene with correct visibility into the z-buffer, without modifying the color buffer.
- For each $\Psi \in \Omega$:
  - Set a single directional light source, with its direction=Ψ, and with its intensity=k_{max}.
  - For each surface s for which a DLM $m_{s,\Psi}$ exists (the $m_{s,\Psi}$ of Ψ are stored in a list for that):
    - Set the texture that contains $m_{s,\Psi}$ as current drawing texture if it is not currently set.
    - Set the material of s as current drawing material if it is not currently set.
    - Draw s only in those pixels where it is visible according to the z-buffer, illuminated by the directional light source, and modulated by $m_{s,\Psi}$, and add the resulting fragments to the color buffer.

Note that the number of state switches is minimized, because usually many DLMs are drawn before the direction of the light source, the current texture, or the current material has to be changed. This is important for the performance, because state switches cause expensive stalls of the hardware rendering pipeline. If each surface would have its individual set of light directions, then the light source direction would have to be changed for each DLM, which would cause very many stalls.

4 Results

Figure 1 shows a snapshot of an interactive walkthrough that has been rendered with directional light maps. A large part of the scene consists of soft glossy surfaces with specular exponent 10 (floor, table, chairs). Only a small part of the scene (ceiling, back wall) receives direct illumination, the rest of the scene is completely indirectly illuminated.

We have used n_{Ω}=14 and k_{max}=2. The directional light maps have required 3.5 MB texture memory. In total 77,000,000 light paths (3,500,000 light paths per CPU) have been shot in a parallel implementation of the photon tracing preprocessing step, which took 7.2 minutes on a cluster of 11 PCs with dual 1 GHz Pentium3s in a 100 MBit Ethernet network. In this parallel implementation each CPU generates directional light maps for the whole scene. The directional light maps from all CPUs are then accumulated to achieve the final directional light maps which are used in the interactive walkthrough.

During the interactive walkthrough the scene has been rendered at a frame-rate of 44-57 Hz on a PC with a 900 MHz Thunderbird CPU and a GeForce2 GTS graphics card with 32 MB frame buffer and texture memory under OpenGl.

Fig. 1: Snapshot of a walkthrough in a globally illuminated soft glossy scene that has been rendered at interactive frame-rates with directional light maps.

Acknowledgments

This work has been supported by the Austrian Science Fund (FWF) project no. P13600-INF.

References

[Arv86] J. Arvo. *Backward Ray Tracing. Developments in Ray Tracing.* Siggraph 86 course notes

[BH+99] R. Bastos, K. Hoff, W. Wynn, A. Lastra. *Increased Photorealism for Interactive Architectural Walkthroughs*. Symposium on Interactive 3D Graphics 1999

[GG+96] S. J. Gortler, R. Grzeszczuk, R. Szeliski, M. F. Cohen. *The Lumigraph*. Siggraph 96 p. 43

[Hei99] W. Heidrich, H.-P. Seidel. *Realistic, Hardware-accelerated Shading and Lighting*. Siggraph 99 p. 171

[Hei01] W. Heidrich. *Interactive Display of Global Illumination Solutions for Non-diffuse Environments - A Survey*. Computer graphics forum vol. 20 no. 4 p. 225, 2001

[LH96] M. Levoy, P. Hanrahan. *Light Field Rendering*. Siggraph 96 p. 31

[MRP98] G. Miller, S. Rubin, D. Ponceleon. *Lazy decompression of surface light fields for precomputed global illumination*. Eurographics workshop on rendering 1998 p. 281

[NSI99] K. Nishino, Y. Sato, K. Ikeuchi. *Eigen-Texture Method*. IEEE Computer Vision and Pattern Recognition 1999 vol. 1 p. 618

[Shi92] P. Shirley. *Nonuniform Random Point Sets via Warping*. Graphics Gems III p. 80, 1992

[SSS95] M. Stamminger, P. Slusallek, H.-P. Seidel. *Interactive Walkthroughs and Higher Order Global Illumination*. Modeling, Virtual Worlds, Distributed Grapics p. 121, 1995

[SS+00] M. Stamminger, A. Scheel, X. Granier, F. Perez-Cazorla, G. Drettakis, F. Sillion. *Efficient Glossy Global Illumination with Interactive Viewing*. Computer Graphics Forum vol. 19 no. 1 p. 13, 2000

[SB97] W. Stürzlinger, R. Bastos. *Interactive Rendering of Globally Illuminated Glossy Scenes*. Eurographics workshop on rendering 1997 p. 93

[Stü98] W. Stürzlinger. *Calculating Global Illumination for Glossy Surfaces*. Computers & Graphics vol. 22 no. 2-3 p. 175, 1998

[WH+97] B. Walter, P. M. Hubbard, P. Shirley, D. Greenberg. *Global Illumination Using Local Linear Density Estimation*. ACM Transactions on Graphics vol. 16 no. 3 p. 217, 1997

[WA+97] B. Walter, G. Alppay, E. Lafortune, S. Fernandez, D. P. Greenberg. *Fitting Virtual Lights For Non-Diffuse Walkthroughs*. Siggraph 97 p. 45

[WA+00] D. N. Wood, D. I. Azuma, K. Aldinger, B. Curless, T. Duchamp, D. H. Salesin, W. Stuetzle. *Surface Light Fields for 3D Photography*. Siggraph 2000 p. 287