# Meshless Hierarchical Radiosity on the GPU

{Michael.Zollhoefer, Marc.Stamminger}@informatik.uni-erlangen.de

LGDV
COMPUTER GRAPHICS • ERLANGEN

## Abstract

Meshless radiosity [1, 2] is a radiosity method that is based on a point-based and hierarchical discretization of the scene. This better decouples the runtime complexity from the geometric complexity of the scene and allows for an adaptive high-quality simulation of the diffuse global light transport. In this paper, we analyze the bottlenecks of this approach and examine the possibilities for an efficient and parallel implementation of this paradigm on the GPU. We show how by modifying the hierarchical data structures and the computation of the transport operator, a highly efficient GPU-based solution can be achieved which is by orders of magnitude faster and allows to compute high-quality global illumination solutions within seconds.

## Parallelization of the Transport

The adaptive computation of the transport operator starts by computing gathering integrals at the coarsest level basis functions. Thereafter, an oracle decides if further refinement is required.

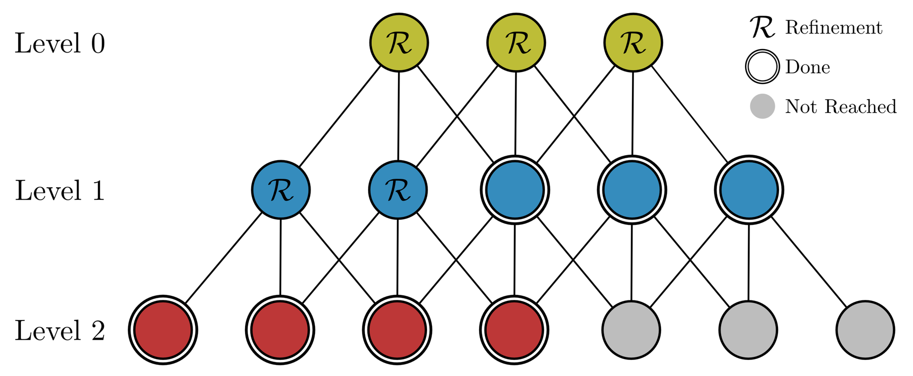


We use a breadth-first traversal of the hierarchy and compute all gathering integrals on common hierarchy levels in parallel. A parallelization across multiple levels is not possible, because the refinement oracle introduces a dependency to the previous level.

## Implementation of the Transport

All gathering rays in the Monte Carlo integration step and all per hit-point computations are independent.

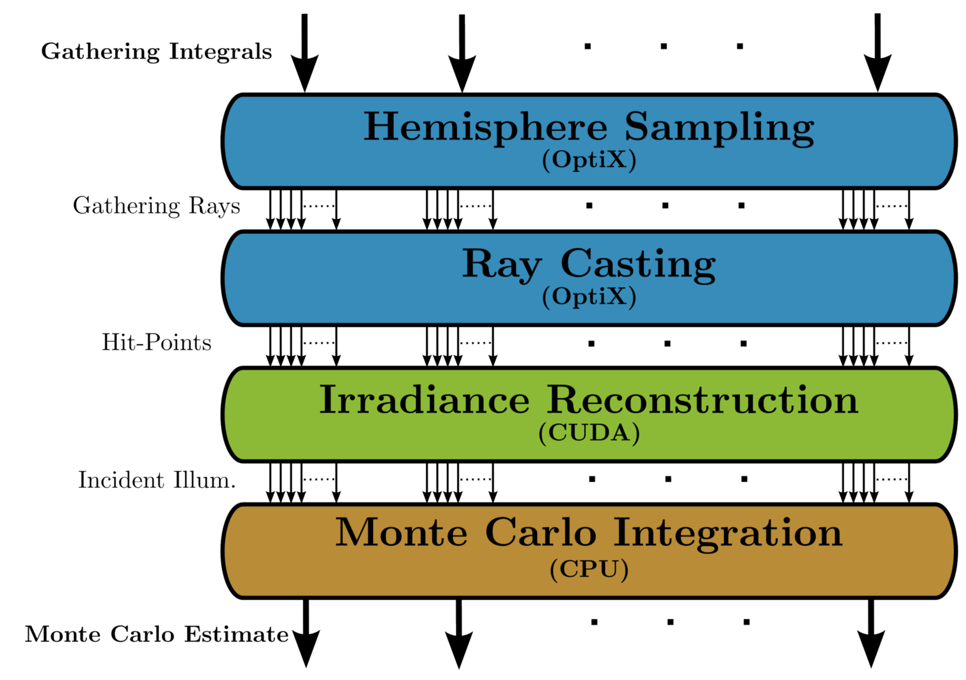


This allows to compute the transport using the presented GPU-based pipeline.

## Meshless Hierarchy

A meshless hierarchy [1, 2] consisting of $m$ levels allows to represent illumination at a point $\mathbf{p}$ in the following way:

$$F_m(\mathbf{p}) = \sum_{l=0}^{m-1} \sum_{j=0}^{N_l} \alpha_j^l \cdot B_j^l(\mathbf{p}).$$

Thereby, the $l$-th level is given by $N_l$ basis functions $B_j^l$ and the coefficients $\alpha_j^l$ encode the illumination. Its point-based and hierarchical nature allows to decouple the runtime complexity from the geometric complexity of the scene and to locally adapted the resolution of the computations. The important design decision in the original work is to store absolute values on the coarsest hierarchy level and to use delta coefficients on all other levels.

## Modified Meshless Hierarchy

We store absolute values on all hierarchy levels. Therefore, the modified version of the hierarchy encapsulates multiple absolute Shepard Approximations of the illumination. This allows to eliminate the expansive computation of the energy transported to the parents which has to be subtracted in the original approach. To keep the absolute coefficients consistent, we propagate the gathered illumination through the hierarchy.

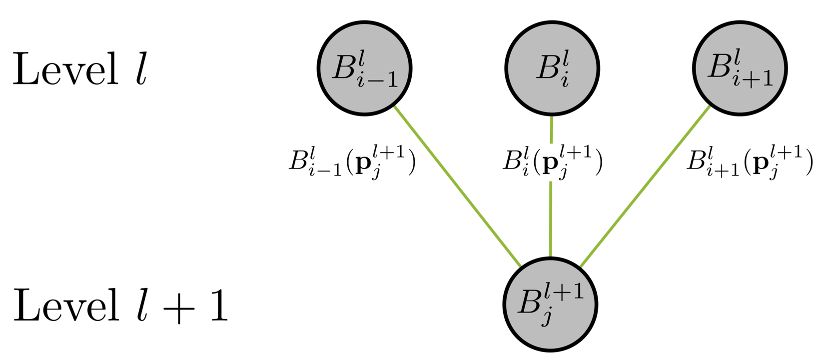


In addition, we apply an approximate reconstruction scheme which speeds up the computation of the light transport. This is similar to the approach used by Christensen [3] to speed up photon mapping.

## References

[1] Lehtinen J., Kontkanen M. Z. J., Sillion E. T. F. X., Aila T.: Technical Report TML-B7, Publications in Telecommunications Software and Multimedia, Helsinki University of Technology, Meshless Finite Elements for Hierarchical Global Illumination.

[2] Lehtinen J., Zwicker M., Turquin E., Kontkanen J., Durand F., Sillion F. X., Aila T.: A meshless hierarchical representation for light transport. *ACM Trans. Graph. 27*, 3 (2008), 1–9.

[3] Christensen P. H.: Faster photon map global illumination. *Journal of Graphics Tools 4* (1999), 1–10.

## Results

| | CPU Orig. | | | GPU Orig. | | | GPU Ours | | |
|---|---|---|---|---|---|---|---|---|---|
| | **Cornell** | **Happy** | **Sponza** | **Cornell** | **Happy** | **Sponza** | **Cornell** | **Happy** | **Sponza** |
| **1. Bounce** | 18.8s | 8m 19s | 11m 11s | 1.0s | 14.3s | 24.4s | 0.5s | 3.3s | 3.5s |
| **2. Bounce** | 8.7s | 2m 33s | 37.5s | 0.5s | 4.0s | 1.5s | 0.2s | 0.9s | 0.4s |
| **3. Bounce** | 2.7s | 39.0s | 22.6s | 0.2s | 1.3s | 0.9s | 0.1s | 0.4s | 0.3s |
| **Integrals** | 14k | 66k | 47k | 14k | 66k | 47k | 15k | 66k | 46k |
| **Total Time** | 34.6s | 12m 1s | 13m 19s | 2.1s | 20.7s | 29.6s | 1.1s | 5.3s | 5.1s |
| **Speedup** | 1x | 1x | 1x | 16.5x | 34.8x | 27.0x | 31.5x | 136.0x | 157.7x |

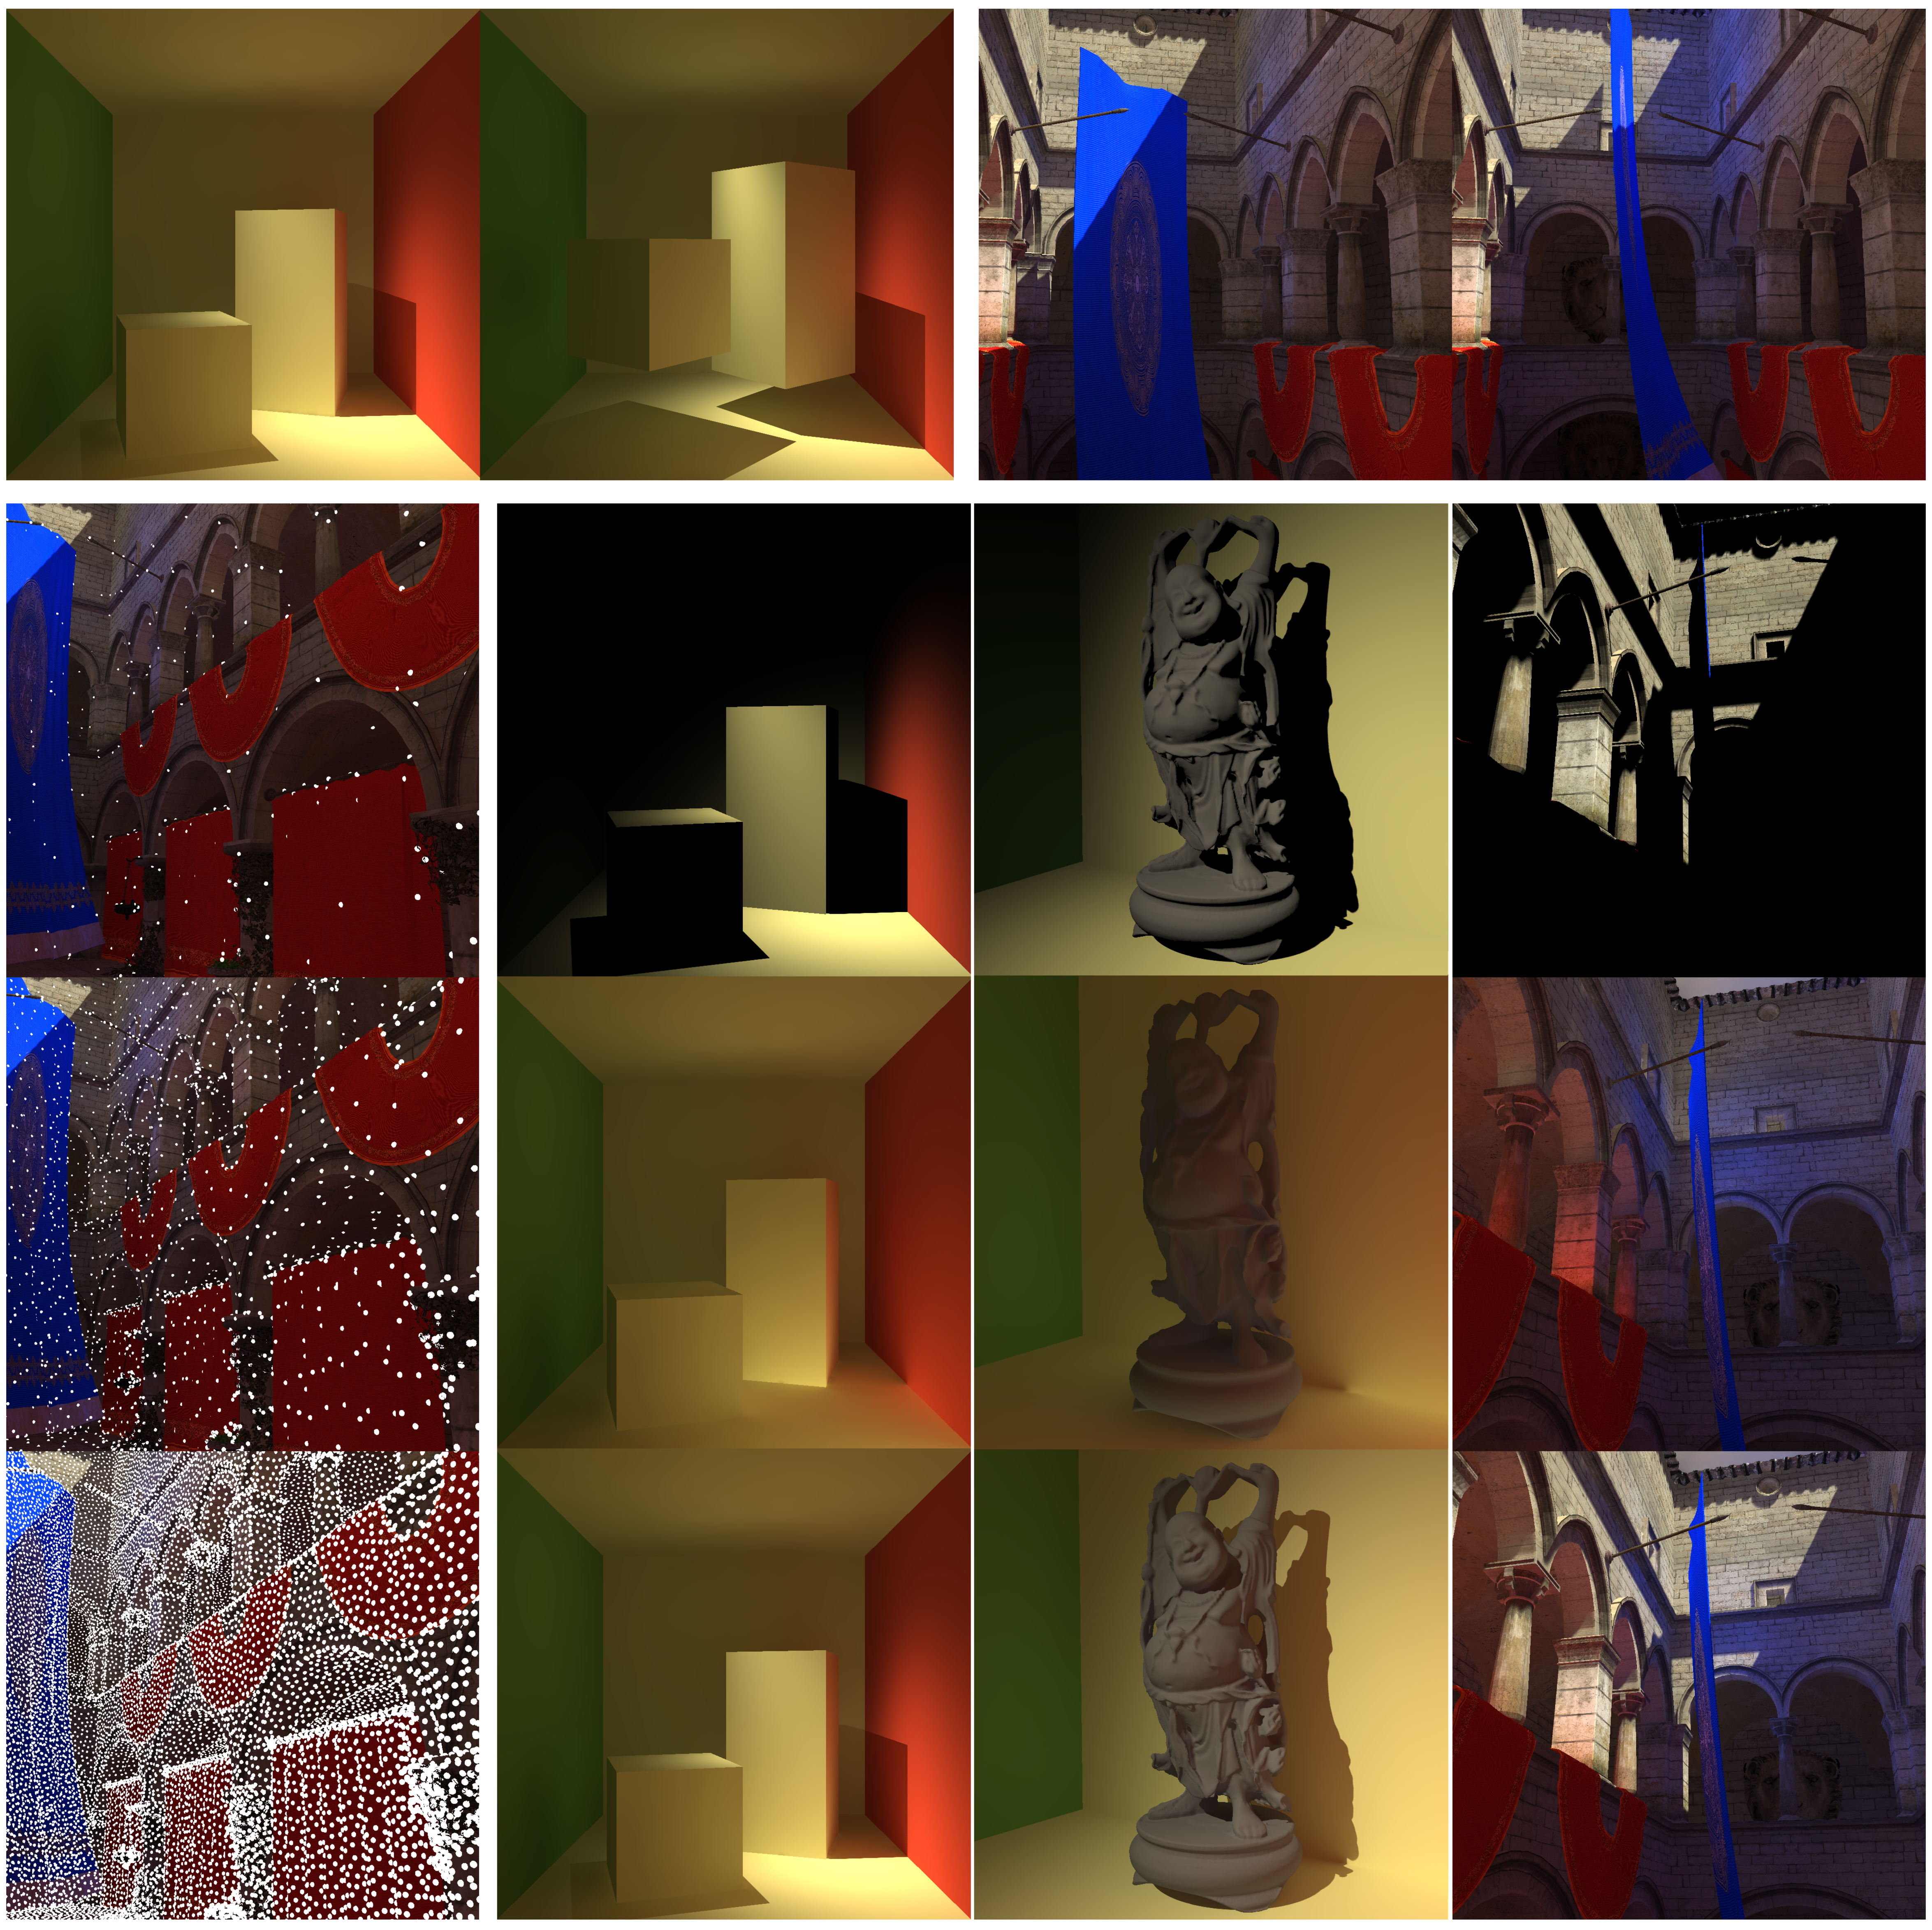

## Acknowledgements

We thank Cornell University (Cornell Box), Stanford University (Happy Buddha) and Crytek (Crytek Sponza) for making their models publicly available. This work was partly funded by the German Research Foundation (DFG) under grant STA-662/3–1.