DOI: 10.1111/cgf.12948
Eurographics Symposium on Rendering 2016
E. Eisemann and E. Fiume
(Guest Editors)

Volume 35 (2016), Number 4

Line Sampling for Direct Illumination

Niels Billen and Philip Dutré

KU Leuven – University of Leuven

Abstract

Computing direct illumination efficiently is still a problem of major significance in computer graphics. The evaluation involves an integral over the surface areas of the light sources in the scene. Because this integral typically features many discontinuities, introduced by the visibility term and complex material functions, Monte Carlo integration is one of the only general techniques that can be used to compute the integral. In this paper, we propose to evaluate the direct illumination using line samples instead of point samples. A direct consequence of line sampling is that the two-dimensional integral over the area of the light source is reduced to a one-dimensional integral. We exploit this dimensional reduction by relying on the property that commonly used sampling patterns, such as stratified sampling and low-discrepancy sequences, converge faster when the dimension of the integration domain is reduced. We show that, while line sampling is generally more computationally intensive than point sampling, the variance of a line sample is smaller than that of a point sample, resulting in a higher order of convergence.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Ray Tracing

1. Introduction

Evaluating direct illumination is particularly challenging, since it often shows high-frequency illumination effects such as sharp shadows and bright specular reflections. This is in contrast to indirect illumination which usually has low-frequency content that becomes smoother with additional reflections.

Direct illumination is expressed as an integral of the incident lighting, material reflectance and visibility integrated over the area of the light sources. Because this integral in general cannot be evaluated analytically, Monte Carlo integration is used as an unbiased technique to compute the integral numerically. However, for complex shadows and highly specular materials, there can be a considerable amount of noise. We can reduce the noise by improving the efficiency of the Monte Carlo estimator in two complementary ways: using clever sampling schemes which increase the order of convergence and/or reducing the time to evaluate a sample.

In this paper, we follow the former approach and express direct illumination as a one-dimensional integral over a space of line segments. As a result, the integration domain of the direct illumination is reduced from two dimensions to a single dimension which, combined with appropriate sample distributions, allows us to achieve higher orders of convergence.

The contributions of this paper are the following:

- We propose a practical algorithm which calculates the direct illumination using line samples;
- We evaluate the performance and efficiency of line sampling compared to point sampling in a variety of scenes;
- We provide theoretical insights and empirically verify that, with proper sampling distributions, line sampling can achieve a higher order of convergence than point sampling.

2. Related Work

Monte Carlo integration for ray tracing was introduced to replace the aliasing artefacts, introduced by the usage of regular sampling patterns, with stochastic noise [Coo86]. Different sampling strategies have since been developed, because generating random samples independently is generally sub-optimal as the samples tend to clump together (approximately $\sqrt{N}$ out of N samples lie in clumps [Caf98]).

Stratification tries to resolve this issue by subdividing the integration domain into disjoint strata, and generating a random sample in each stratum. Theoretically, it can be shown that stratification will never perform worse than random sampling [Caf98] and that it has considerably faster order of convergence for well-behaved functions [Mit96], [RAMN12].

Quasi-Monte Carlo integration does not rely on random samples but uses a well-chosen deterministic sequence of points samples, where the sample positions are correlated to reduce clumping [Nie92] [Caf98]. Due to the improved coverage of the integration domain, these point sequences can achieve a higher order of convergence than independent point samples.

In this paper, we primarily focus on stratification and low discrepancy sequences, for a general overview of Monte Carlo integration we refer to [KW08].

Linear light sources represent an infinite collection of point light sources which are collocated on a line segment. Typically, computing the direct illumination from a linear light source involves two parts: resolving the visibility and evaluating the shading.

The visibility can be resolved by identifying the lit, umbra and penumbra regions of the scene by projecting the scene geometry onto itself using the endpoints of the linear light source as the center of projection [NON85] [BP93]. Another approach is to use shadow maps, computed at the endpoints of the linear light source, to estimate the visibility [HBS00]. The discontinuities in the visibility along the linear light source can also be approximated numerically [OF99]. A regular grid can be employed to identify the potential occluders of a linear light source [PA91].

For an extensive survey on evaluating the shading from a linear light source using point samples, we refer to [OF01]. Analytical equations to evaluate the shading of diffuse linear light source for purely diffuse and specular materials are described in [NON85] [Pic92] [BP93].

Analytical methods exist for computing the illumination from an unoccluded area light source with diffuse emission [Lam60] [Arv95a] and arbitrary piece-wise polynomial emission [Arv95b]. In order to take visibility into account, the visible parts of the area light are computed by clipping the projection of all the occluders away from the area light source [NN85].

Applications of line sampling in computer graphics which, apart from direct illumination, make use of line samples include, analytical scanline rendering [JP00], the computation of single scattering [SZLG10] and multiple scattering [NNDJ12] in participating media and the reconstruction of high-quality motion blur [GBAM11]. A sampling technique for line segments which maintains blue-noise properties was presented by [SZG*13].

3. Monte Carlo integration of direct illumination

Direct illumination from a single area light source is expressed by the following integral:

$$L(x \to \Theta) = \int_{A_{\text{light}}} L(y \to x)\, f_r(\Theta \leftrightarrow \overrightarrow{yx})\, V(x,y)\, G(x,y)\, \mathrm{d}A_y \quad (1)$$

where L is the radiance emitted from a point y on the light source towards the shading point x, f_r is the bidirectional reflectance distribution function (BRDF), $G(x,y)$ represents the geometric form factor between the shading x and a point y on light source and $V(x,y)$ is the visibility function which is equal to 1 when x and y are mutually visible and 0 otherwise [DBBS06].

When N independent samples $y_1, y_2, y_i, ..., y_n$ are randomly chosen using a probability density function $p(y_i)$, $L(x \to \Theta)$ can be approximated by:

$$\hat{L}(x \to \Theta) = \frac{1}{N} \sum_{i=1}^{N} \frac{L(y_i \to x)\, f_r(\Theta \leftrightarrow \overrightarrow{y_i x})\, V(x,y_i)\, G(x,y_i)}{p(y_i)} \quad (2)$$

The standard deviation of this estimator converges according to $\mathcal{O}\left(N^{-\frac{1}{2}}\right)$ in terms of the number of samples. While the convergence is quite slow, requiring four times the number of samples to halve the error, it has the advantage that the convergence rate is maintained even in the presence of discontinuities and irrespective of the dimension of the integrand.

Stratification: the order of convergence be improved significantly by informed placement of the point samples over the integration domain. Stratifying the samples, by subdividing the integration domain in disjoint strata and generating a random sample in each stratum, distributes the samples more evenly in the integration domain. Stratification increases the convergence rate to $\mathcal{O}\left(N^{-\frac{1}{2}-\frac{1}{s}}\right)$ for functions with a bounded first derivative and to $\mathcal{O}\left(N^{-\frac{1}{2}-\frac{1}{2s}}\right)$ when the function is only piece-wise continuous [Mit96], s being the dimension of the integration domain.

Quasi-Monte Carlo integration: a higher order of convergence can be attained by using quasi-random (also called low-discrepancy) point samples, which are generated deterministically. The point samples in such a low-discrepancy point sequence are correlated to fill the integration domain as uniformly as possible. The convergence rate $\mathcal{O}\left((\log N)^s/N\right)$ is given by the Koksma-Hlawka inequality [Nie92], however the effectiveness of low-discrepancy sequences decreases in the presence of discontinuities [Caf98].

The order of convergence for both stratified sampling and low-discrepancy sampling increases when the dimension of the integration domain is reduced. In this paper, we aim to use this property and reduce the dimension of the integration domain by writing the direct illumination integral as a one-dimensional integral in line space. In Section 4, we show how the direct illumination can be reduced to a one-dimensional integral. In Section 5, we introduce our algorithm to evaluate the direct illumination using line samples. Finally, we show that, while line samples are more computationally intensive to evaluate, a higher quality image can be generated faster due to a higher order of convergence.

4. Light source sampling using line samples

Our goal is to reduce the variance of the Monte Carlo estimator, by reducing the dimensionality of the integration domain. We achieve this by representing the area light sources as integrals over line light sources.

We assume the light source to be triangular (but this can be generalised to any convex, planar geometry). First, we project the triangle onto an arbitrary vector in the plane of the light source $\overrightarrow{L_\perp}$, resulting in a closed interval bounded by a and b on $\overrightarrow{L_\perp}$. Each point $u \in [a,b]$ represents an infinite line in the plane of the light source along the direction $\overrightarrow{L_d}$ which is orthogonal to $\overrightarrow{L_\perp}$. Such an infinite line intersects the boundary of the light source at two points which span a line segment of length $l(u)$, with direction $\overrightarrow{L_d}$, where either intersection point can be picked as the origin $\overrightarrow{L_o}(u)$ for the line segment (Figure 1).

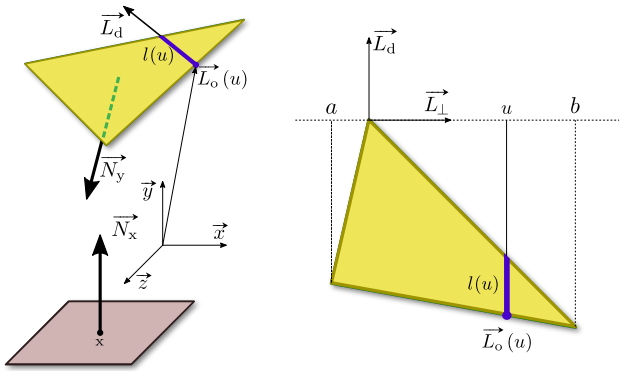

Figure 1: *Parametrisation of a line light (blue line) on a triangular area light source (yellow triangle). The left image shows the parameterisation in three dimensions, where a point x with normal N_x is illuminated by an area light source with normal N_y. The right image shows the parameterisation within the plane of the area light source.*

This allows us to explicitly write Equation (1) as:

$$L(x \to \Theta) = \int_{A_{\text{light}}} L(y \to x)\, f_r(\Theta \leftrightarrow \vec{yx})\, V(x,y)\, G(x,y)\, \mathrm{d}A_y \tag{3}$$

$$= \int_a^b \int_0^{l(u)} L(p(u) \to x)\, f_r\left(x, \Theta \leftrightarrow \overrightarrow{p(u)x}\right) V(x,p(u))\, G(x,p(u))\, \mathrm{d}u\mathrm{d}t \tag{4}$$

$$\equiv \int_a^b L_{\text{line}}(u, x \to \Theta)\, \mathrm{d}u \tag{5}$$

where $p(u) = \vec{L_o}(u) + t \cdot \vec{L_d}(u)$ is a point on the line segment for an offset u along the $\vec{L_\perp}$ axis.

By using a fixed orientation for the projection, the two-dimensional integral over the area of the light source is now expressed as a one-dimensional integral over the shading contribution from infinitely many parallel line segments. Our goal is to apply Monte Carlo integration to the one-dimensional integral, while evaluating the contributions of the line segments analytically.

5. Algorithm

In this section we will discuss our algorithm to compute the direct illumination using line samples. Our algorithm proceeds as follows:

1. Generate a line sample on the light source by sampling $u \in [a,b]$ (Section 5.1);
2. Determine the visible segments of the line sample using an appropriate acceleration structure (Section 5.2);
3. Evaluate the shading analytically for the visible parts of the line sample (Section 5.3).

5.1. Line sample generation

To evaluate the direct illumination using line samples, we first have to choose a direction for the projection of the light source $\vec{L_\perp}$. Second, we need to properly sample the offset u to generate a line segment on the light source.

5.1.1. Line orientation sampling

We can distinguish three possible methods to choose the direction to project the light source, and hence, the orientation of the line sample (see Figure 2).

Fixed direction: we keep the projection orientation $\vec{L_\perp}$ fixed for each light source (i.e. all the line samples on a light source share the same orientation for each shading point x). While this results in an unbiased estimate of the direct illumination, artefacts become apparent, which are most visible in the penumbra regions of the scene (see the renders and insets of Figure 2). When the projection direction is orthogonal to a shadow boundary, a region of the interval $[a,b]$ contains line segments with zero contribution, introducing variance. When the projection direction is parallel with the shadow boundary, every line segment in the interval $[a,b]$ has a non-zero contribution, resulting in less variance. This can be seen in the insets of Figure 2, where the right side of the penumbra (where the projection direction is parallel to the shadow boundary) is nearly converged, while the left side (where the projection direction is orthogonal to the shadow boundary) remains noisy.

Uniform direction: we can independently pick the projection orientation $\vec{L_\perp}$ uniformly on the unit circle for every line sample. However, this would result in a convergence rate $\mathcal{O}\left(N^{-\frac{1}{2}}\right)$. The expected value of the Monte Carlo estimator for Equation (5) is the same for every random projection direction. By choosing a uniform orientation projection direction for every line sample, we are just averaging the values of N independent estimators, each evaluated with only a single line sample.

While it is possible to apply appropriate sampling to the combination of the projection direction and the line offset, this would we be equivalent to sampling a two-dimensional integral, preventing us from achieving a higher order of convergence.

Parallel per shading point: we can assign a random projection direction to each light source for every shading point x. Line samples used to evaluate a single shading point are parallel to each other. This enables us to properly distribute the offsets of the parallel line samples (e.g. using stratification, low-discrepancy sequences, ...), allowing us to achieve a higher order of convergence. Furthermore, the artefacts present when the projection direction is fixed are replaced by noise, which is visually less objectionable, because the line samples of each shading point will have a different orientation. Therefore, we prefer this method to choose the projection direction in the implementation of our algorithm.

5.1.2. Line offset sampling

A simple way to chose the offsets of the line samples is to uniformly distribute them within the interval $u \in [a,b]$. However, since longer line samples will have a higher contribution to the image than shorter segments, this would introduce additional variance as shown in Figure 3.

To resolve this issue, we importance sample the line segments with respect to their length $l(u)$ on the light source. Our importance sampling reduces the noise significantly, since every point on the light source now has an equal probability to be located on a line

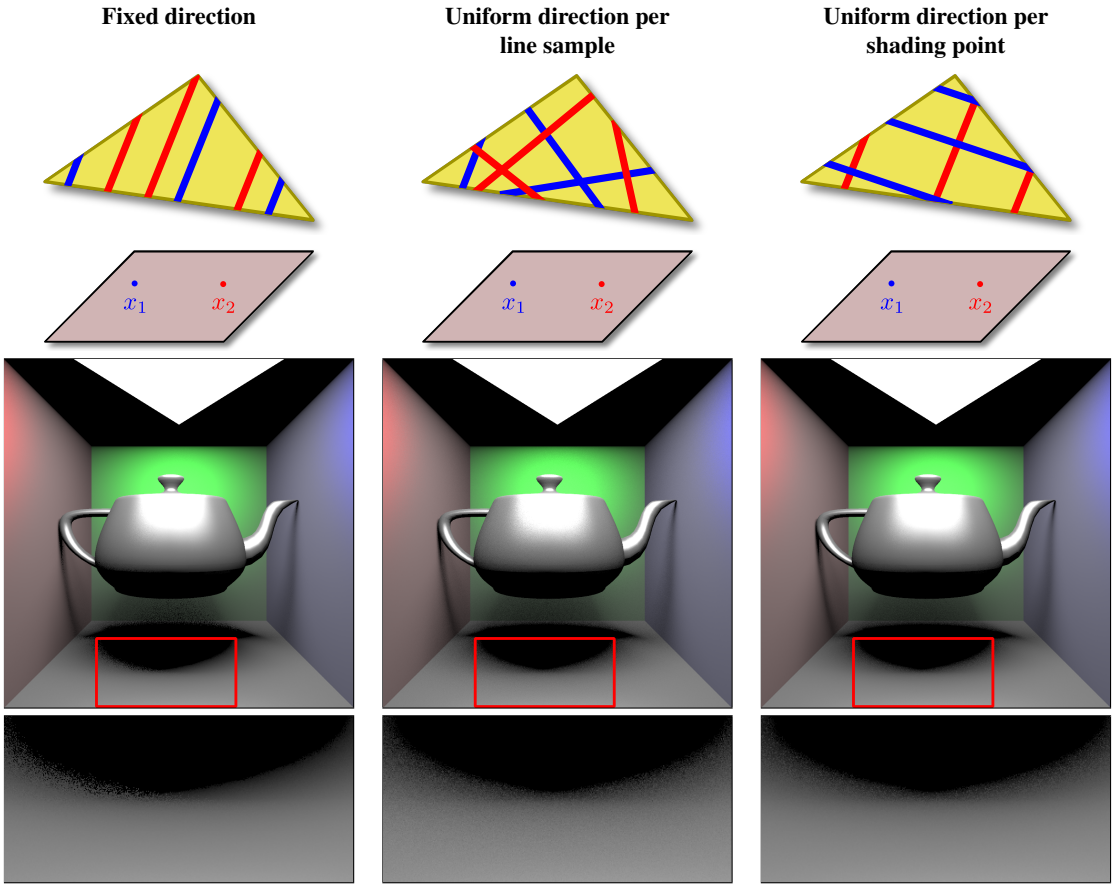

Figure 2: *Influence of the choice of the projection direction. The top row shows the directions of the line samples for two different shading points (in red and blue respectively). The scene with the glossy teapot has been rendered using the three line sampling methods.*

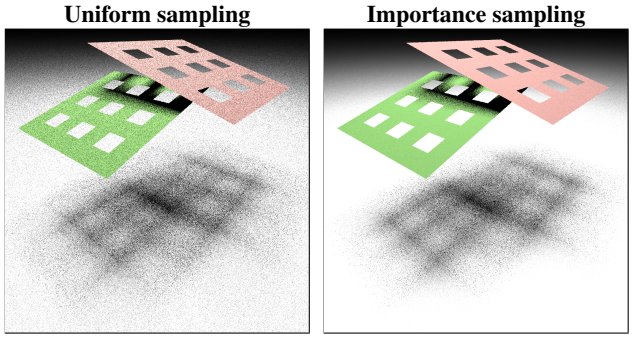

Figure 3: *Uniform versus importance sampling of the offset of the line segments.*

sample. Importance sampling can be implemented trivially for triangular light sources by noting that the probability density function for sampling u will also have a triangular shape. If the projection of the vertices of the triangle onto $\overrightarrow{L_\perp}$ are equal to a, b and c, with $a < c < b$, then the pdf will be zero at a and b and reach its maximum of $2/A_{\text{light}}$ in c.

5.2. Visibility evaluation

Before we can evaluate the shading analytically, we have to determine the parts of the line sample which are visible for the shading point. Only the geometry of the scene which intersects the triangle spanned by the shading point and the endpoints of the line sample (the *shadow triangle*) can potentially occlude the line sample (Figure 4). To efficiently find the intersecting geometry, we use a Bounding Volume Hierarchy (BVH) as an acceleration structure [WMG*07].

The BVH is traversed by recursively descending into the nodes whose bounding box intersects with the shadow triangle [AM05]. When a leaf node is encountered, we perform intersection tests between the shadow triangle and the triangles stored in the leaf node [Möl97]. The *intersection lines*, resulting from the intersection between a the shadow triangle and an occluding triangle, are then back-projected on the line sample on the light source to determine which part of the light sample is blocked.

We perform clipping efficiently in the parameter space t of a point p on the line sample with $p = \overrightarrow{L_o} + t \cdot \overrightarrow{L_d}$. We start with a fully visible line sample. Every back-projected intersection between the shadow triangle and a triangle in the scene clips a part from the

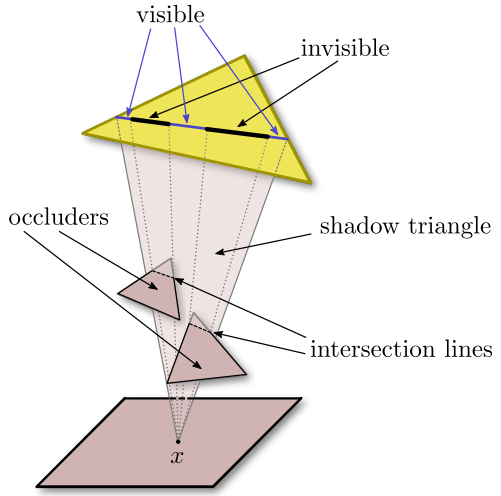

Figure 4: *Visibility determination of a line segment. Once, all the intersection lines between the shadow triangle and the geometry in the scene are found, the intersection lines are projected back onto and clipped from the line sample. After clipping, only line segments which are visible to x remain.*

Figure 5: *The number of line segments in which a line sample is cut after visibility evaluation for two scenes shown in Figure 6. The Cornell box contains many disjoint occluders, resulting in many separate line segments for which the direct illumination has to be evaluated.*

still visible intervals, which is performed in logarithmic time, because we use an binary tree over the visible intervals of the light source. Scenes containing many disjoint occluders cut a line sample in many smaller line segments. This is illustrated in Figure 5, which shows the number of disjoint visible line segments in which a line sample was cut by occluders for each pixel.

The complexity of finding all the triangles overlapping the shadow triangle equals $\mathcal{O}(\log(n)+m)$, where n is the number of geometrical primitives in the scene and m the number of triangles which overlaps with the shadow triangle. Therefore, the visibility query becomes more computationally intensive for complex scenes, where many triangles overlap with the shadow query.

5.3. Shading evaluation

To evaluate Equation (5) as a one-dimensional integral, we need an analytic expression for the direct illumination emitted by the line segments which remain visible to x after the visibility evaluation. We derive analytic solutions for diffuse line light sources illuminating diffuse (Section 5.3.1) and glossy (Section 5.3.2) materials. Similar derivations have been done in previous work on linear light sources [NON85] [Pic92] [BP93], however, they assume the light source to have a uniform emission over all outgoing directions. In contrast, we have to take normal of the light source into account, because our line samples are sampled from an area light source.

5.3.1. Diffuse material

A diffuse material has a BRDF equal to $f_r = k_d/\pi$. The only non-constant factor in the direct illumination for a completely visible diffuse light source is the geometry term. We can simplify Equation (5) by translating the frame of reference for x to be at the origin:

$$L_{\text{line}}(x \to \Theta) = L_{\text{light}} \frac{k_d}{\pi} \int_0^l G(p)\,\mathrm{d}t \tag{6}$$

$$= -L_{\text{light}} \frac{k_d}{\pi} \int_0^l \frac{\left(\vec{p}\cdot\vec{N_x}\right)\left(\vec{p}\cdot\vec{N_y}\right)}{||\vec{p}||^4}\mathrm{d}t \tag{7}$$

where $\vec{N_x}$ and $\vec{N_y}$ are the normals of the shading point and the light source respectively.

The term $\left(\vec{p}\cdot\vec{N_y}\right)$ can be reduced to $\left(\vec{L_o}\cdot\vec{N_y}\right)$ because the direction of the light $\vec{L_d}$ is always located in the plane of the light source and therefore orthogonal to the normal $\vec{N_y}$:

$$L_{\text{line}}(x \to \Theta) = -L_{\text{light}} \frac{k_d}{\pi} \left(\vec{L_o}\cdot\vec{N_y}\right) \int_0^l \frac{\left(\vec{L_o}+t\cdot\vec{L_d}\right)\cdot\vec{N_x}}{||\vec{L_o}+t\cdot\vec{L_d}||^4}\mathrm{d}t \tag{8}$$

which has a closed form solution. For the remainder of the derivation, we refer to Appendix A.

5.3.2. Phong material

A Phong material has a reflectance function:

$$f_r = k_s \frac{(n+2)\left(\vec{px}\cdot\vec{R}\right)^n}{2\pi} \tag{9}$$

where $\vec{R}$ is the perfect mirrored reflection of the incoming direction Θ and n the specular exponent. The direct illumination for a Phong material is then equal to:

$$L_{\text{line}}(x \to \Theta) = L_{\text{light}} \frac{n+2}{2\pi} k_s \int_0^l \frac{\left(\vec{px}\cdot\vec{N_x}\right)\left(\vec{px}\cdot\vec{N_y}\right)\left(\vec{px}\cdot\vec{R}\right)^n}{||\vec{px}||^{4+n}} dt \tag{10}$$

To simplify this integral, we follow the derivation of [PA91] and first perform a change in coordinate systems so that the shading point x is located at the origin and the line segment lays completely in the xy-plane. Next, we perform a change of variables and represent all vectors in spherical coordinates (θ, φ, r), to replace the integral along the length of the line segment by an integral over the angle subtended by the line segment. Because the line segment is

contained in the xy-plane, the substitution is equal to:

$$\tan(\theta) = \frac{\overrightarrow{L_{o_y}} + t \cdot \overrightarrow{L_{d_y}}}{\overrightarrow{L_{o_x}} + t \cdot \overrightarrow{L_{d_x}}} = \frac{\overrightarrow{L_{o_y}} + t \cdot \sin(\theta_L)}{\overrightarrow{L_{o_x}} + t \cdot \cos(\theta_L)} \tag{11}$$

$$r = \sqrt{\left(\overrightarrow{L_{o_x}} + t \cdot \cos(\theta_L)\right)^2 + \left(\overrightarrow{L_{o_y}} + t \cdot \sin(\theta_L)\right)^2} \tag{12}$$

where $\theta_L = \tan^{-1}\left(L_{d_y}, L_{d_x}\right)$. We find the Jacobian for the variable substitution by taking the derivative of Equation (11):

$$dt = \frac{r^2}{\overrightarrow{L_{o_x}} \sin(\theta_L) - \overrightarrow{L_{o_y}} \cos(\theta_L)} d\theta \tag{13}$$

Performing the change of variables to Equation (10) results in:

$$L_{\text{line}}(x \to \Theta) = \frac{\sin\left(\varphi_{\vec{N_x}}\right) \sin\left(\varphi_{\vec{N_y}}\right)}{\overrightarrow{L_{o_x}} \sin(\theta_L) - \overrightarrow{L_{o_y}} \cos(\theta_L)} \cdot \int_{\theta_{\min}}^{\theta_{\max}} \cos\left(\theta - \theta_{\vec{N_x}}\right) \cos\left(\theta - \theta_{\vec{N_y}}\right) \cos\left(\theta - \theta_{\vec{R}}\right)^n d\theta \tag{14}$$

By applying trigonometric identities to Equation (14), the equation can be written as a sum of two types of integrals:

$$\int \cos(\theta)^n \sin(\theta) d\theta = \frac{-\cos(\theta)^{n+1}}{n+1} \tag{15}$$

$$\int \cos(\theta)^n d\theta = \frac{\cos(\theta)^{n-1} \sin(\theta)}{n} + \frac{n-1}{n} \int \cos(\theta)^{n-2} d\theta \tag{16}$$

where the first has a definite form and the second has to be evaluated recursively. For the full derivation, we refer to Appendix B.

6. Discussion

6.1. Results

We have implemented our algorithms in PBRTv3 [PHW16] and evaluated algorithm on a number of scenes, both with large light sources (Figure 6) and with small light sources (Figure 10).

(a) **Grid planes:** this scene features a complex shadow pattern being cast by simple geometry on a surface.
(b) **Eurographics logo:** shows the specular reflection of the Eurographics logo from a surface with a large specular exponent.
(c) **Cornell box:** shows complex shadow patterns being cast by a large set of occluders in a Cornell box.
(d) **Sponza:** a scene of moderate complexity featuring bump mapping.

We compare our line sampling algorithm against point sampling using a number of Monte Carlo point sampling distributions.

Uniform sampling

Both point sampling and line sampling converge with the inverse square root of the number of samples. For both the scenes containing large and small light sources, line sampling has a consistently lower mean squared error when compared to point sampling for an equal amount of samples (see Figure 7 and Figure 11). This is to be expected, since a line sample contains infinitely many points and thus contains more information about the integrand than a single point sample.

However, if we compare the mean squared error versus the rendering time, we can see that point sampling usually achieves lower errors in a less time (see Figure 8 and Figure 12). While the variance of line samples is lower, they require more time to evaluate resulting in a lower efficiency.

Stratified sampling

As stated in previous work [Mit96], stratification of the samples results in higher convergence rates. Furthermore, stratified sampling suffers from the curse of dimensionality, implying that the convergence rate degrades for higher dimensional integrals. This can be verified from the convergence graphs in both scenes with large and small light sources (see Figure 7 and 11), which show that line sampling converges faster than point sampling.

Due to the longer rendering times of line sampling, we can see that for more complex scenes, point sampling initially achieves lower errors in less time (see Figure 8 and Figure 11). However, due to the significantly faster convergence rate, line sampling is able to overtake point sampling.

Low-discrepancy sampling

Due to the more uniform coverage of the integration domain, low-discrepancy sequences increase the order of convergence when compared to purely random distribution. This can be verified from Table 9 and Table 13, which report the order of convergence the standard deviation in terms of the number of samples for the scenes with large and small light sources respectively.

However, low-discrepancy line sampling only outperforms stratified sampling for the Eurographics logo and the Sponza. The Grid scene and Triangle soup scene feature more complex visibility events, which introduce a large amount of discontinuities in the direct illumination integral. This explains the lower convergence rate, because low-discrepancy sequences converge slower for integrands for which the variability is higher [Nie92].

Visibility evaluation

The rendering time increases when the light sources are larger. Line samples on large light sources span larger shadow triangles in the scene, increasing the time to evaluate the visibility because more intersection tests have to be performed. This can be seen by comparing the number of intersections performed against the bounding boxes of the BVH and the triangles of the scene in Table 1.

Phong shading

Analytical evaluation of the Phong shading requires evaluating a recursive function which has an algorithmic complexity of $\mathcal{O}(n)$, where n is the Phong exponent. As a result, highly glossy materials require prohibitive amounts of computation time to evaluate. To resolve this issue, we tabulate and interpolate the possible values of Equation (16). This can be done efficiently using a non-uniform point set with a higher density near the origin, since Equation (16) quickly converges to a constant for high Phong exponents.

6.2. Limitations

Currently, we are limited to materials for which the direct illumination along a line segment can be evaluated analytically. However, materials with a more complex appearance can be approximated by a weighted sum of generalised cosine lobes [LFTG97] [NDM05] and/or other basis functions which can be evaluated analytically.

Evaluating the visibility dominates the total rendering time as 90% of the time is spent traversing the acceleration structure. Furthermore, our algorithm performs worse for scenes where a large amount of geometry blocks a light source. To evaluate the visibility of a line sample, all the potential blockers have to be found. The complexity of finding all the intersections equals $\mathcal{O}(\log(n)+m)$. Therefore, in scenes with large amounts of small geometry overlapping the shadow triangles, the complexity of the visibility evaluation will become nearly linear.

7. Conclusion

We have presented a novel algorithm to evaluate the direct illumination using line instead of point samples. Line sampling reduces the dimension of the integration domain for the direct illumination from a two-dimensional integral to a one-dimensional integral. We have exploited this dimensional reduction by using the theoretical property that common sample distributions, such as stratified sampling and low-discrepancy sequences, have a higher order of convergence for lower dimensional integrals. Our results verify this behaviour and show that, while evaluating a line sample is more computationally intensive than point sampling, line sampling achieves a higher order of convergence than point sampling.

8. Future work

In this work, we have chosen to parameterise the area light source with parallel line segments. It remains future work to investigate different parameterisations of the light source. An example of such a parameterisation generates line segments which all have the same offset, but have random orientations (i.e., all line segments share a common origin).

The efficiency of the visibility estimation can be improved by using geometrical proxies [CLF*03] [SSLL14], hierarchical representations of the geometry [CNS*11], or by tagging fully occluded nodes and changing the traversal of the acceleration structure [DKH09].

We have only implemented diffuse and Phong materials, however, any material which is analytically integrable along a line segment is compatible with our framework. Spherical Gaussian are interesting basis functions for materials since they have simple analytical expressions for the inner product, convolution and integration.

9. Acknowledgements

We would like to thank our anonymous reviewers for their comments and helpful suggestions. The Sponza scene is used in this paper was originally created by Marko Dabrovic. Niels Billen is funded by the Agency for Innovation by Science and Technology in Flanders (IWT).

A. Diffuse material — Derivation

The contribution of a line sample, sampled from a diffuse light source, to the shading of a point x with a diffuse BRDF equals:

$$L_{\text{line}}(x \to \Theta) = -L_{\text{light}} \frac{k_{\text{d}}}{\pi} \left(\vec{L_{\text{o}}} \cdot \vec{N_{\text{y}}}\right) \int_0^l \frac{\left(\vec{L_{\text{o}}} + t \cdot \vec{L_{\text{d}}}\right) \cdot \vec{N_{\text{x}}}}{||\vec{L_{\text{o}}} + t \cdot \vec{L_{\text{d}}}||^4} \text{d}t \tag{17}$$

By applying the following substitutions:

$$\begin{aligned} A &= \vec{L_{\text{o}}} \cdot \vec{N_{\text{x}}} & B &= \vec{L_{\text{d}}} \cdot \vec{N_{\text{x}}} & C &= \vec{L_{\text{o}}} \cdot \vec{L_{\text{o}}} \\ D &= \vec{L_{\text{o}}} \cdot \vec{L_{\text{d}}} & E &= \vec{L_{\text{d}}} \cdot \vec{L_{\text{d}}} = 1 & F &= \vec{L_{\text{o}}} \cdot \vec{N_{\text{y}}} \end{aligned} \tag{18}$$

we can write Equation (17) to:

$$L_{\text{line}}(x \to \Theta) = -L_{\text{light}} \frac{k_{\text{d}}}{\pi} F \int_0^l \frac{A + tB}{\left(C + 2Dt + t^2\right)^2} dt \tag{19}$$

which has a closed form solution:

$$\begin{aligned} &L_{\text{line}}(x \to \Theta) = \\ &\quad L_{\text{light}} \frac{k_{\text{d}}}{\pi} \frac{(A - BD)}{2\left(C - D^2\right)^{\frac{3}{2}}} \left(\tan^{-1}\left(\frac{D}{\sqrt{C - D^2}}\right) - \tan^{-1}\left(\frac{D + l}{\sqrt{C - D^2}}\right)\right) \\ &\quad - L_{\text{light}} \frac{k_{\text{d}}}{\pi} F \frac{l\left(BC(C + l)\right) + A\left(C - Dl - 2D^2\right)}{2C\left(C - D^2\right)\left(l^2 + 2Dl + C\right)} \end{aligned} \tag{20}$$

B. Phong material — Derivation

The contribution of a line sample, sampled from a diffuse light source, to the shading of a point x with a Phong BRDF equals:

$$\begin{aligned} L_{\text{line}}(x \to \Theta) = &\frac{\sin\left(\varphi_{\vec{N_{\text{x}}}}\right) \sin\left(\varphi_{\vec{N_{\text{y}}}}\right)}{\overrightarrow{L_{\text{o}_{\text{x}}}} \sin(\theta_{\text{L}}) - \overrightarrow{L_{\text{o}_{\text{y}}}} \cos(\theta_{\text{L}})} \\ &\int_{\theta_{\min}}^{\theta_{\max}} \cos\left(\theta - \theta_{\vec{N_{\text{x}}}}\right) \cos\left(\theta - \theta_{\vec{N_{\text{y}}}}\right) \cos\left(\theta - \theta_{\vec{R}}\right)^n \text{d}\theta \end{aligned} \tag{21}$$

To solve the integral, we perform another change of variables with $u = \theta - \theta_{\vec{R}}$ and use trigonometric identities to rewrite the product of cosine to a sum:

$$\int_{u_{\min}}^{u_{\max}} \cos(u)^n \cos\left(u + \theta_{\vec{R}} - \theta_{\vec{N_{\text{x}}}}\right) \cos\left(u + \theta_{\vec{R}} - \theta_{\vec{N_{\text{y}}}}\right) \text{d}u \tag{22}$$

$$\begin{aligned} &= \int_{u_{\min}}^{u_{\max}} \cos(u)^n \\ &\quad \cdot \left(\cos(u) \cos\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{x}}}}\right) - \sin(u) \sin\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{x}}}}\right)\right) \\ &\quad \cdot \left(\cos(u) \cos\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{y}}}}\right) - \sin(u) \sin\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{y}}}}\right)\right) \text{d}u \end{aligned} \tag{23}$$

$$\begin{aligned} &= \cos\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{x}}}}\right) \cos\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{y}}}}\right) \int_{u_{\min}}^{u_{\max}} \cos(u)^{n+2} \text{d}u \\ &\quad - \cos\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{x}}}}\right) \sin\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{y}}}}\right) \int_{u_{\min}}^{u_{\max}} \cos(u)^{n+1} \sin(u) \text{d}u \\ &\quad - \sin\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{x}}}}\right) \cos\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{y}}}}\right) \int_{u_{\min}}^{u_{\max}} \cos(u)^{n+1} \sin(u) \text{d}u \\ &\quad + \sin\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{x}}}}\right) \sin\left(\theta_{\vec{R}} - \theta_{\vec{N_{\text{y}}}}\right) \int_{u_{\min}}^{u_{\max}} \cos(u)^n \sin(u)^2 \text{d}u \end{aligned} \tag{24}$$

where $u_{\min} = \theta_{\min} - \theta_{\vec{R}}$ and $u_{\max} = \theta_{\max} - \theta_{\vec{R}}$

Each of the four terms in Equation (24) can be written into a form which is equal to either Equation (15), which has a closed form solution, or either to Equation (16), which has to be evaluated recursively.

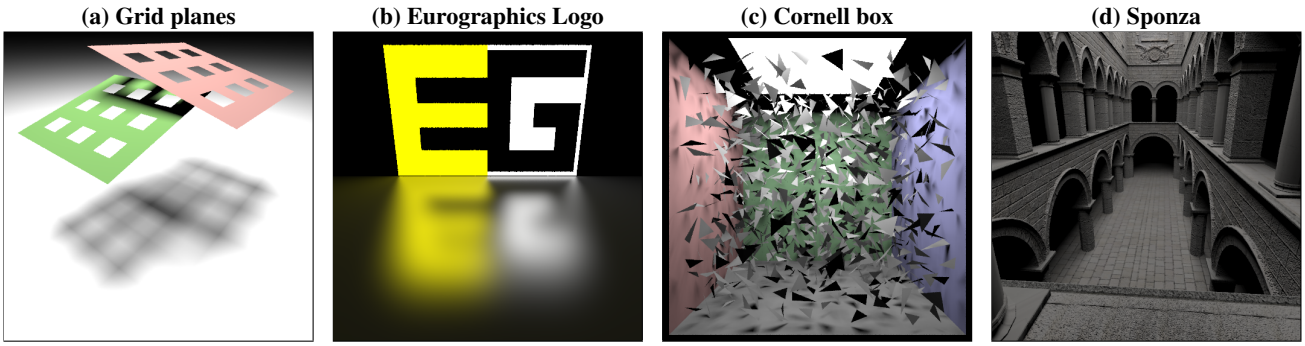

Figure 6: *The scenes used to evaluate our line sampling algorithm containing large light sources.*

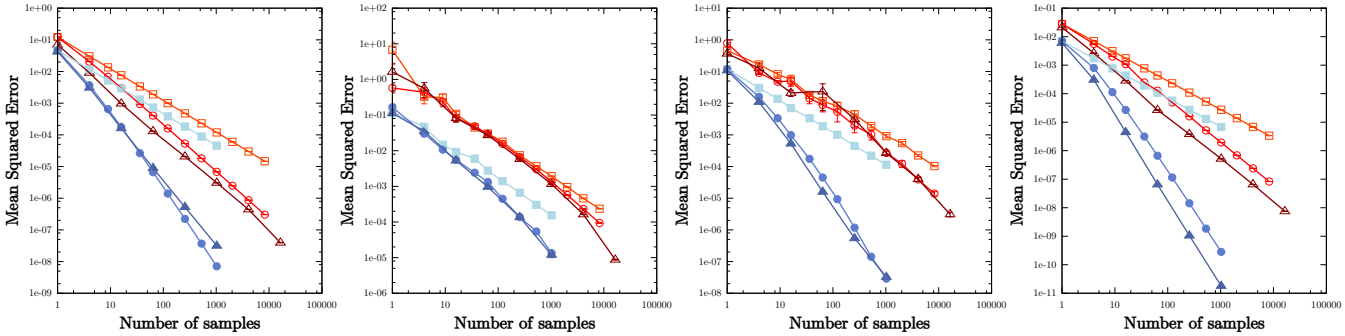

Figure 7: *The mean squared error of the different sampling distributions versus an increasing number of samples.*

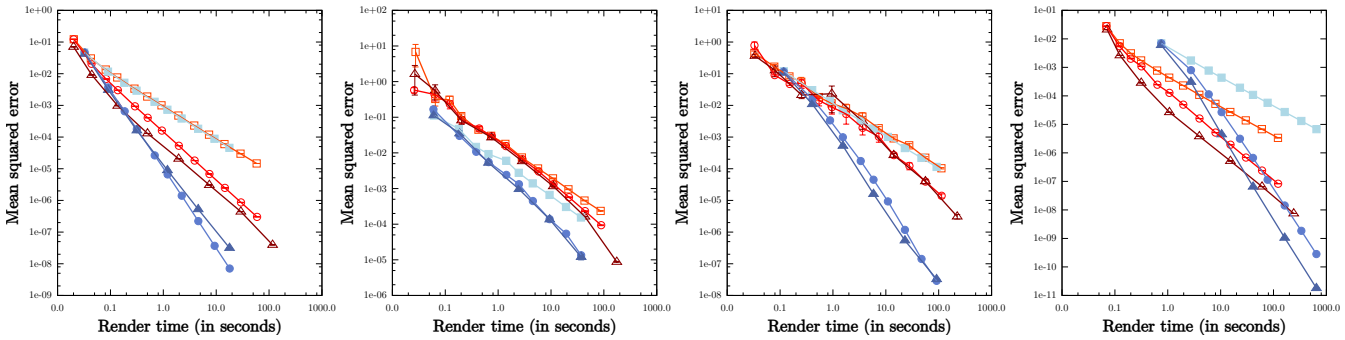

Figure 8: *The mean squared error of the different sampling distributions versus the rendering time.*

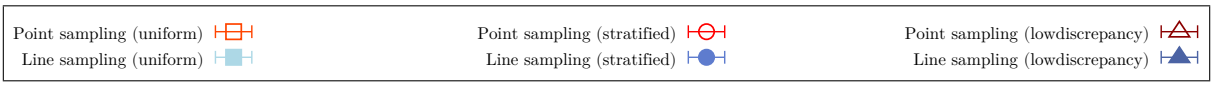

Scene	Random		Stratified		Low-discrepancy	
	Point sampling	**Line sampling**	**Point sampling**	**Line sampling**	**Point sampling**	**Line sampling**
Grid planes	-0.500006	-0.499698	-0.722607	-1.152465	-0.728064	-1.025048
Eurographics logo	-0.528350	-0.483650	-0.511009	-0.660526	-0.600950	-0.661047
Triangle soup	-0.466542	-0.502162	-0.580416	-1.130108	-0.589068	-1.115833
Sponza	-0.500072	-0.499975	-0.720497	-1.269119	-0.763655	-1.444029

Figure 9: *Convergence order of the standard deviation the in terms of the number of samples for the various sampling distributions.*

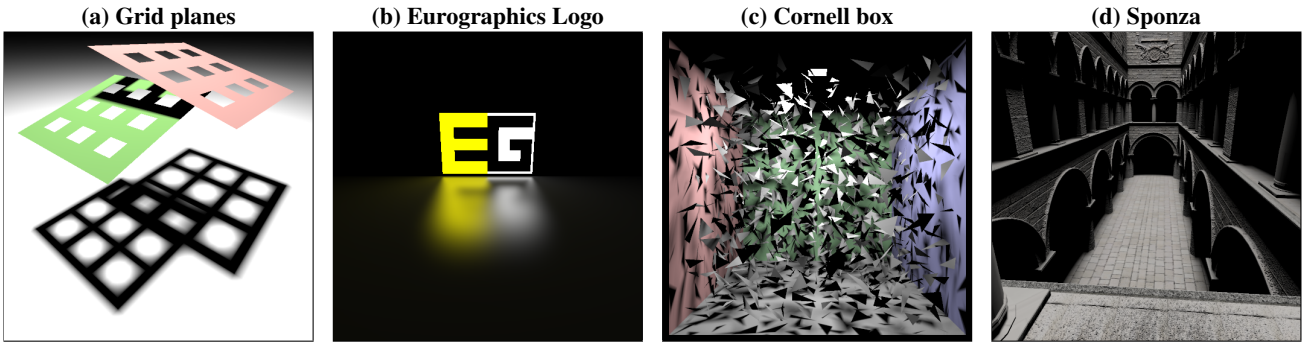

Figure 10: *The scenes used to evaluate our line sampling algorithm containing small light sources.*

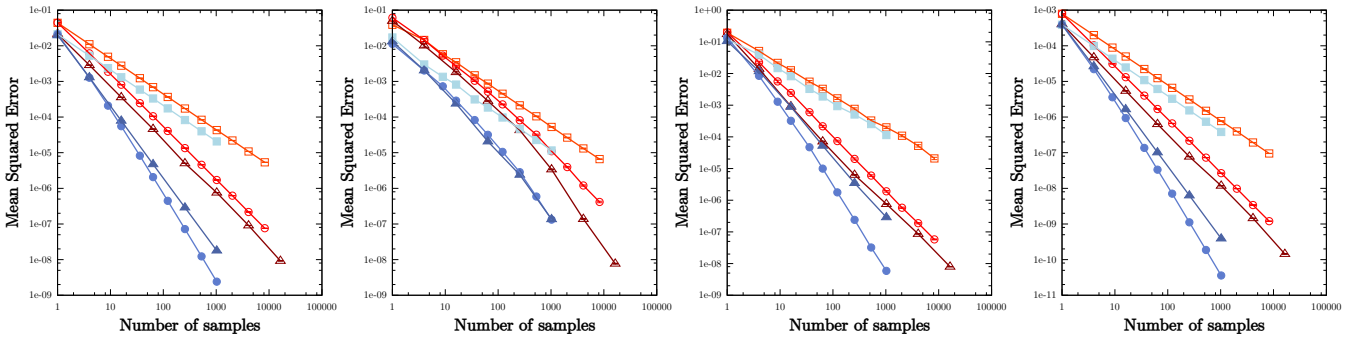

Figure 11: *The mean squared error of the different sampling distributions versus an increasing number of samples.*

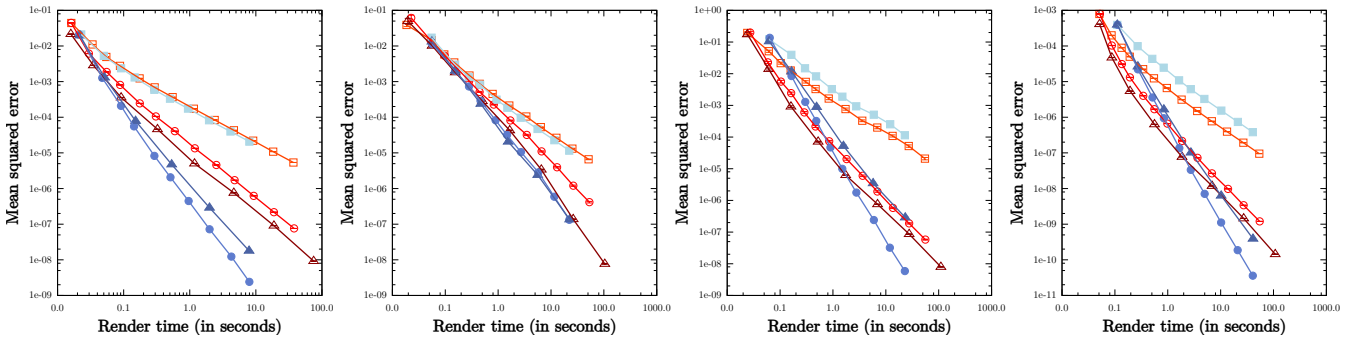

Figure 12: *The mean squared error of the different sampling distributions versus the rendering time.*

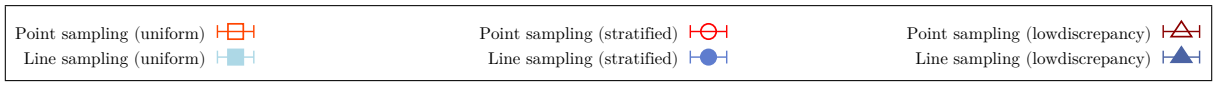

Scene	Random		Stratified		Low-discrepancy	
	Point sampling	**Line sampling**	**Point sampling**	**Line sampling**	**Point sampling**	**Line sampling**
Grid planes	-0.499750	-0.500003	-0.737551	-1.165623	-0.752114	-1.005696
Eurographics logo	-0.493513	-0.514205	-0.660507	-0.820198	-0.803129	-0.822268
Triangle soup	-0.502810	-0.506269	-0.845316	-1.245989	-0.866152	-0.945898
Sponza	-0.500227	-0.500084	-0.743929	-1.180712	-0.757675	-0.999242

Figure 13: *Convergence order of the standard deviation the in terms of the number of samples for the various sampling distributions.*

Scene	Grid Planes		Eurographics logo		Triangle soup		Sponza	
Light source size	**Small**	**Large**	**Small**	**Large**	**Small**	**Large**	**Small**	**Large**
Render time	7.9 s	11.0 s	22.0 s	22.7 s	22.7 s	58.0 s	39.7 s	455.7 s
# shadow triangle–box intersections	1462.0 M	1926.4 M	1929.3 M	1929.3 M	7686.7 M	13848.6 M	7809.1 M	53729.2 M
# shadow triangle–triangle intersections	136.3 M	268.2 M	412.7 M	415.5 M	344.9 M	1238.2 M	1498.3 M	21474.2 M
# shadow triangle–box intersections per second	185.8 M	174.8 M	87.7 M	85.1 M	339.4 M	238.9 M	196.6 M	117.9 M
# shadow triangle–triangle intersections per second	17.3 M	24.3 M	18.8 M	18.3 M	15.2 M	21.4 M	37.7 M	47.1 M

Table 1: *Performance of the visibility evaluation for our test scenes rendered with small and large light sources. We used 1024 line samples to evaluate the direct illumination for each shading point. The table shows the render time, the number of intersections between shadow triangles and the bounding boxes of the BVH nodes, the number of intersections between shadow triangles and the geometry of the scene. Finally, we have also included the average number of intersection tests per second.*

References

[AM05] AKENINE-MÖLLER T.: Fast 3d triangle-box overlap testing. In *ACM SIGGRAPH 2005 Courses* (New York, USA, 2005), SIGGRAPH '05, Association for Computing Machinery. doi:10.1145/1198555.1198747. 4

[Arv95a] ARVO J.: *Analytic methods for simulated light transport*. PhD thesis, Yale University, 1995. 2

[Arv95b] ARVO J.: Applications of irradiance tensors to the simulation of non-lambertian phenomena. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques* (New York, USA, August 1995), SIGGRAPH '95, Association for Computing Machinery, pp. 335–342. doi:10.1145/218380.218467. 2

[BP93] BAO H., PENG Q.: Shading models for linear and area light sources. *Computers & Graphics 17*, 2 (1993), 137–145. doi:10.1016/0097-8493(93)90097-S. 2, 5

[Caf98] CAFLISCH R. E.: Monte carlo and quasi-monte carlo methods. *Acta Numerica 7* (January 1998), 1–49. doi:10.1017/S0962492900002804. 1, 2

[CLF*03] CHRISTENSEN P. H., LAUR D. M., FONG J., WOOTEN W. L., BATALI D.: Ray differentials and multiresolution geometry caching for distribution ray tracing in complex scenes. *Computer Graphics Forum 22*, 3 (2003), 543–552. doi:10.1111/1467-8659.t01-1-00702. 7

[CNS*11] CRASSIN C., NEYRET F., SAINZ M., GREEN S., EISEMANN E.: Interactive indirect illumination using voxel cone tracing. *Computer Graphics Forum 30*, 7 (2011), 1921–1930. doi:10.1111/j.1467-8659.2011.02063.x. 7

[Coo86] COOK R. L.: Stochastic sampling in computer graphics. *ACM Transactions on Graphics (TOG) 5*, 1 (January 1986), 51–72. doi:10.1145/7529.8927. 1

[DBBS06] DUTRÉ P., BALA K., BEKAERT P., SHIRLEY P.: *Advanced Global Illumination*. AK Peters Ltd, 2006. doi:10.1201/b10632. 2

[DKH09] DJEU P., KEELY S., HUNT W.: Accelerating shadow rays using volumetric occluders and modified kd-tree traversal. In *Proceedings of the Conference on High Performance Graphics 2009* (New York, NY, USA, 2009), HPG '09, Association for Computing Machinery, pp. 69–76. URL: http://doi.acm.org/10.1145/1572769.1572781, doi:10.1145/1572769.1572781. 7

[GBAM11] GRIBEL C. J., BARRINGER R., AKENINE-MÖLLER T.: High-quality spatio-temporal rendering using semi-analytical visibility. *ACM Transactions on Graphics (TOG) 30*, 4 (July 2011), 54:1–54:12. doi:10.1145/2010324.1964949. 2

[HBS00] HEIDRICH W., BRABEC S., SEIDEL H.-P.: Soft shadow maps for linear lights. In *Rendering Techniques 2000: Proceedings of the Eurographics Workshop, Czech Republic, June 26–28, 2000* (Vienna, June 2000), Péroche B., Rushmeier H., (Eds.), Springer Vienna, pp. 269–280. doi:10.1007/978-3-7091-6303-0_24. 2

[JP00] JONES T. R., PERRY R. N.: Antialiasing with line samples. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000* (London, UK, 2000), Péroche B., Rushmeier H., (Eds.), Springer Vienna, pp. 197–206. doi:10.1007/978-3-7091-6303-0_18. 2

[KW08] KALOS M. H., WHITLOCK P. A.: *Monte Carlo methods, Second Revised and Enlarged Edition*, 2 ed. Wiley-VCH, 2008. doi:10.1002/9783527626212. 2

[Lam60] LAMBERT J. H.: *Photometria sive De mensura et gradibus luminis colorum et umbrae*. Leipzig verlag von Wilhelm Engelmann, Augsburg, 1760. 2

[LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques* (New York, USA, 1997), SIGGRAPH '97, ACM Press/Addison-Wesley Publishing Co., pp. 117–126. doi:10.1145/258734.258801. 7

[Mit96] MITCHELL D. P.: Consequences of stratified sampling in graphics. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques* (1996), SIGGRAPH '96, Association for Computing Machinery, pp. 277–280. doi:10.1145/237170.237265. 1, 2, 6

[Möl97] MÖLLER T.: A fast triangle-triangle intersection test. *Journal of Graphics Tools 2* (1997), 25–30. doi:10.1080/10867651.1997.10487472. 4

[NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of brdf models. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques* (Aire-la-Ville, Switzerland, Switzerland, 2005), Eurographics Symposium on Rendering 2005, Eurographics Association, pp. 117–126. doi:10.2312/EGWR/EGSR05/117-126. 7

[Nie92] NIEDERREITER H.: *Random Number Generation and quasi-Monte Carlo Methods*. Society for Industrial and Applied Mathe-

matics, Philadelphia, Pennsylvania, USA, 1992. doi:10.1137/1.9781611970081. 1, 2, 6

[NN85] NISHITA T., NAKAMAE E.: Continuous tone representation of three-dimensional objects taking account of shadows and interreflection. *SIGGRAPH '85 Proceedings of the 12th annual conference on Computer graphics and interactive techniques 19*, 3 (July 1985), 23–30. doi:10.1145/325165.325169. 2

[NNDJ12] NOVÁK J., NOWROUZEZAHRAI D., DACHSBACHER C., JAROSZ W.: Virtual ray lights for rendering scenes with participating media. *ACM Transactions on Graphics 31*, 4 (jul 2012), 60:1–60:11. doi:10.1145/2185520.2185556. 2

[NON85] NISHITA T., OKAMURA I., NAKAMAE E.: Shading models for point and linear sources. *ACM Transactions on Graphics 4*, 2 (April 1985), 124–146. doi:10.1145/282918.282938. 2, 5

[OF99] OUELLETTE M. J., FIUME E.: Approximating the location of integrand discontinuities for penumbral illumination with area light sources. In *Rendering Techniques '99: Proceedings of the Eurographics Workshop in Granada, Spain, June 21–23, 1999* (Vienna, June 1999), Lischinski D., Larson G. W., (Eds.), Springer Vienna, pp. 213–224. doi:10.1007/978-3-7091-6809-7_19. 2

[OF01] OUELLETTE M. J., FIUME E.: On numerical solutions to one-dimensional integration problems with applications to linear light sources. *ACM Transactions on Graphics (TOG) 20*, 4 (October 2001), 232–279. doi:10.1145/502783.502785. 2

[PA91] POULIN P., AMANATIDES J.: Shading and shadowing with linear light sources. *Computers & Graphics 15*, 2 (1991), 259–265. doi:10.1016/0097-8493(91)90079-W. 2, 5

[PHW16] PHARR M., HUMPHREYS G., WENZEL J.: Pbrt v3, 2016. http://www.pbrt.org/. 6

[Pic92] PICOTT K. P.: Extensions of the linear and area lighting models. *IEEE Computer Graphics and Applications 12*, 2 (March 1992), 31–38. doi:10.1109/38.124286. 2, 5

[RAMN12] RAMAMOORTHI R., ANDERSON J., MEYER M., NOWROUZEZAHRAI D.: A theory of monte carlo visibility sampling. *ACM Transactions on Graphics (TOG) 31*, 5 (September 2012), 121:1–121:16. doi:10.1145/2231816.2231819. 1

[SSLL14] SILVENNOINEN A., SARANSAARI H., LAINE S., LEHTINEN J.: Occluder simplification using planar sections. *Computer Graphics Forum 33*, 1 (2014), 235–245. doi:10.1111/cgf.12271. 7

[SZG*13] SUN X., ZHOU K., GUO J., XIE G., PAN J., WANG W., GUO B.: Line segment sampling with blue-noise properties. *ACM Transactions on Graphics (TOG) 32*, 4 (July 2013), 127:1–127:14. doi:10.1145/2461912.2462023. 2

[SZLG10] SUN X., ZHOU K., LIN S., GUO B.: Line space gathering for single scattering in large scenes. *ACM Transactions on Graphics (TOG) 29*, 4 (July 2010), 54:1–54:8. doi:10.1145/1778765.1778791. 2

[WMG*07] WALD I., MARK W. R., GÜNTHER J., BOULOS S., IZE T., HUNT W., PARKER S. G., SHIRLEY P.: State of the art in ray tracing animated scenes. In *Eurographics 2007 - State of the Art Reports* (2007), Schmalstieg D., Bittner J., (Eds.), The Eurographics Association. doi:10.2312/egst.20071056. 4