

Fast Multi-image-based Photon Tracing with Grid-based Gathering

Yun Fei and Bin Wang

School of Software, Tsinghua University, P. R. China



ABSTRACT

We developed a real-time solution for approximate global illumination, where the temporal cost in complex scenes is dramatically reduced. Our approach initially traces photon-rays with multiple cube-maps, and then gathers the irradiance of photons using uniform grids filled with low-order spherical harmonics. Numerous global illumination effects can be rendered efficiently in our framework.

Problem Formation

Ray-tracing and light-gathering are the two key steps in two-pass rendering algorithms, such as photon mapping, etc. Recent method focusing on multi-image-based photon mapping [1] works well in simple scenes, but has serious limitations in complex scenes due to three reasons: a) they intersect a photon-ray iteratively with distance imposters (cube maps), where, however, the initialization can be far from accurate, which leads the iterative refinement to be failed (Figure 1); b) they intersect a photon-ray with all the distance imposters in the scene, which is costly and unnecessary; c) the method for radiance estimation called splatting, is not efficient when there are large amount of photons, since splatting millions of photons will need the rasterization of billions of pixels.

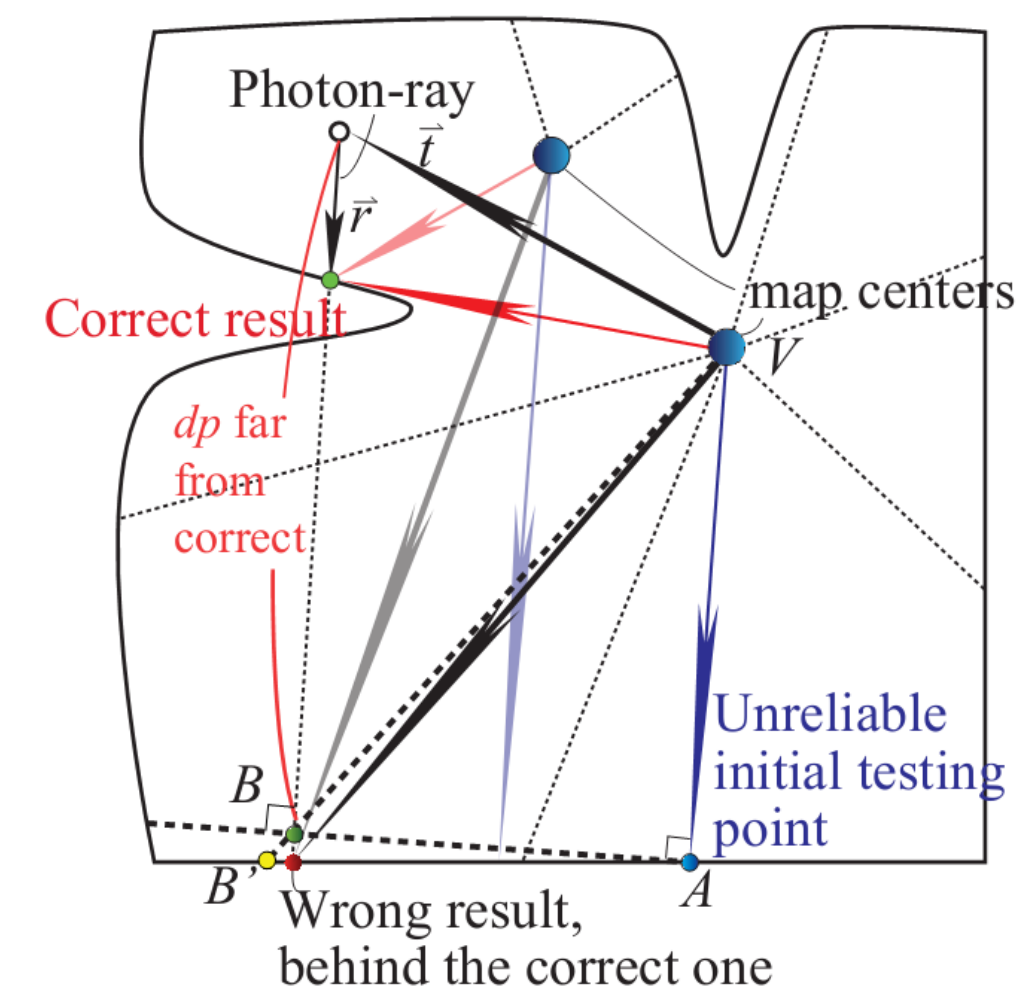


Figure 1: the photon ray \vec{r} is intersected with the distance imposter. With an unreliable initialization point A chosen, the wrong result B would be generated by existing approaches. Even with two cube maps used (two blue points), the correct result is still unapproachable.

Introduction

We propose an effective solution focusing on ray-tracing and intensity-estimating in complex scenes. Our first contribution is a heuristic ray-tracing scheme that initializes in a more reasonable way, providing a reliable beginning for the later refinement; afterwards only one or two cube-maps are used by each photon-ray, which increases the efficiency of ray-cube-map intersection over an order of magnitude. Our second contribution is an approximate grid-based light-gathering technique for all-frequency effects. Each photon is represented with a spherical harmonics (SH) vector, which is later stored into two types of uniform grids. One grid called light-gathering volume, or LGV, is a 3D volume used to gather smoothed light such as diffuse lighting, and the other one, called screen-space LGV (SSLGV), is a screen-space grid used to gather more detailed light such as caustics.

Improved Multi-image-based Ray-tracing

To solve the failure case in the existing algorithm and to improve the efficiency, we propose a new solution, based on two observations that: 1) the intersection point is only strongly related to the ambient objects near to the photon-ray, and 2) the existing methods failed since their initialization is not reliable. We detailed our method in two stages:

1) use a cascaded partition scheme to find an intersection close to be accurate, used as the initialized point (Figure 2). In this step, the cube-map closest to the photon-ray origin is chosen as the distance imposter. We then subdivide the 3D plane (the yellow




Figure 2: We divide the fan shape area (yellow) supported by the center-photon vector and the ray direction into several intervals, and find a position with the minimum energy as a "reliable initial testing point" (point I in the figure), then we refine this point to the correct result.

region) into several intervals. Then we minimize an energy across these intervals. Then we check whether the energy changing between the energy in a previous loop to the current one is sufficiently small. If it is not, we uniformly subdivide the newly computed 3D fan-shape again (the blue region, for example) and minimize the energies; otherwise, we exit the loop.

2) From the initialized point we find an exact solution following the work proposed by Szirmay-Kalos et al.[2], using the cube-map closest to the initial testing point.

The Energy Function:

To compute the energy function, we firstly fetch two distance imposters from the cube-map, in the directions of the two boundaries, denoted as \vec{p}_i and \vec{p}_{i+1} . Then we individually compute functions $f(\vec{p}_i)$ and $f(\vec{p}_{i+1})$ (defined below), and add them up as the energy function.

$f(\vec{p})$ is defined as the difference volume (Figure 3) estimated between a testing point \vec{p} and the correct intersection. It is designed based on the two key observations below:

1) the distance from the position of a photon \vec{i} to a point on the distance imposter \vec{p} should be as close as to the length of its projection on the ray direction, specifically:

$$d(\vec{p}) = \left| |\vec{p} - \vec{i}| - (\vec{p} - \vec{i}, \vec{r}) \right| \approx 0$$

2) both of the terms, $|\vec{p} - \vec{i}|$ and $(\vec{p} - \vec{i}, \vec{r})$, should be small.

To fulfill these two conditions, we compute the volume of a triangle frame around $\vec{p} - \vec{i}$ and its projection on \vec{r} .

$$f(x) = \begin{cases} \left(|\vec{p} - \vec{i}| - (\vec{p} - \vec{i}, \vec{r}) \right)^2 & \left(|\vec{p} - \vec{i}| + (\vec{p} - \vec{i}, \vec{r}) \right), (\vec{p} - \vec{i}, \vec{r}) > 0 \\ \infty & \text{otherwise} \end{cases}$$




Figure 3: The difference volume showed with translucent blue boxes. When the $|\vec{p} - \vec{i}|$ gets smaller, $f(x)$ is determined by $\left(|\vec{p} - \vec{i}| - (\vec{p} - \vec{i}, \vec{r}) \right)^2$, which is the thickness of the triangle frame.

Light Gathering Volume

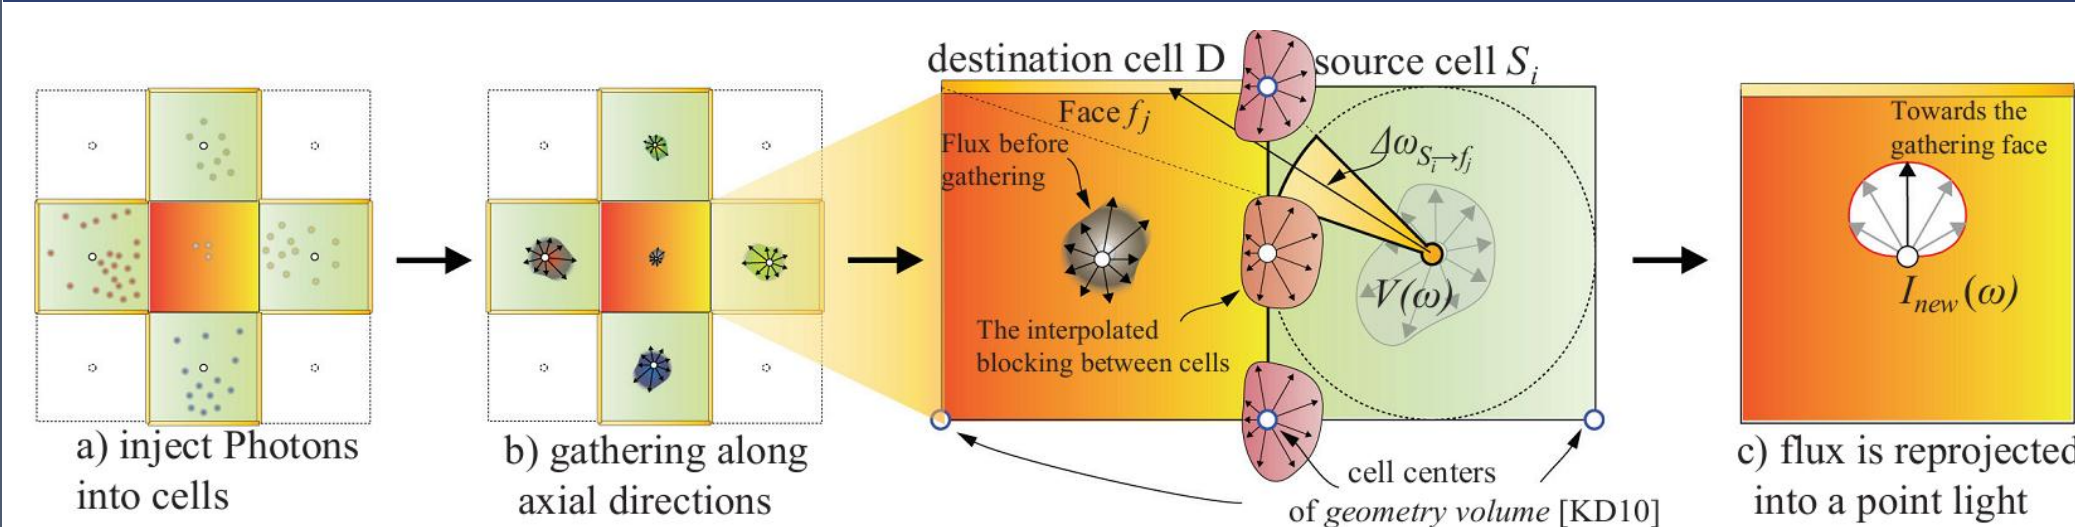


Figure 4: The photon injection, light gathering and reprojection, the figure is adapted from [4].

In this section, we introduce our novel strategy for photon gathering. Our approach can be detailed into three phases: 1) photon injection, 2) light-gathering, and 3) scene illumination (Figure 4).

The Light Gathering Volume (LGV) is a 3D texture covering the entire scene. After ray-traced, photons are injected into the LGV via two stages: 1) we transform the position that a photon-ray intersects with objects into the coordinate of texel in the 3D texture, and 2) we accumulate the SH-approximation of intensity into the texel. Each texel stores the SH-approximations of the radiance contributed by injected photons. For each photon j , the radiance L_j is calculated with the following formula (which can be derived from the Kajiya's rendering equation [3]):

$$L_j = \frac{AC_0^2}{4N} \Phi_p(\theta_j, \vec{n})^+ f_r(x)$$

A : the total area of the scene; N : the amount of photons cast; C_0 : a variant bandwidth factor [1]; θ_j : the normalized incident direction of the photon. Φ_p : the power of the photon injected; $f_r(x)$: the BRDF around point x ; \vec{n} : the normal of surface hit by the photon.

The radiance L_j is approximated by spherical harmonics (SH) vectors [4]. Multiple SH-vectors each of which corresponds to one photon are accumulated with additive blending. After injection, the radiance in each texel can represent the local illumination around the position of the texel (in world space). Then for each texel, we gather intensity from its six neighborhoods adjacent to it following Kaplanyan's light propagation framework [5].

Screen-space Light Gathering Volume

High frequency details such as caustics need higher order of SH-vectors. When using a lattice filled with SH-approximation in low-order, the resolution of the grid has to be increased. This increasing can be unacceptable for high frequency details under both spatial and temporal considerations. Fortunately, the high-resolution computation is much less costly to be taken in screen-space. Therefore, we propose a 2D grid to assist the light-gathering process (Figure 5).

Injection: Photons are transformed into eye-space and filtered, according to whether they are caustics photons or not. Then they are injected into a 2D texture in the manner similar to LGV.

Gathering: For each pixel in the eye-space, the SH-approximation is extracted from its 4-neighbored pixels and is projected to the directions to the centroid pixel. The direction is computed using world space coordinate recorded in the previously rendered position buffer.

Illumination: After several times of iteration, the SSLGV-texture is rendered with projection that estimates the intensity along the surface normal.




Figure 5: During the photon injection, the light contributed by each photon is accumulated in the axis perpendicular to the clipping plane. During gathering, SH-approximations in the yellow cells (pixels) are accumulated into the current grey cell (pixel), and the one in the red cell (pixel) is declined for its distance to the current pixel is larger than some given radius.

Limitations for Future Works

Our method shares some artifacts that are pervasive in existing lattice-based methods, specifically temporal discontinuity and light bleeding. Besides, the rendering quality of some flattened "middle"-frequency caustics is still very photon-dependent where tracing billions of photons is still unavailable in real-time. Therefore some techniques for importance sampling may be introduced into our pipeline in the future.




Figure 6: Without pre-computation, our technique renders dynamic global illumination effects in a complex scene in real-time on a GTX480.




Figure 7: Comparisons of the results. Our result is comparable to the referenced results rendered in two hours with V-Ray, as well as the result rendered in 24 minutes with PBRT.

References

- [1] Yao, C., Wang, B., Chan, B., Yong, J., and Paul, J.-C. Multi-image based photon tracing for interactive global illumination of dynamic scenes. *Comput. Graph. Forum* 29, 4 (2010), 1315–1324.
- [2] Szirmay-Kalos, L., Aszodi, B., Lazanyi, I., and Premecz, M. Approximate ray-tracing on the gpu with distance imposters. *Comput. Graph. Forum* 24, 3 (2005), 695–704.
- [3] Kajiya, J.T. The rendering equation, In Proceedings of the SIGGRAPH '86, pp. 143–150.
- [4] Sloan, P.-P. Stupid spherical harmonics (sh) tricks. In *Lectures of the GDC '08 (2008)*, pp. 1–41.
- [5] Kaplanyan, A., and Dachsbacher, C. Cascaded light propagation volumes for real-time indirect illumination. In *Proceedings of the I3D '10 (2010)*, pp. 99–107.

Contact: E-mail: fyun@acm.org, wangbins@tsinghua.edu.cn