Density Estimation on Delaunay Triangulations

Leif P. Kobbelt, Marc F. Stamminger, Hans-Peter Seidel

University of Erlangen, Computer Graphics Group
Am Weichselgarten 9. 7, 91058 Erlangen, Germany
Email: kobbelt@informatik.uni-erlangen.de

Abstract

Density Estimation is a very popular method to compute global illumination solutions in a complex virtual scene. After simulating the reflection paths of a large number of photons in the scene, the illumination at a surface point is approximated by estimating the density of photon hits in the point's surrounding. In this paper we describe a novel approach to compute such a density approximation based on Delaunay triangulation and mesh modification techniques.

1 Introduction

Global illumination computations in scenes of high complexity are one of the most challenging tasks in realistic image synthesis. The simulation of light transport through a virtual world usually requires lengthy computations and large memory. Nevertheless, the subtle indirect lighting effects which can only be captured by such methods are essential if realistic images are to be obtained.

Ray tracing is still the mostly used method in commercial rendering systems for high quality rendering [10, 3]. To some extend, indirect illumination can be computed by stochastically sending recursive rays (distribution ray tracing), but especially diffuse surfaces would require an exploding number of rays to get reasonable results without too much noise caused by the stochastic sampling scheme. Many extensions of ray tracing have been proposed to speed up such diffuse interreflections and to reduce the noise effects in the solutions.

One of these approaches is to shoot rays not only from the view point but also from the light sources into the scene [1]. Such rays are usually interpreted as photons, leaving the light source, being reflected in the scene several times and finally being absorbed at some surface. This pass is often referred to as *Particle Tracing* [2]. If the transport paths of a large set of photons are traced, the distribution of the photons' points of absorption and their incident directions give a scattered approximation of the illumination distribution in the scene. By mixing this information with the result of normal "forward" eye ray tracing, most global illumination effects can be simulated. Since the photon pass is view independent, its results can be reused to render different views of the same scene.

A very popular method that bases on these ideas is the *Density Estimation* algorithm proposed by Shirley et al. [8, 9]. This algorithm transforms the set of photon hits obtained by particle tracing into triangle meshes that approximate the irradiance in the scene. This reduces the problem to a density estimation problem, because the task is to convert a map of photon hits to a map of photon density values, which are then interpreted as irradiance.

In the density estimation method of Shirley et al., several data conversions have to be performed during the transformation, including splatting, resampling and eventually a mesh reduction and smoothing step. The data is converted between different representations until the final triangle mesh containing the irradiance approximation is obtained. In this paper we describe a density approximation approach that starts with an initial triangulation of the photon hits and then performs transformations on this single representation. The goal is to achieve a more robust, adaptive method with predictable effects of parameters that are set by the user.

2 Shirley's Algorithm

The density estimation method proposed by Shirley et al. is depicted for a one-dimensional patch in Figure 1. The algorithm takes as input the set of photon hits on a single surface (Figure 1a). The particle tracer used to create the photons guarantees that all photons have equal energy. In order to transform these photon samples into density values, the energy carried by each photon is spread on the surface by a splat centered on the photon hit (b). A crucial parameter is the size of the splat for each photon. Shirley et al. simply approximate the mean distance between all photons on the surface to determine a constant splat size allover the patch.

The splatted density function is then uniformly sampled (c) and converted into a uniform triangle mesh of irradiance values. Finally, mesh reduction techniques are applied to this mesh to smooth the obtained density distribution and also to reduce the memory requirements (d).

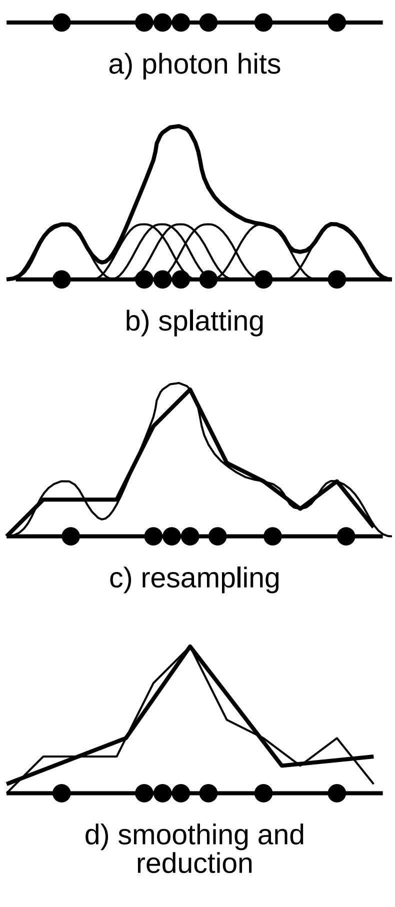

Figure 1: Shirley et al.'s Density Estimation

The method has the advantage that it is very easy to parallelize. The particle tracing pass consists of the simulation of usually millions of independent photon paths. Also the density estimation can be applied to each surface independently. Memory is mainly required to store the photon hits. Because all photons are created and processed sequentially, slow background storage like a hard disk can be used for this purpose.

Nevertheless, the method also has drawbacks. The splatting with uniform splats yields satisfactory results only for splat sizes which are about one order of magnitude larger than the photon distance. As a consequence, detail information is blurred by such big splats and homogenous regions still exhibit clearly noticable mid frequency noise in the solution. Adjusting the splat size parameter for a complete scene is crucial for the final quality. Because even for carefully selected splat sizes the results are visually unpleasant, postprocessing is necessary to improve the quality of the solutions. The postprocessing step is similar to the problem of scattered data approximation and hence similar techniques are applied.

3 Density Estimation

Splatting corresponds to the convolution of Dirac impulses located at the impact sites with a fixed smoothing kernel function. The physical model behind this way of estimating the local density of photons is that each photon carries the same amount of energy which is diffusely reflected by the patch in the vicinity of the hit. Hence the irradiance at one specific point on the patch increases when the photons come in more densely since the supports of more kernel functions overlap locally and contribute to the resulting value.

Another interpretation of the particle tracing approach is motivated by the mere definition of irradiance which is energy per area. To effectively estimate the density we can either keep a certain unit area fixed and count the number of photon being absorbed, or we can estimate the area that is illuminated by one single photon carrying unit energy. If more photons come in nearby, each has to illuminate a smaller area and hence irradiance increases.

The second model can be evaluated effectively by computing the Voronoii diagrams for the impact sites. Each photon H is assigned the region of all surface points which lie closer to H than to any other photon. This is the region which is illuminated by H. The higher the density of photons the smaller the Voronoii regions and hence the higher the local irradiance.

We can easily reconstruct the irradiance over the whole patch by constructing the Delaunay triangulation (which is the dual graph to the Voronoii diagram) and linearly interpolate the irradiance values at the vertices (= photon impact sites).

There are several drawbacks to this straight forward approach. First, the non-uniformity of the sample distribution causes severe noise artifacts in the reconstruction. Second, the number of triangles in the Delaunay triangulation is about twice the number of photons and hence much higher than actually needed for proper reconstruction. Moreover, the local resolution depends on the irradiance value only and not on significant illumination features like shadow boundaries.

Shirley solves this problem by resampling the superposition of the kernel splats on a uniform grid. This part of the algorithm is quite expensive since the kernel function has to be evaluated $O(nr^2/h^2)$ times where h is the step size of the grid, n is the number of photons, and r is the radius of the kernel function's support.

Having already constructed a Delaunay triangulation of the photon sites, we can exploit efficient algorithms for mesh optimization. First, we have to define a target resolution h, i.e. the precision up to which we want to recover illumination features in the irradiance distribution. This corresponds to the step size of the resampling grid in Shirley's approach. Our goal is to obtain an intermediate mesh with all triangles having approximately the size h.

We consider a Delaunay triangulation with absorbed photons being associated with the vertices. Each photon/vertex has unit energy. The initial triangulation usually tends to have much smaller triangles in light regions and larger triangles in dark regions. The triangulation for a simple example of a patch with a dark region on the left and a brighter region on the right can be seen in Figure 2a.

In the first step of the algorithm, we insert new vertices into the triangulation where triangles are too large. Due to energy preservation, the new vertices have to be associated with zero energy photons. A simple iterative algorithm bisects all edges in the triangulation that are longer than the prescribed h and updates the Delaunay triangulation locally [6, 5]. The resulting triangulation for our simple example is depicted in Figure 2b.

The first step yields a triangulation with all edges shorter than h but in regions of high photon density the edges are much too short. In the second step, we hence apply a custom tailored mesh decimation algorithm [7] which removes vertices iteratively and retriangulates the remaining holes. The potential candidates for removal are rated according to an application dependent fairness criterion. In our case we define the quality of a mesh to be good, if all triangles are as close as possible to equilateral. In each step we remove that vertex which improves the global fairness the most. This criterion leeds to meshes with strong coherence in the length of the edges. Empirically we could observe that the standard deviation of the edge lengths from the average stays below one percent for a wide range of resolutions. The incremental mesh reduction stops when the average edge length reaches h. Figure 2c shows the obtained triangulation with equal edge lengths.

At the begining of step two, the triangulation consists of original vertices/photons with unit energy and auxiliary vertices without energy. During the removal of vertices, we have to keep track of the total energy on the patch. When a vertex V is removed from the mesh, its energy E is distributed among the neighboring vertices V_i on its crown. To account for the relative location of the vertex V within the polygon $[V_i]$, we use inverse weighing

$$E_i \;\leftarrow\; E_i + \frac{E}{\sum_j (\varepsilon + d_i)/(\varepsilon + d_j)}$$

with $d_i = \|V_i - V\|$ and some $\varepsilon > 0$.

After the mesh decimation step, we end up with a mesh having a rather uniform distribution of vertices and strong coherence of the edge lengths around an average of h. Each vertex is assigned an amount of energy which reflects the number of vertices that have been removed in its vicinity. Hence vertices that lie in a region which initially had a high density of photons have more energy than vertices from darker regions. Dividing the vertex energies by the area of the corre-

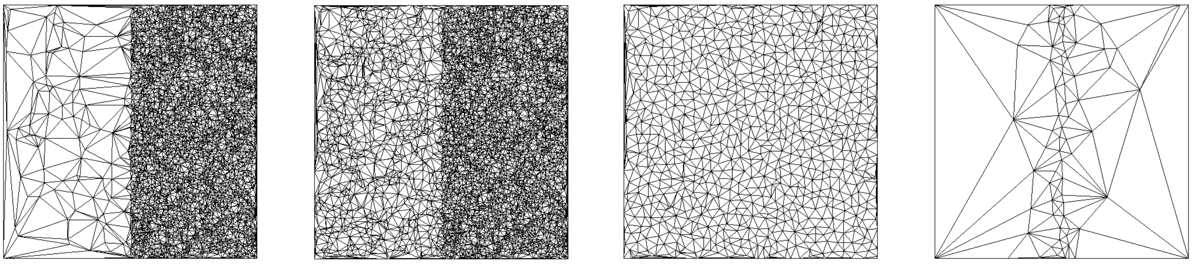

(a) Triangulation of the photon hits

(b) Insertion of zero energy photons at long edges

(c) Decimation up to target resolution

(d) Irradiance based mesh reduction

Figure 2: Triangle mesh transformations along a shadow boundary

sponding Voronoii region in the coarse triangulation yields an irradiance value for each vertex and interpolation across the Delaunay triangles gives a continuous irradiance field.

To further reduce the number of triangles in the mesh, we can assign the irradiance values as texture attributes to the vertices and, again, apply mesh decimation. This time we replace the aspect ratio fairness criterion (cf. step two) by a criterion which punishes strong changes in the color [4]. The final mesh has a minimum number of triangles while still being faithful to the original distribution. The result for our example can be seen in Figure 2d.

4 Conclusion

In this paper we showed a novel way to transform a map of photon hits obtained from a particle tracer to a photon density approximation. In opposite to the approach of Shirley et al.[8, 9], the method immediately starts with a triangulation using the photon hits and then successively transforms this mesh into an adaptive, smooth and visually satisfactory irradiance mesh. The method does *not* use splatting, and thus avoids the resulting blurring and mid frequency noise effects. Furthermore, the method is easier to use, because the parameters to steer the density estimation are more intuitive and the difficult selection of the splatting parameters is avoided.

References

[1] James Arvo. Backward ray tracing, developments in ray tracing. *ACM Siggraph '86, Course Notes*, pages 259–263, 1986.

[2] James Arvo and David Kirk. Particle transport and image synthesis. *Computer Graphics (SIGGRAPH '90 Proceedings)*, 24(4):63–66, August 1990.

[3] Andrew Glassner. *An Introduction to Ray Tracing.* Academic Press, 1989.

[4] Hugues Hoppe. Progressive meshes. *Computer Graphics (SIGGRAPH '96 Proceedings)*, pages 99–108, 1996.

[5] R. Klein, G. Liebich, and W. Straßer. Mesh reduction with error control. In *IEEE Visualization '96, Conference Proceedings*, pages 311–318, 1996.

[6] Reinhard Klein. *Netzgenerierung impliziter und parametrisierter Kurven und Flächen in einem objektorientierten System.* PhD thesis, University of Tübingen, 1995.

[7] Leif P. Kobbelt, Swen Campagna, and Hans-Peter Seidel. Mesh reduction revisited. Technical Report 6, University of Erlangen, 1997.

[8] Peter Shirley, Bretton Wade, Philip M. Hubbard, David Zareski, Bruce Walter, and Donald P. Greenberg. Global illumination via density-estimation. In *Rendering Techniques '95 (Proceedings Sixth Eurographics Workshop on Rendering)*, pages 219–230. Springer, June 1995.

[9] Bruce Walter, Philip M. Hubbard, Peter Shirley, and Donald F. Greenberg. Global illumination using local linear density estimation. *ACM Transactions on Graphics*, 16(3):217–259, July 1997.

[10] T. Whitted. An improved illumination model for shaded display. *Computer Graphics (Special SIGGRAPH '79 Issue)*, 13(3):1–14, August 1979.