

Implementing Lighting Grid Hierarchy for Self-illuminating Explosions

Can Yuksel
DreamWorks Animation

Cem Yuksel
University of Utah

Introduction

This document includes some implementation related details that are omitted in the paper and practical considerations. While the implementation details presented here can be deduced from the information provided in the paper, we include them here for completeness and aiding an implementation oriented reader.

Lighting Grid Hierarchy

For building the lighting grid hierarchy, we distribute the intensity $\mathbf{I}_{0,j}$ of each light in \mathbb{S}_0 to the vertices of the corresponding grid cell using trilinear weights w_{ij} of grid vertices i of level ℓ . The trilinear weights can be written as

$$w_{ij} = \max(0, \min(1, \langle \mathbf{1} - |\mathbf{p}_{0,j} - \mathbf{q}_{\ell,i}|/h_{\ell} \rangle)), \quad (1)$$

where $\mathbf{q}_{\ell,i}$ is the position of the grid vertex i at level ℓ , the operation $|\cdot|$ denotes the absolute values of the vector components, and the operation $\langle \cdot \rangle$ provides the product of the vector components.

During lighting estimation we use blending functions to combine the contributions from each level of the hierarchy. Let $r_{\ell} = \alpha h_{\ell}$ be the influence radius of level ℓ , determined by a user-defined parameter α . We define the blending functions as

$$B_{\ell}(d) = \begin{cases} 0, & \text{if } d \leq \frac{r_{\ell}}{2}, \\ \frac{2d}{r_{\ell}} - 1, & \text{if } \frac{r_{\ell}}{2} < d \leq r_{\ell}, \\ 2 - \frac{d}{r_{\ell}}, & \text{if } r_{\ell} < d \leq 2r_{\ell}, \\ 0, & \text{if } 2r_{\ell} < d. \end{cases} \quad (2)$$

This blending function linearly increases from zero at $\frac{r_{\ell}}{2}$ to 1 at r_{ℓ} , and linearly decreases back to zero at $2r_{\ell}$. The two exceptions are the blending functions for the first level

$$B_0(d) = \begin{cases} 1, & \text{if } d \leq r_0, \\ 2 - \frac{d}{r_0}, & \text{if } r_0 < d \leq 2r_0, \\ 0, & \text{if } 2r_0 < d. \end{cases} \quad (3)$$

and the last level

$$B_{\ell_{max}}(d) = \begin{cases} 0, & \text{if } d \leq \frac{r_{\ell}}{2}, \\ \frac{2d}{r_{\ell}} - 1, & \text{if } \frac{r_{\ell}}{2} < d \leq r_{\ell}, \\ 1, & \text{if } r_{\ell} < d. \end{cases} \quad (4)$$

This way, the influence region of the last level does not diminish for distances beyond $r_{\ell_{max}}$.

Shadow Maps for Lighting Grids

For efficiently approximating the filtered density values at any point, we can pre-filter the density field. Let ρ_0 be the input density field from the simulation data. We generate the filtered density field ρ_{ℓ} (with the same resolution as ρ_0) using a pyramidal convolution filter of size h_{ℓ} . Note that ρ_{ℓ} can be efficiently computed by reading 27 values directly from $\rho_{\ell-1}$ on a $3 \times 3 \times 3$ lattice with $h_{\ell-1}$ separation, such that the density value of ρ_{ℓ} at coordinates $[x, y, z]$ can be computed as

$$\rho_{\ell}(x, y, z) = \sum_{i,j,k \in \{-h_{\ell-1}, 0, h_{\ell-1}\}} \lambda_{\ell}(i, j, k) \rho_{\ell-1}(x+i, y+j, z+k), \quad (5)$$

where the pyramidal filter function is

$$\lambda_{\ell}(i, j, k) = \frac{1}{8} \left(1 - \frac{i}{h_{\ell}}\right) \left(1 - \frac{j}{h_{\ell}}\right) \left(1 - \frac{k}{h_{\ell}}\right). \quad (6)$$

Using the pre-filtered density fields, the filtered density value at a point \mathbf{x} with filter size δ can be approximated as

$$\rho(\mathbf{x}, \delta) \approx \left(\frac{\delta}{h_{\ell-1}} - 1\right) \rho_{\ell}(\mathbf{x}) + \left(2 - \frac{\delta}{h_{\ell-1}}\right) \rho_{\ell-1}(\mathbf{x}), \quad (7)$$

which is a linear interpolation of the pre-filtered samples $\rho_{\ell}(\mathbf{x})$ for $h_{\ell-1} \leq \delta \leq h_{\ell}$.

Note that the choice of a pyramidal convolution filter is not arbitrary. The weights of the pyramidal filter coincide with the trilinear weights used while building the lighting grid hierarchy.

Practical Considerations

Explosion simulations are often computed at a high resolution for achieving high visual detail. On the other hand, self-illumination of the explosion may not require the same resolution for producing visually similar results. In practice, it is possible to completely eliminate the first few lighting levels by defining the minimum lighting resolution $\ell_{min} > 0$. In that case, the high-frequency details of the lighting from lower (finer) levels of the hierarchy are lost, but the overall lighting can still be approximated using the remaining levels. Considering the fact that the lower (finer) levels of the hierarchy include a lot more lights, merely eliminating level 0 can provide a substantial saving in computation cost. The examples we present in this paper, however, include all levels (i.e. $\ell_{min} = 0$) and they do not take advantage of this simplification.

Eliminating first few levels of the hierarchy is particularly effective for multiple scattering computation. Since multiple scattering typically produces more smooth illumination changes as compared to single scattering, the lower (finer) levels of the hierarchy are not as necessary for estimating light that goes through multiple scattering. By defining a larger ℓ_{min} for multiple scattering, the precomputation time can be reduced considerably. If a greater ℓ_{min} parameter is used for multiple scattering as compared to single scattering, we cannot simply merge the lights. In this case, we must store a separate set of lights for level ℓ_{min} of multiple scattering and use Equation 3 as its blending function (replacing r_0 with $r_{\ell_{min}}$). Higher (coarser) levels of the hierarchy, however, can be merged by generating them from both single and multiple scattered light sets. The examples in this paper do not include this simplification either.

It is also tempting to remove very dim light sources from \mathbb{S}_0 for reducing the computation cost. However, this can lead to visible flickering in the final results. This is because even if an individual light has insignificant intensity, multiple such lights can collectively represent strong enough illumination to be visible. Therefore, removing such weak lights may end up removing a substantial portion of the illumination. Nonetheless, it is possible to remove such weak lights only from the lower (finer) levels of the lighting hierarchy without introducing visible errors in the final results. This way, the cumulative illumination of multiple weak light sources can be preserved at higher (coarser) levels for relatively distant illumination, and the insignificant individual contributions of these lights can be eliminated from the lighting computation of lower (finer) levels.