Thorsten Roth
Institute of Visual Computing
Bonn-Rhein-Sieg University of Applied Sciences
thorsten.roth@h-brs.de

Lightfield Information for Fast Previews of Global Illumination

Introduction

Although around since the 1980s, global illumination is still one of the major challenges in computer graphics. Global illumination tries to realistically compute the (physically-based) distribution of light in a virtual scene. This scene is given by its geometrical description and according material properties. Light distribution is generally described by the rendering equation:

$$L_o(p,\omega_o) = L_e(p,\omega_o) + \int_\Omega L_i(p,\omega_i)\, f(p,\omega_i,\omega_o)(\omega_i \cdot n_p)\, d\omega_i$$

As this is a high-dimensional integral generally impossible to compute analytically, numerical methods from the area of Monte Carlo Simulation are often employed to achieve results accurate enough for display. This means that a vast amount of random light rays is shot through the scene in order to determine the light energy arriving at the virtual camera. However, depending on the specific scene, this process can still take hours for a single frame even on high-end hardware because of complex effects such as caustics.

Goals, Methods and Challenges

The main idea here is to employ light field data measured throughout the rendering process for approximate visualizations of the virtual scene. This means that at each ray-object intersection found during the rendering process, the directional energy at this point is stored. However, during a Monte Carlo process, this is not yet the integrated energy leaving a point into this direction, as realistic materials are not perfectly specular and reflect energy coming from one direction to a set of outgoing directions. Instead, there may be thousands of rays for one single direction, as each of these rays belongs to one specific path that has been sampled.

Thus, the data achieved during the rendering process still has to be processed by a clustering algorithm, merging similar positions and directions in a way accounting for specific material properties such as reflectivity. By doing this, the amount of data should also be greatly reduced.

Reconstruction of the illumination is currently done by creating a g-buffer consisting of geometric information about the primary rays' hit points and finding the best matches in the light field data. This is done using an approximate k-nearest-neighbor search where the discovered neighbors are weighted according to their respective distance from the queried point.

The main challenge is to design an algorithm which yields a performance high enough to quickly be able to preview a globally illuminated scene from dynamic perspectives without explicitly rendering all of them with the global illumination algorithm itself. Afterwards, to get a high-quality image, an error measure has to be defined in order to determine which pixels already give a good approximation for the final image and which ones need more computing time (i.e., more samples) for satisfactory results.

Acknowledgements

This work is supported by a scholarship of the Institute of Visual Computing at the Bonn-Rhein-Sieg University of Applied Sciences.

Supervisors

Dr. Yongmin Li
Brunel University
London

Prof. Dr. André Hinkenjann
Bonn-Rhein-Sieg University of Applied Sciences
Sankt Augustin

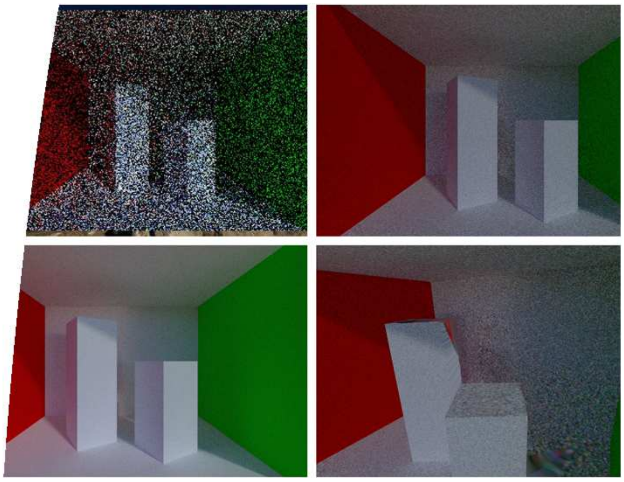

Figure 1: Top left: Approx. 230 images with a resolution of 320x240 pixels have been generated while moving the camera around horizontally and vertically in front of the Cornell Box; **Top right:** 1680x1050 image reconstructed from this data (64 neighbors); **Bottom right:** 1680x1050 reconstruction of a different perspective with visible artifacts due to sparse sampling (64 neighbors); **Bottom left:** Scene rendered in 1680x1050 with 64 samples per pixel

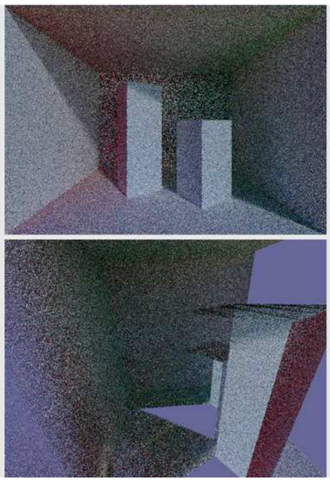

Figure 2: OpenGL-based Point Cloud Visualization of the generated data; **Top:** Front view, well-sampled; **Bottom:** View from inside the box; It is visible that the area behind the box has not been sampled at all due to the perspective used during sample generation; note that the missing colors of the walls are due to texture colors not being multiplied into the light field values. This is done during the reconstruction process in order to preserve detail in textured areas.

BRSU
Graduate Institute

Graduate Institute
gi@h-brs.de
www.gi.h-brs.de

INSTITUTE OF
Visual Computing

Hochschule
Bonn-Rhein-Sieg
University of Applied Sciences