Vol. 14, No. 2: 25–31

Plane Sampling for Light Paths from the Environment Map

Holger Dammertz and Johannes Hanika
Ulm University

Abstract. We present a method to start light paths on an emitting environment map in the context of Monte Carlo global illumination with image-based lighting or physical sky models. With this technique, we can efficiently render unbiased caustics from these kinds of lights, for example using bi-directional path tracing. Additionally, it is now possible to use algorithms like multiple importance sampling, photon mapping, and instant radiosity correctly with environment maps.

1. Introduction

When simulating light transport using Monte Carlo methods, all kinds of paths connecting light sources and the sensor must be generated. Even though all paths will eventually be found by a simple path tracer (starting at the sensor and recursively prolonging the path until a light source is found), some effects are best captured with some specialized sampling technique. A light tracer, for example, starting the random walk at the light sources and deterministically connecting to the sensor, is well suited to render caustics. Other global illumination methods, such as photon mapping [Jensen 96], bidirectional path tracing [Veach and Guibas 94], or point light source-based methods like instant radiosity [Keller 97] rely on paths originating from the light sources. While starting paths on geometric light sources is common practice, it is not as easy to start a path from an emitting environment map with a correct sampling density.

1086-7651/09 $0.50 per page

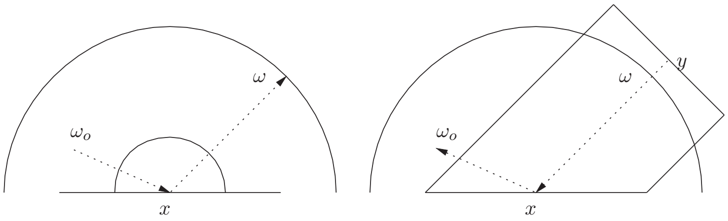

Figure 1. Two ways to sample distant illumination: on the left, ω is obtained by prolonging an eye path ending in x; on the right, a point y outside the scene is needed to trace a direction ω towards x.

2. Background

Analytic sky models or environment maps in the form of captured images [Debevec 98] approximate distant illumination by an infinitely large sphere around the scene. Thus, the radiance $L(\omega)$ emitted from the environment only depends on the direction ω for any point in the scene.

In a classical path tracing setting, the direct illumination is evaluated by sampling a direction at a point x and evaluating the visibility of the environment map (see Figure 1 (left)). In a Monte Carlo renderer this direction is picked according to a carefully chosen probability density function in order to minimize variance.

In a bi-directional algorithm it is not as straightforward, as we need to start a path from outside the scene with a correct and known probability density function. For example, in an instant radiosity-based renderer the indirect illumination is approximated by distributing point lights via a random walk from the light sources. When using environment maps, the direct illumination can be evaluated as described above. For indirect illumination, sample points to start the light paths have to be generated. This is illustrated on the right-hand side of Figure 1 where the desired sampling point is marked as y. The problem is similar to the sampling of parallel light sources with a correct sampling density that was described in [Pharr and Humphreys 04].

3. Starting Light Paths on the Environment Map

Environment maps can be interpreted as an infinite collection of parallel light sources (one for each direction). To start a particle path from the environment we divide the problem into first sampling a direction ω and then sampling an appropriate starting point y outside the scene for the directional light source associated with ω.

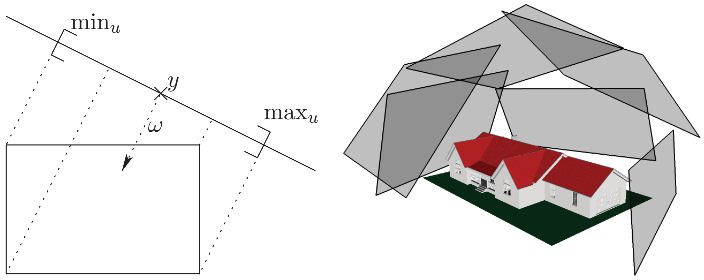

Figure 2. Illustration of the sampled plane in 2D (left), and visualization of six sampled planes for a 3D scene (right).

Algorithmically, we sample the direction ω on the environment map according to a probability density $p(\omega)$, for example $p(\omega) = 1/(4\pi)$. Then, we create an orthonormal basis (u, v, ω) around ω and project the eight corners of the bounding box of the scene onto the plane through the origin spanned by (u, v), as illustrated in Figure 2. From these projections, we only keep a two-dimensional bounding box in the coordinate system of the plane (`minu, maxu, minv, maxv`). This quad is pushed outside the bounding box of the scene. The exact distance does not matter, as the quad is used as a parallel light source. In practice, we project one vertex of the scene bounding box (or the center) onto $-\omega$ and shift this point in the direction $-\omega$ by the sum

```
omega = envmap_sample();
inv_p = 1/p(omega);     // account for p(omega)
L     = envmap_radiance(omega);
(u,v) = get_onb(omega); // create orthonormal basis (u,v,omega)
for(int c=0;c<8;c++)    // project scene bounding box
{
  dotu = dotproduct(aabb.corner[c], u);
  dotv = dotproduct(aabb.corner[c], v);
  minu = min(minu, dotu); maxu = max(maxu, dotu);
  minv = min(minv, dotv); maxv = max(maxv, dotv);
}
float dist = dotproduct(aabb.center, -omega)
          + aabb.width[0] + aabb.width[1] + aabb.width[2];
```

Figure 3. Pseudocode for the generation of the quad from a sampled environment direction. This computation can be reused for several samples.

```
y = - dist * omega + u*(minu + (maxu-minu)*rand1)
                   + v*(minv + (maxv-minv)*rand2);
inv_p *= (maxu-minu)*(maxv-minv); // compensate for p(y) = 1/A
particle.energy = L*inv_p;
particle.y = y;
particle.omega = omega;
```

Figure 4. Random sampling of the particle starting point on the precomputed quad.

of the three bounding-box widths. pseudocode for the plane generation can be found in Figure 3. Such a quad is easy to sample. The starting point y can be chosen uniformly on this quad, which results in the simple probability density $p(y) = 1/|A_y|$, where $|A_y|$ is the area of the quad. See Figure 4 for pseudocode of this operation.

The estimator remains unbiased as long as the quad is large enough to cover all of the scene. In order to maximize the number of rays actually intersecting the scene, the quad should be as small as possible.

In addition to enabling the use of algorithms relying on light paths (instant radiosity, photon maps, etc.), the advantage of the new technique becomes apparent when the environment map has a high dynamic range and importance sampling is used, i.e., $p(\omega) \sim L(\omega)$. We use sample warping [Clarberg et al. 05] to sample ω according to the luminance of the environment map.

4. Comparison

With the new technique it is now possible to efficiently create caustic paths from the environment. Other unbiased techniques generate these paths as well, but at a much lower convergence rate. For an illustration of this significant difference, see Figure 5. The environment map contains a sun with pixels 1000 times brighter as in the rest of the image. All three path samplers use the same importance sampling [Clarberg et al. 05] to evaluate direct light, thus the same sharp shadow appears in all images. However, using this kind of importance sampling for direct light only fails to capture caustics. Metropolis sampling helps, but will waste a lot of samples in the sun, as the probability is proportional to the radiance contributed to the image. Of course Metropolis sampling will perform better when a bi-directional path sampler is used and when the sun is not directly visible, especially in difficult areas of the image such as caustics from multiple inner reflections seen through a mirror.

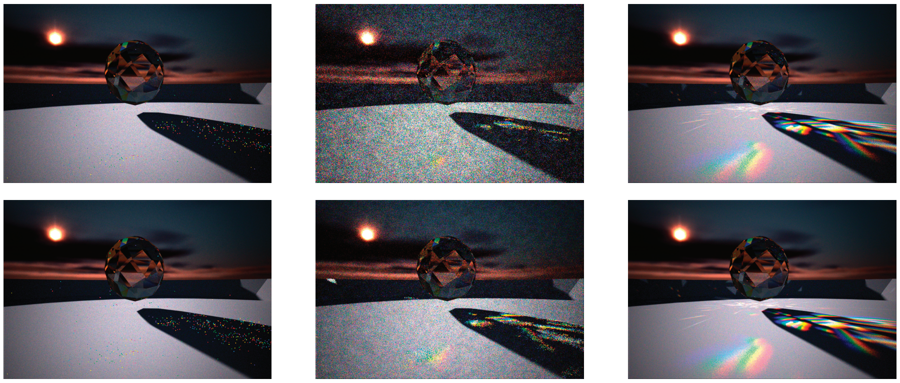

Figure 5. Comparison of plain backward path tracing from the eye (left), the same path tracer with Metropolis sampling (middle), and bi-directional path tracing with a light tracing pass using plane sampling (right). Each example on the top row has equal computation time (about five minutes on an AMD opteron). The path tracing achieves 342 samples in that time; the Metropolis version can only evaluate 105; the bi-directional method has 231 samples per pixel. The bottom row shows a comparison of equal sample count (500 samples per pixel).

5. Discussion

The plane sampling technique is a specialized solution to efficiently capture prominent effects, such as caustics, when image-based lighting is used. It is also most useful when bi-directional algorithms, such as instant radiosity or photon mapping, that rely on paths started from the light sources are used. A bi-directional path tracer could work around this by employing path tracing for these effects, but efficiency will be much lower if the environment is non-diffuse.

However, there are cases in which other techniques are better. For example, a mostly diffuse environment such as a cloudy sky is best sampled by a path tracer, since importance sampling by surface reflection can be done. If only direct lighting is of interest, combined sampling of incoming radiance and surface reflection [Clarberg and Akenine-Möller 08] is the way to go.

The simple method described above to sample a quad, which is the 2D bounding box of a projected 3D bounding box, can easily be extended to choose (u, v) to minimize the size of the quad, or be replaced by sampling, for example, a disc (the projected bounding sphere) as was done in [Pharr and Humphreys 04]. We chose quads because most render systems have the axis-aligned bounding box of the scene readily available, and it is then straightforward to extend the technique to projection maps to be used with photon mapping or instant radiosity.

Acknowledgments. The authors would like to thank the reviewers for their helpful suggestions. Additionally they want to thank mental images GmbH for funding parts of this research, and Matthias Raab and Leonhard Grünschloß for many discussions.

References

[Clarberg and Akenine-Möller 08] Petrik Clarberg and Tomas Akenine-Möller. "Practical Product Importance Sampling for Direct Illumination." *Computer Graphics Forum (Proc. of Eurographics 2008)* 27:2 (2008), 681–690.

[Clarberg et al. 05] Petrik Clarberg, Wojciech Jarosz, Tomas Akenine-Möller, and Henrik Wann Jensen. "Wavelet Importance Sampling: Efficiently Evaluating Products of Complex Functions." *ACM Transactions on Graphics* 24:3 (2005), 1166–1175.

[Debevec 98] Paul Debevec. "Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-Based Graphics with Global Illumination and High Dynamic Range Photography." In *Proceedings of SIGGRAPH 98, Computer Graphics Proceedings, Annual Conference Series*, edited by Michael Cohen, pp. 189–198. Reading, MA: Addison Wesley, 1998.

[Jensen 96] Henrik Wann Jensen. "Global Illumination Using Photon Maps." In *Rendering Techniques '96 (Proc. of the Seventh Eurographics Workshop on Rendering)*, pp. 21–30. Aire-la-Ville: Switzerland: Eurographics Association, 1996.

[Keller 97] Alexander Keller. "Instant Radiosity." *In Proceedings of SIGGRAPH 97, Computer Graphics Proceedings, Annual Conference Series*, edited by Turner Whitted, pp. 49–56. Reading, MA: Addison Wesley, 1997.

[Pharr and Humphreys 04] Matt Pharr and Greg Humphreys. *Physically Based Rendering: From Theory to Implementation.* San Francisco: Morgan Kaufmann Publishers Inc., 2004.

[Veach and Guibas 94] Eric Veach and Leonidas Guibas. "Bidirectional Estimators for Light Transport." In *Rendering Techniques '94 (Proc. of the Fifth Eurographics Workshop on Rendering)*, pp. 147–161. Aire-la-Ville, Switzerland: Eurographics Association, 1994.

Web Information:

http://jgt.akpeters.com/papers/DammertzHanika09/

High-resolution images, the Blender file used to synthesize the images in Figure 5, and a video can be found at http://www.uni-ulm.de/in/mi/graphics/plane-sampling.html.

Holger Dammertz, Ulm University, James-Franck Ring, 89081 Ulm
(holger.dammertz@uni-ulm.de)

Johannes Hanika, Ulm University, James-Franck Ring, 89081 Ulm
(johannes.hanika@uni-ulm.de)

Received November 25, 2008; accepted in revised form September 1, 2009.