Precomputed Light Sets for Fast High Quality Global Illumination

Johannes Günther, Ingo Wald, and Hans-Peter Seidel
MPI für Informatik, Saarbrücken, Germany*

Figure 1: Overview of our method demonstrated on the OFFICE scene: a) Virtual Point Lights (VPLs, yellow) are generated by a random walk starting from the light sources, approximating global illumination. Then a set of Surface Sampling Points (SSPs, blue and red) is generated, each of them storing precomputed visibility information. b) Using only the precomputed illumination of the nearest SSP during rendering leads to artifacts: The voronoi regions are clearly visible. c) *Evaluating* the VPLs using the precomputed visibility information results in fast computation of high quality global illumination images. d) The SHIRLEY 10 scene with area 100 lights and about 15000 VPLs raytraced in 11.3 seconds on a *single* PC (dual Opteron 2.4 GHz) at video resolution, including soft shadows and indirect illumination

1 Introduction

Instant Radiosity [Keller 1997] is an elegant algorithm to approximate global illumination by a number of *Virtual Point Lights (VPLs)* that are generated by a random walk starting from the light sources (Figure 1a). For complex settings however hundreds or even thousands VPLs are required and all these VPLs need to be evaluated and their visibility calculated during computation of each pixel in order to minimize artifacts and to obtain smooth shadows. Ray tracing is well suited for the visibility determination [Wald et al. 2002] as the high number of lights complicate the use of shadow maps due to the huge memory consumption.

In this sketch we propose a technique to significantly reduce the number of VPLs to be considered for each pixel by determining the visibility and illumination information in a preprocessing step. Using our method we can quickly produce high quality global illumination renderings even from scenes with complex illumination.

2 Local Light Sets

We divide the scene into small cells and evaluate the set of VPLs seen from each cell in a preprocessing step. The cells are defined as the intersection of the scene geometry and voronoi regions to a set of Surface Sample Points (SSPs). This approach is very flexible in means of supported surface descriptions: polygonal, implicit (NURBS, ISO Surfaces) and fractal descriptions can be covered.

The SSPs by itself are generated by the intersection points from uniformly distributed rays with the scene geometry (Figure 1a). These points are then sorted into a *kd*-tree for fast retrieval during rendering when we need to get the corresponding cell to a point to be shaded and its information.

Supersampling: Taking only the SSP of one cell to sample the VPLs will obviously result in undersampling artifacts e.g. if a shadow border from one VPL crosses the cell and the SSP is in the shadow that VPL will be wrongly regarded as not visible. To compensate for such errors we sample the VPLs from more points within each cell. This approach has the additional advantage that we can now also determine the variance in the visibility of each VPL used to classify the VPLs into two groups: VPLs that are visible from every sample point (fully visible VPLs) and VPLs that are visible from some points (partly visible VPLs). During rendering we can skip the visibility determination for fully visible VPLs while partly visible VPLs are still sampled. This allows for making the cells relatively big (their projection can easily cover many pixels) and nevertheless capture fine shadow details (Figure 1b and 1c).

Importance Sampling of VPLs: Especially in scenes with little occlusion such as the OFFICE scene (Figure 1c) the number of fully visible VPLs of each cell can be very high. To evaluate their combined contribution efficiently we can apply importance sampling, which is particularly suited as we already know the expected contribution of each VPL from the preprocessing step. Consequently taking the expected contribution as a probability density function for importance sampling leads to a very low variance.

3 Results and Conclusion

For the OFFICE scene we generated 1430 VPLs from 3 primary (area) light sources and 50k cells with 25 sample points on average. With a single PC (dual Opteron 2.4 GHz) this scene can be rendered in 12.4 seconds at a resolution of 640×480. Using the preprocessing information reduces the number of shadow rays by a factor of 17 compared to evaluating all VPLs. In scenes with high occlusion (Figure 1d) the benefit is even higher: With about 15k VPLs from 100 primary lights and 150k cells the number of shadow rays are reduced by a factor of 69 and the render time by a factor of 45.

Up to now we demonstrated our method only with diffuse surfaces. However the big advantage is that we still evaluate the VPLs and thus the incoming light directions at each point to be shaded are known. As a consequence it should be easy to integrate non-diffuse BRDFs, as preliminary results suggest.

References

KELLER, A. 1997. Instant Radiosity. *Computer Graphics (Proceedings of ACM SIGGRAPH)*, 49–56.

WALD, I., KOLLIG, T., BENTHIN, C., KELLER, A., AND SLUSALLEK, P. 2002. Interactive Global Illumination using Fast Ray Tracing. *Rendering Techniques*, 15–24. (Proceedings of the 13th Eurographics Workshop on Rendering).

*e-mail: {guenther,wald,hpseidel}@mpi-inf.mpg.de