# LightCluster – Clustering Lights to Accelerate Shadow Computation

Daniel Wiesenhütter, Andreas Klein, Alfred Nischwitz

Munich University of Applied Sciences

## Introduction

Shadows are an important part of a visualization and give the viewer supplemental details about the appearance of objects. In real-time rendering, shadow mapping [Williams 1978] is a popular approach to compute shadows. However, a shadow map must be computed for each light and thus, the memory and the computation time increases with the number of lights.

We present an approach, called LightCluster, to automatically select representative light sources and accelerate the computation of direct shadows for scenes with many lights. We carefully select light sources as cluster centers and cluster the remaining lights using a minimum distance metric [Wolfowitz 1957]. We represent each cluster by an area light source and use Percentage Closer Soft Shadows [Fernando 2005] to render soft shadows for each cluster.

In our implementation, we use omnidirectional point lights. However, the approach can be adapted for other light types, such as directional or spot lights.

## Related Work

**Real-Time Soft Shadows**

Shadow Mapping [Williams 1978] is a popular method to compute shadows in real-time rendering. The idea is to assume point light sources and to replace the visibility test by comparing depth values from the light's point of view and the observer's point of view. In order to compute shadows for omnidirectional lights, a cube shadow map can be rendered [Gerasimov 2004]. Percentage Closer Filtering (PCF) [Reeves et al. 1987] computes filtered hard shadows by making multiple shadow comparisons within a filter window. This idea is further extended by Fernando [Fernando 2005] with Percentage Closer Soft Shadows (PCSS) to realize shadows with variable sized penumbras. Instead of using a fixed filter per pixel, the filter window is scaled according to a penumbra size. The penumbra size can be estimated by calculating an average blocker depth and using similar triangles.

**Many-Light Methods**

Clustered Visibility [Dong et al. 2009] accelerates the computation of indirect shadows with Reflective Shadow Maps (RSM) [Dachsbacher and Stamminger 2005]. The method uses k-means clustering on the RSM to build clusters of VPL. They interpret each cluster as an area light in order to accelerate the visibility test.

The idea of [Dong et al. 2009] is closely related to our work. In contrast to Clustered Visibility we focus our work on high frequency shadows for direct lighting. As the lights are not distributed by a RSM, a k-means clustering may lead to errors. A central cluster position can result in a representative light source that is occluded by geometry, i.e. located within walls. Our approach uses an existing light source as a cluster center and clusters the remaining lights with a minimum distance metric.

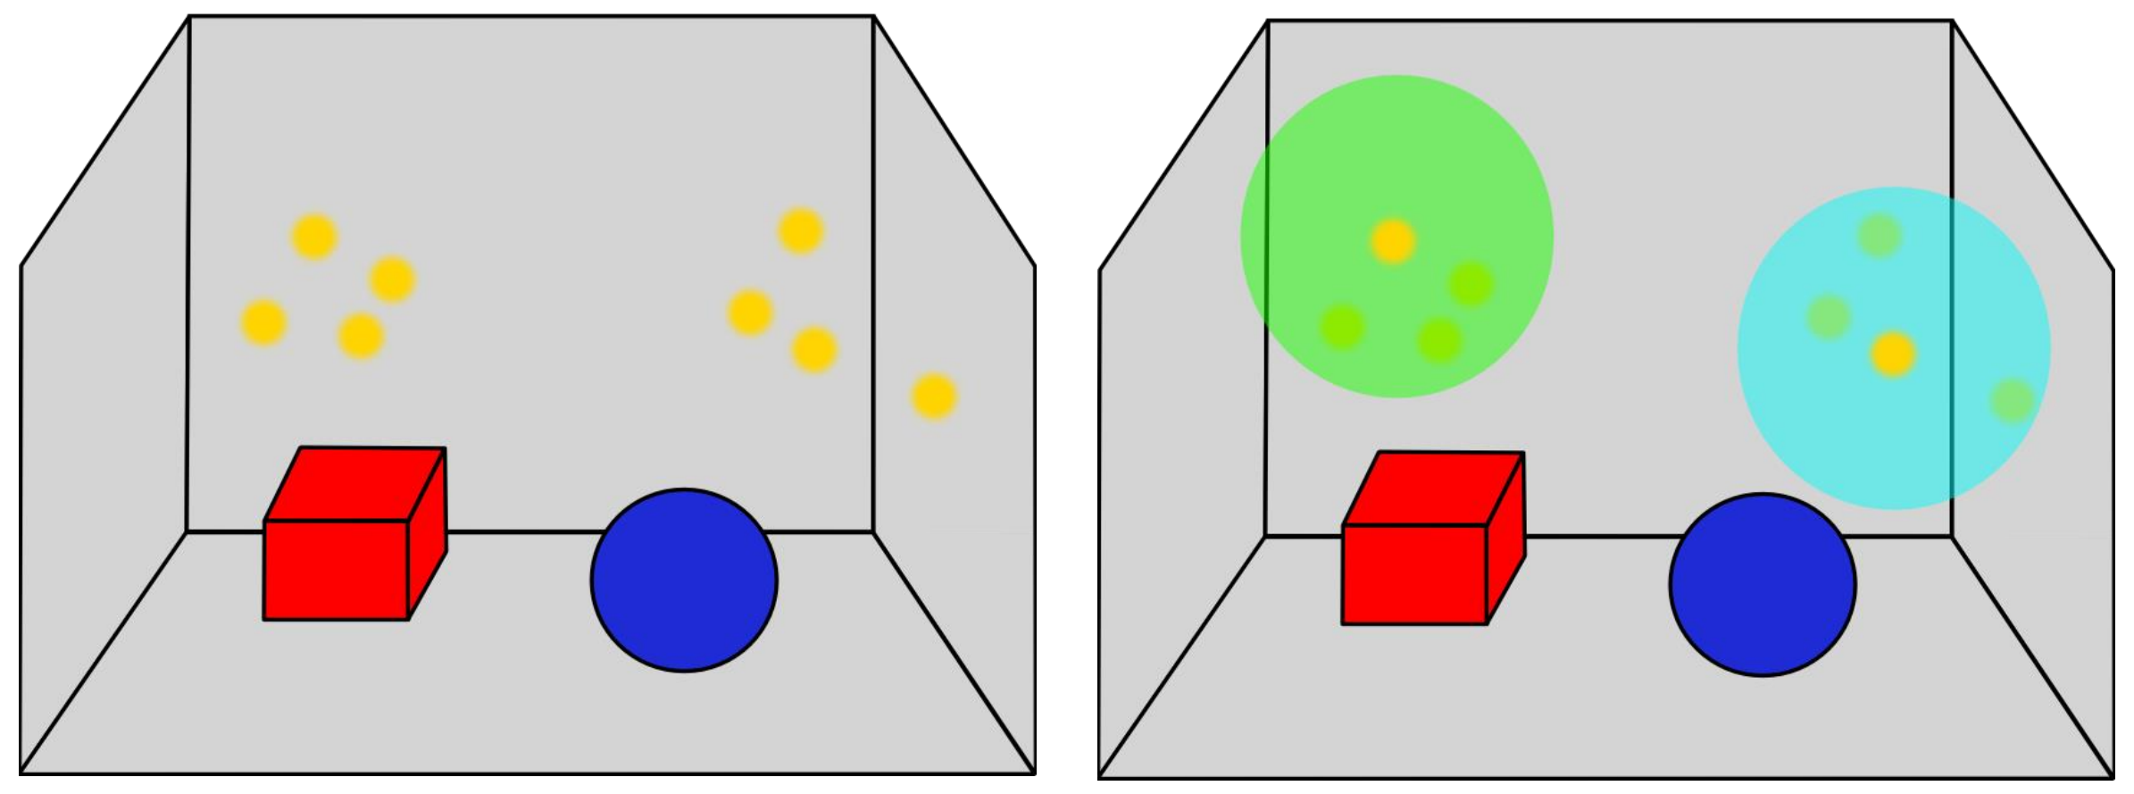

## Minimum-Distance-Clustering

The minimum-distance-cluster algorithm [Wolfowitz 1957] is a hierarchical clustering algorithm and clusters all objects within a user defined minimum distance. The algorithm can be realized in a top-down strategy and proceeds as follows. First, it is assumed that all objects are within a single cluster. The algorithm iterates over all objects of a cluster and splits them into new clusters, if their distance exceeds the defined minimum distance. This step is repeated until no more clusters can be split.

## Our Approach

In order to reduce the error in shadows, we perform a two pass clustering with different metrics in each pass. In addition, the point lights are culled against the viewing frustum. We first select point lights as cluster centers by using the light range and a minimum-distance metric, which is scaled by the camera distance. This allows us to generate smaller clusters and thus, more shadow maps, near the camera position. In the second clustering pass, the remaining point lights are assigned to the nearest cluster centers. Therefore, the error in shadows due to the decreased amount of shadow maps can be reduced.

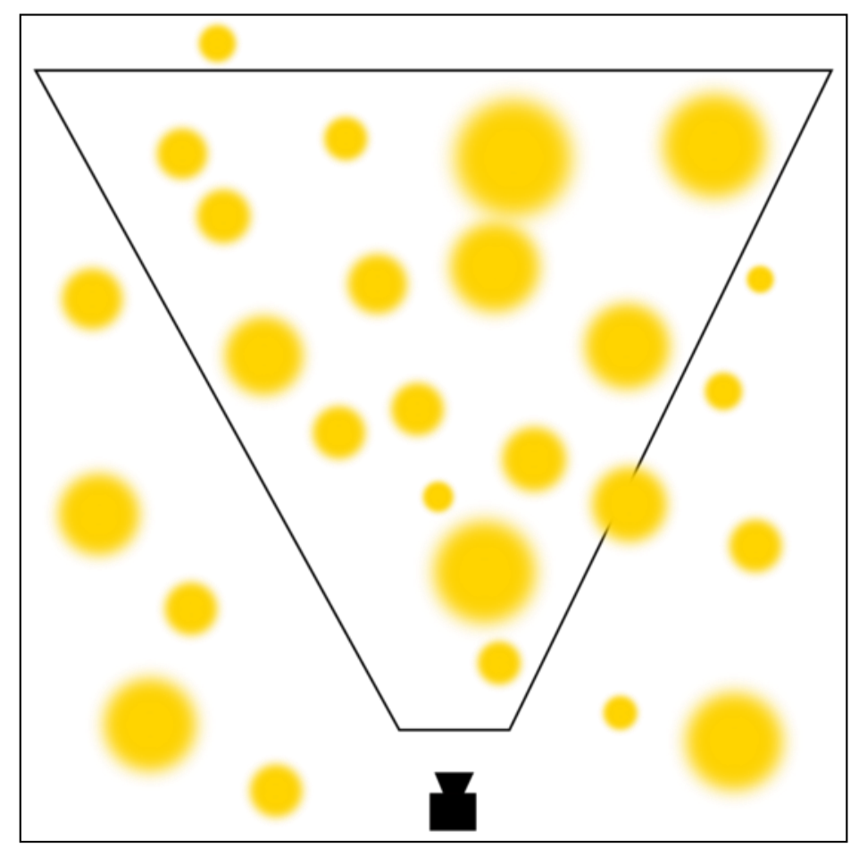

(a) point light distribution

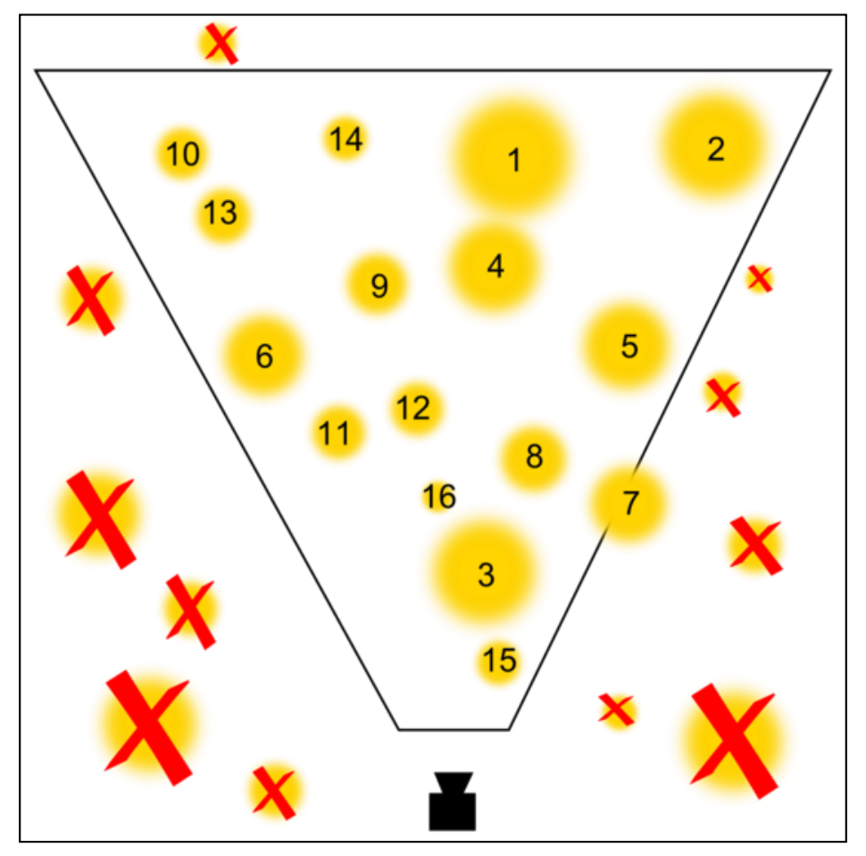

(b) frustum culling of point lights

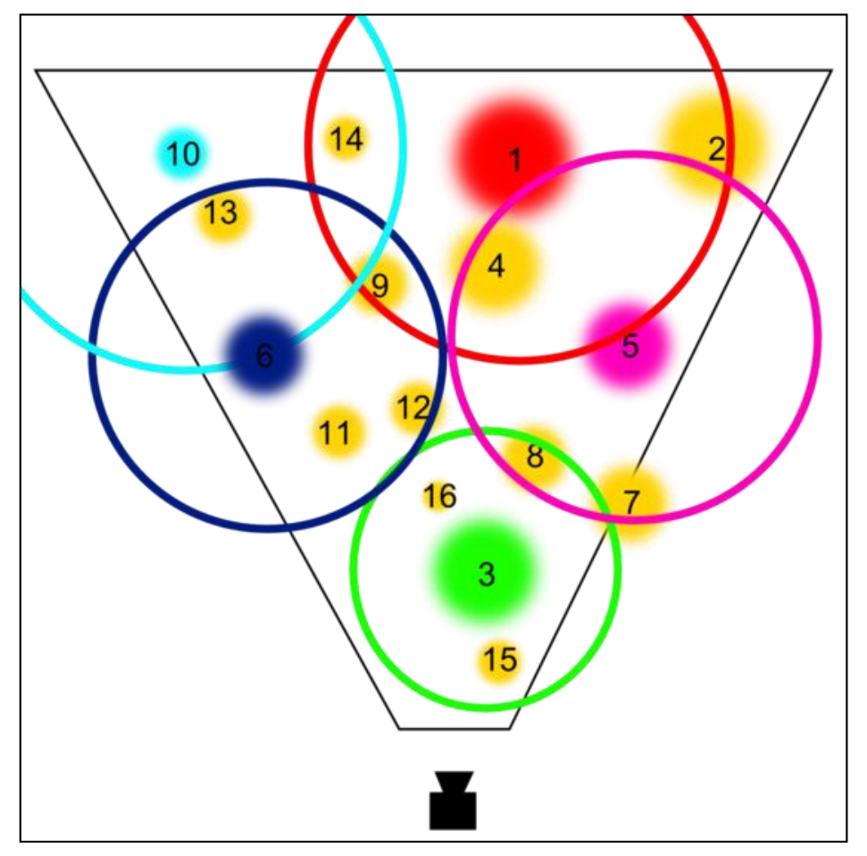

(c) selecting clusters with a minimum distance metric

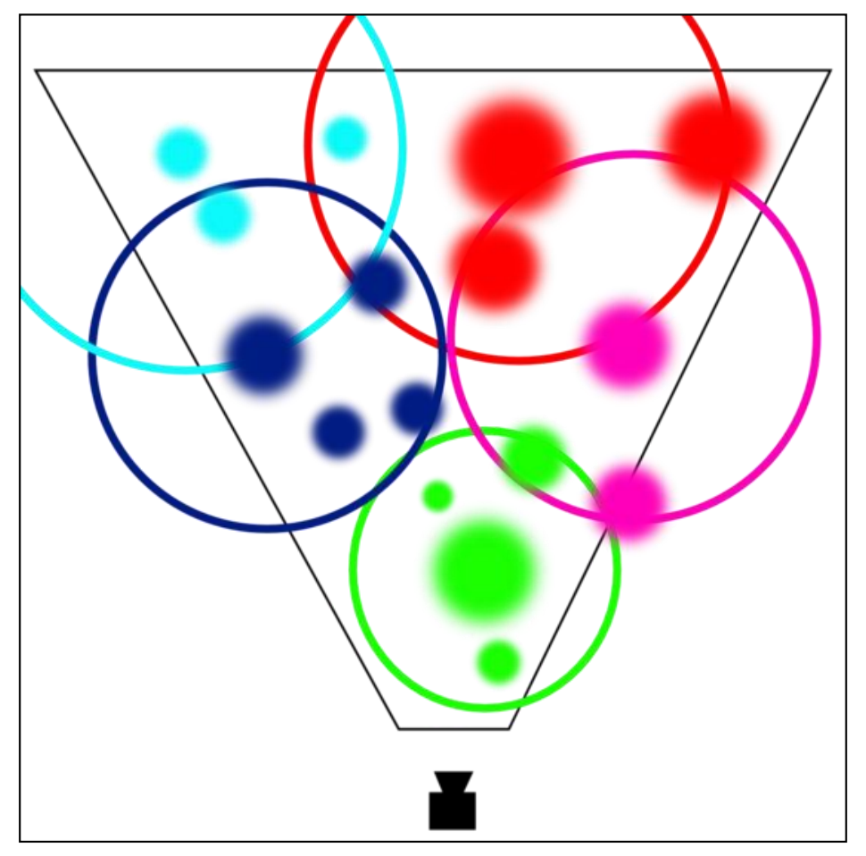

(d) assigning lights to nearest clusters

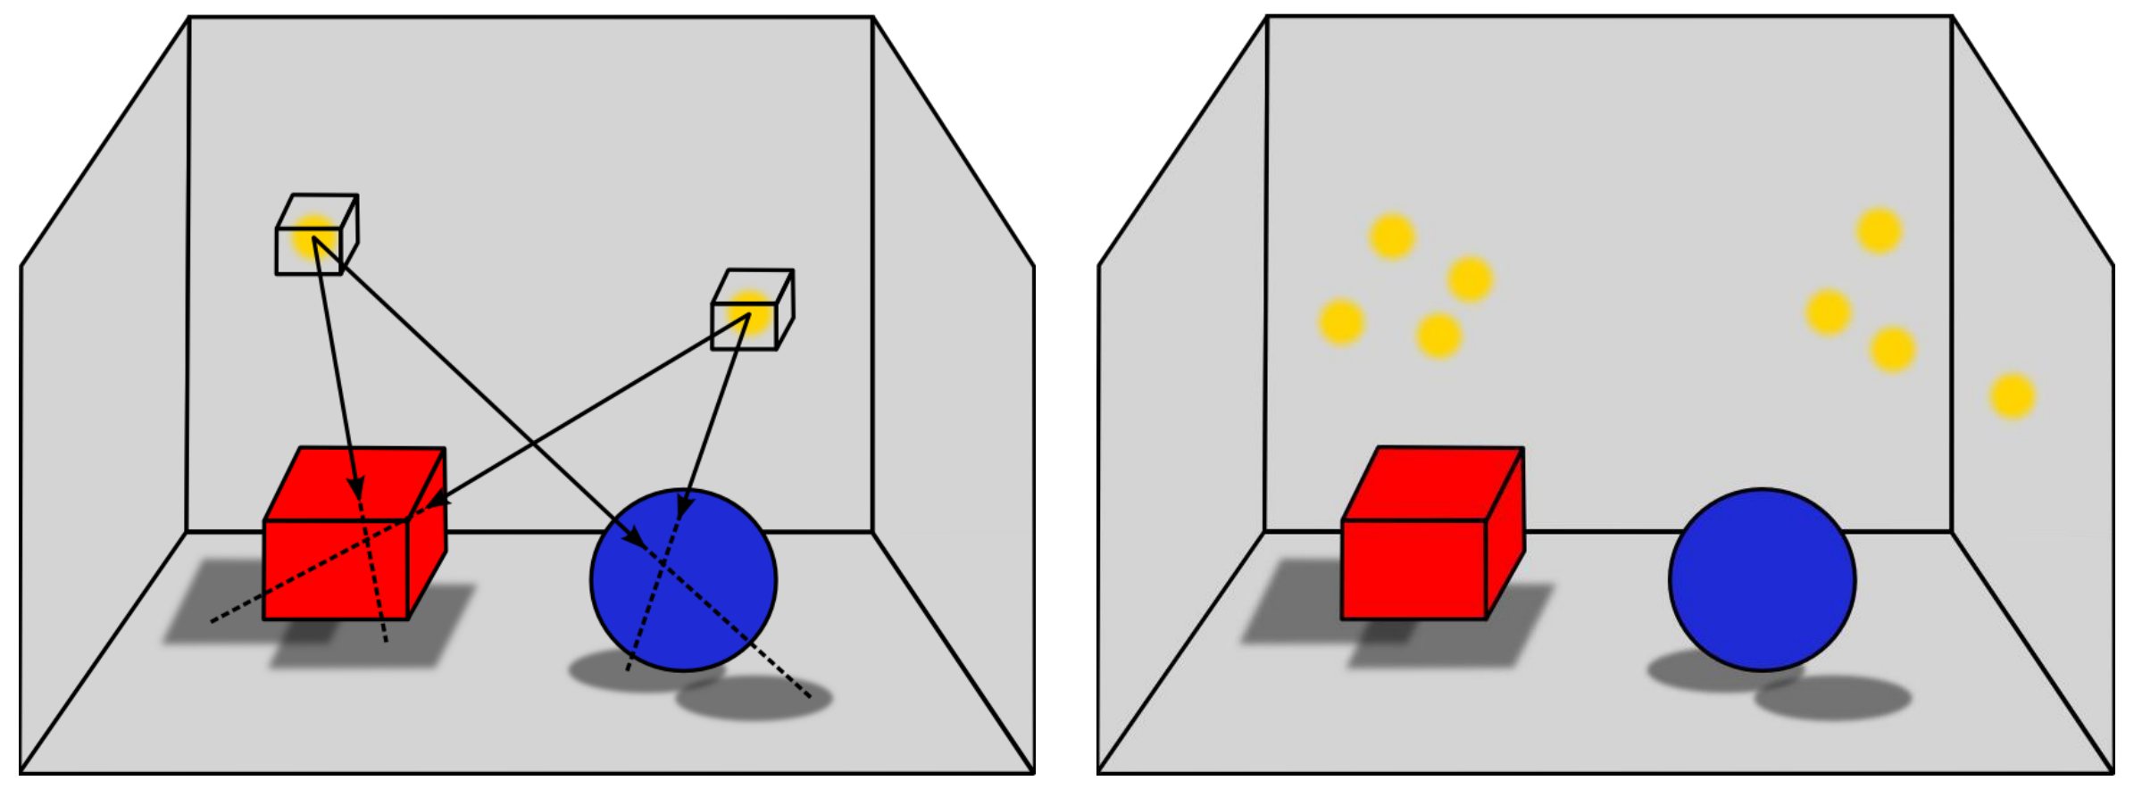

## Compute Shadows

We render a cube shadow map for each cluster and interpret the cluster as a disc-shaped area light source. The radius of the area light source is given by the minimum distance of the cluster. We use this radius to scale the filter window of PCSS. The visibility factor is then stored in a texture for each cluster. This allows us to calculate the shadows iteratively and reduces the texture memory from a cube shadow map to a screen sized texture per cluster. In this way, we sample the cube shadow maps only once for each cluster and avoid additional PCSS sampling for each light source during shading. After the shadow computation, we use the set of visibility textures for shading. We shade the scene with each point light source and modulate the resulting color with the visibility value stored in the texture for the point light's cluster.

## Results

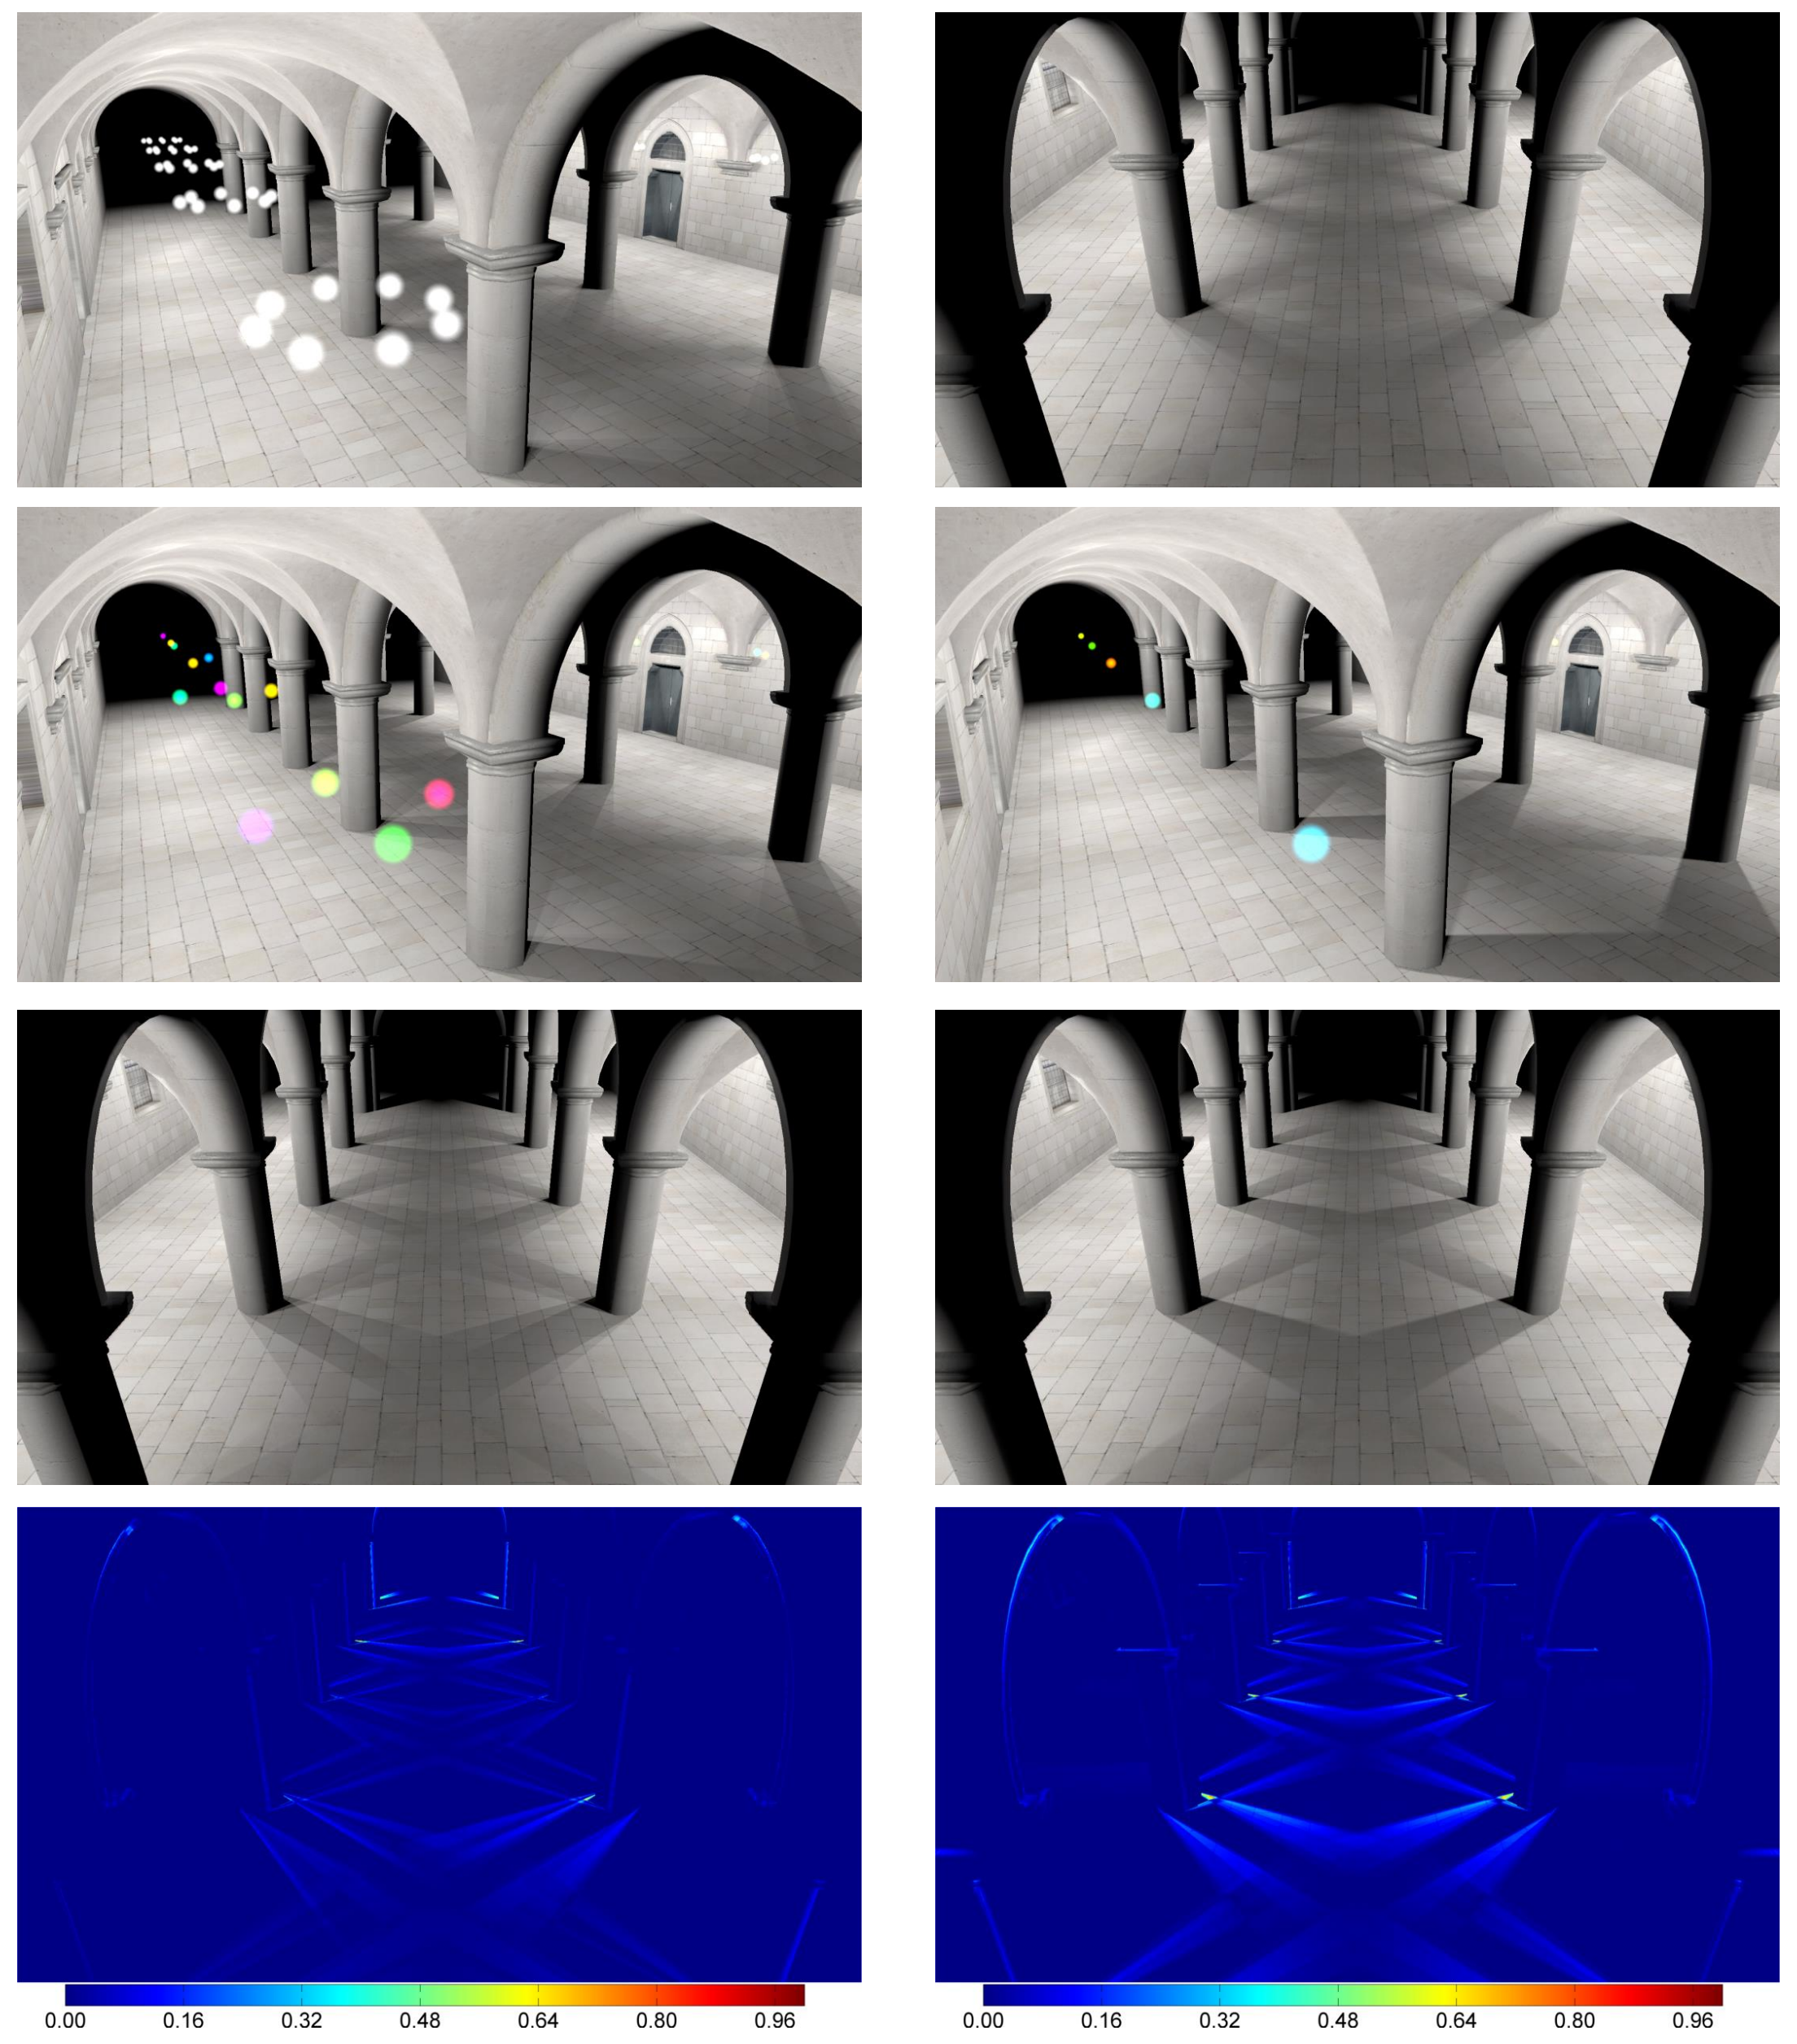

| Dabrovic Sponza Scene | Time (ms) |
|---|---|
| LightCluster with 26 clusters | 43.1 |
| LightCluster with 10 clusters | 20.7 |
| Reference with 80 point lights | 97.8 |

## Results

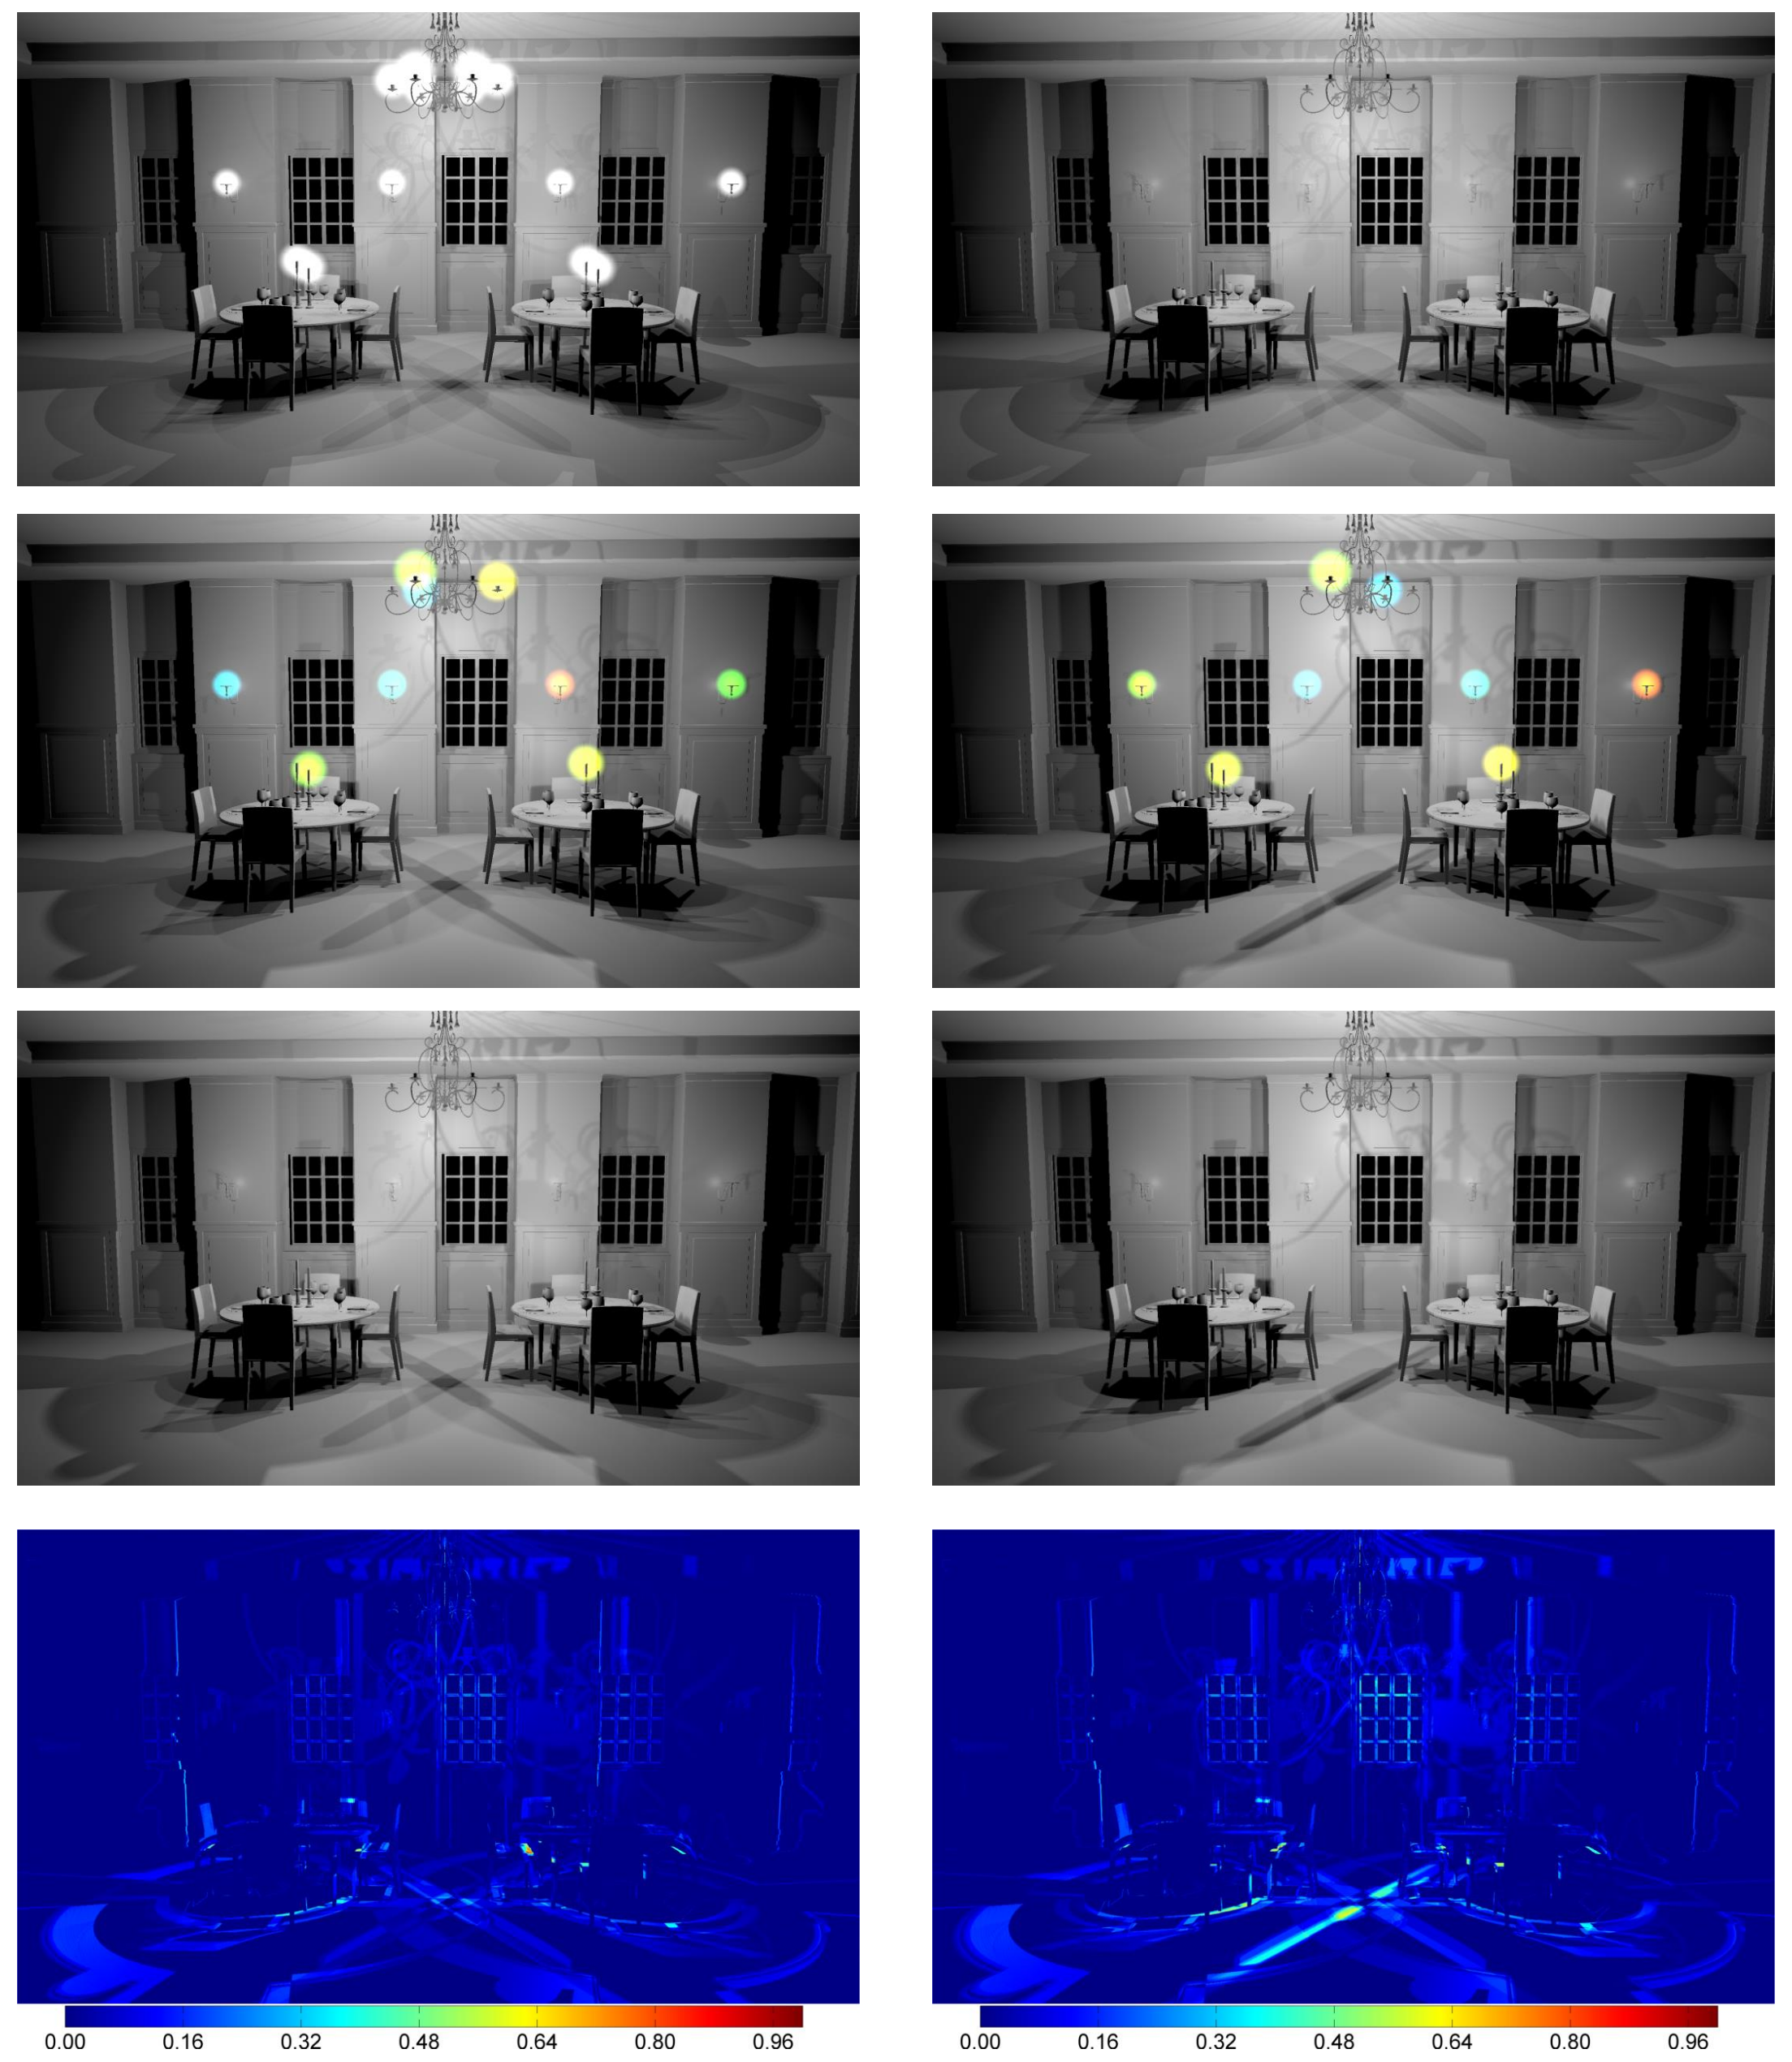

| Restaurant Scene | Time (ms) |
|---|---|
| LightCluster with 9 clusters | 32.9 |
| LightCluster with 8 clusters | 30.5 |
| Reference with 14 point lights | 39.3 |

## References

DACHSBACHER C., STAMMINGER M. Reflective Shadow Maps. In Conf.proc. I3D 2005, 203-231, 2005.

DONG, Z., GROSCH, T., RITSCHEL, T., KAUTZ, J., AND SEIDEL, H.-P. 2009. Real-time indirect illumination with clustered visibility. In Proc. of the VMV, 187 – 196.

FERNANDO, R. 2005. Percentage-closer soft shadows. In ACM SIGGRAPH 2005 Sketches, ACM, New York, NY, USA, SIGGRAPH '05.

GERASIMOV P. S. Omnidirectional Shadow Mapping, GPU Gems, Addison-Wesley, 193-203, 2004.

REEVES W. T., SALESIN D. H., COOK R. L. Rendering antialiased shadows with depth maps. In Conf.proc. SIGGRAPH '87, ACM, 283-291, 1987.

WILLIAMS, L. 1978. Casting curved shadows on curved surfaces. SIGGRAPH Comput. Graph. 12, 3 (Aug.), 270–274.

WOLFOWITZ, J. 1957. The minimum distance method. Annals of Mathematical Statistics 28, 1, 75–88.

Scenes downloaded from:
Dabrovic Sponza: http://graphics.cs.williams.edu/data
Restaurant Scene: http://idst-render.com/scenes.html