DOI: 10.1111/cgf.14583
Eurographics Symposium on Rendering 2022
A. Ghosh and L.-Y. Wei
(Guest Editors)

Volume 41 (2022), Number 4

Once-more scattered next event estimation for volume rendering

Johannes Hanika[1,2], Andrea Weidlich[1,3], Marc Droske[1,4]

[1]Weta Digital [2]Karlsruhe Institute of Technology [3]NVIDIA [4]Unity

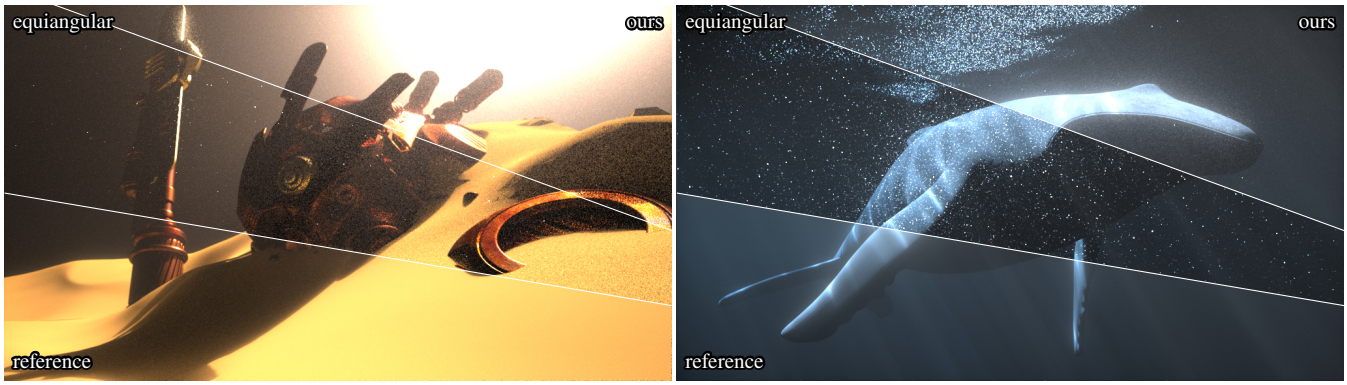

Figure 1: *Two cases where highly forward scattering phase functions deliver a distinct look. Left: atmospheric haze, right: ocean water. The slices compare our specialised sampling technique to equiangular sampling and a reference image. The characteristic glow around the objects submerged in the medium disappears in noise for equiangular sampling, since it is caused by the phase function that is not importance sampled in this case. The images other than the reference are equal sample count and approximately equal time.*

Abstract

We present a Monte Carlo path tracing technique to sample extended *next event estimation contributions in participating media: we consider one additional scattering vertex on the way to the next event, accounting for focused blur, resulting in visually interesting image features. Our technique is tailored to thin homogeneous media with strongly forward scattering phase functions, such as water or atmospheric haze. Previous methods put emphasis on sampling transmittances or geometric factors, and are either limited to isotropic scattering, or used tabulation or polynomial approximation to account for some specific phase functions. We will show how to jointly importance sample the product of an arbitrary phase function with analytic sampling in the solid angle domain and the two reciprocal squared distance terms of the adjacent edges of the transport path. The technique is fast and simple to implement in an existing rendering system. Our estimator is designed specifically for forward scattering, so the new technique has to be combined with other estimators to cover the backward scattering contributions.*

CCS Concepts

• ***Computing methodologies*** → ***Rendering;***

1. Introduction

When rendering computer generated imagery, it is our goal to immerse the viewer in a virtual environment. Due to the sophistication of the human visual system, often times such immersion depends on subtle features in the image that are subconsciously recognised as familiar light transport effects. Volumetric scattering plays a critical role in conveying scale and atmosphere. For example, Figure 1 (right) shows an underwater object, lit by sun rays travelling through the shallow water. The whale appears slightly blurred with distance, and its bright back features some glow around it. These effects are due to scattering in the water between the object and the observer and have a characteristic shape that depends on the phase function of the medium. The effect is not limited to underwater though, it can be seen in hazy atmospheres where it will change the appearance of landscapes and causes glow around the reflections off specular objects (see Figure 1, left).

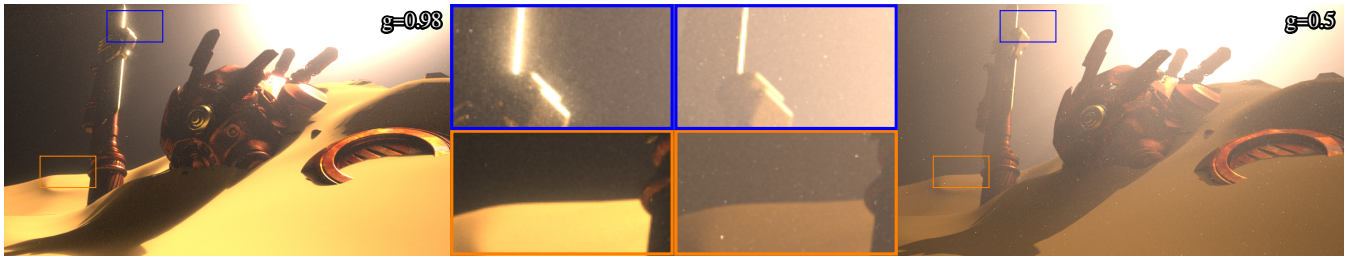

Figure 2: *A visual comparison of the same scene with mean cosine $g = 0.98$ and $g = 0.5$. To avoid challenging light transport paths, a common approach is to widen the forward peak of the phase function, resulting in an unwanted change in contrast and loss of the characteristic glow around highlights. Our new method makes the left simulation practical and efficient.*

Scattering in such media often has a highly-peaked forward component with only small backscattering part. While the backward scattering is comparably uniform and easy to sample, the forward scattering becomes close to Dirac, which makes the scattering especially difficult to sample with existing techniques. Unwanted fireflies are the result, which will cause extremely long render times. The visual importance of peaked phase functions can be seen in Figure 2 where we show a comparison of low order scattering (max. 4 path vertices) with a forward scattering phase function $g = 0.98$ and with only moderately forward scattering $g = 0.5$. As illustrated in Figures 3 and 4, the fireflies are caused by a lack of sampling density around contributions of high throughput: importance sampling the phase function addresses only one peak in the distribution, while a stronger concentration occurs around directions towards the end vertex. As it can be seen, the common practice to reduce the forward scattering to allow for better convergence results in a different appearance and important features are lost.

In this paper we design an estimator specifically for the forward scattering part. We present a simple yet powerful specialised Monte Carlo sampling technique for thin, homogeneous, forward scattering media. The method is a form of *extended next event estimation*, adding a vertex on the emitter or sensor, as well as one additional scattering vertex in between. With this, we analytically sample two geometry terms and any phase function with existing solid angle sampling routine. Our analysis and sampling is based on homogeneous media, but we briefly discuss heterogeneous media too.

2. Background and previous work

Light transport in a volume can be described as path space integral

$$\int_{\mathcal{P}} f(\overline{\boldsymbol{x}})\,\mathrm{d}\overline{\boldsymbol{x}} \tag{1}$$

where $\overline{\boldsymbol{x}} = (\boldsymbol{x}_0, \boldsymbol{x}_1, \cdots \boldsymbol{x}_k) \in \mathcal{P}$ are light transport paths consisting of a list of vertices, connecting the sensors and the light sources. The integrand is the *measurement contribution function* and is a product of multiple relevant terms that we will detail later.

An effective way to solve this equation is with Monte Carlo algorithms. However, their stochastic nature can result in high variance and different approaches exist to decrease noise. Due to the amount of literature we would like to point the reader to Novák et

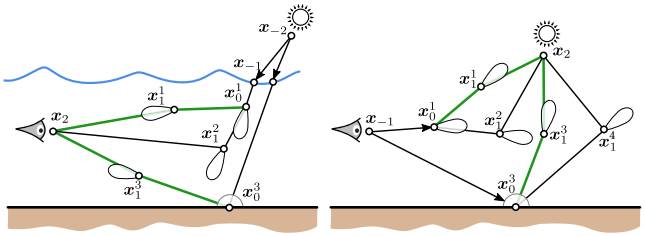

Figure 3: ***Left:*** *Classic light tracing with next event estimation or equiangular sampling samples a direction at a volume vertex $\boldsymbol{x}_0^1$ or surface vertex $\boldsymbol{x}_0^3$ towards $\boldsymbol{x}_1^1$ and $\boldsymbol{x}_1^3$, respectively, typically by importance sampling the BSDF. The high throughput 2-segment connections (green) that bend at these vertices toward the eye and create an intermediate vertex $\boldsymbol{x}_1^1$ or $\boldsymbol{x}_1^3$ are sampled with low probability and thus suffer from variance. Our technique is tailored to sample these connections efficiently.* ***Right:*** *The path tracing case analogously samples such 2-segment connections towards the light.*

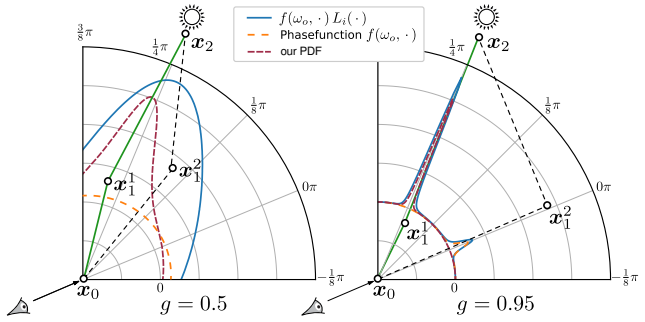

Figure 4: *In the volumetric double scattering case, the outgoing radiance $L_o(\boldsymbol{x}_0)$ is a result of scattering at $\boldsymbol{x}_0$ and an intermediate vertex $\boldsymbol{x}_1$. The main product terms, the one-bounce incoming radiance $L_i(\boldsymbol{x}_0)$ and the phase function f_0, are strongly competing in media with anisotropic phase functions. First sampling a direction according to f_0 and subsequently sampling $\boldsymbol{x}_1$ handles the forward scattering of the phase function well but misses the distinct peak centred around the direction towards $\boldsymbol{x}_2$. Our method is tailored to sample paths (green) corresponding to this peak: the dashed bordeaux-coloured line corresponds to the marginalised solid-angle PDF at $\boldsymbol{x}_0$ of our sampling method.*

Table 1: *Overview of the terms that individual techniques sample. The checkmarks in parentheses indicate that these cannot be done analytically but rely on approximation in the form of tabulation or polynomial fits. In the case of thin homogeneous forward scattering media, the phase function $f_s(\boldsymbol{x}_1)$ and the reciprocal squared distances d_1^{-2} and d_2^{-2} are of highest importance. Previous methods could not sample this combination analytically.*

	f_s	d_1^2	d_2^2	T_1	T_2	$f_0 \cos$
OMNEE (ours)	✓	✓	✓			
std. NEE		✓		✓		✓
[KF12]			✓			✓
[GKH*13]	(✓)	✓	✓			(✓)
[VGGN21]	(✓)		✓		(✓)	

al. [NGHJ18] for a better introduction and thorough review of all aspects that we are not discussing.

(Semi)-analytic solutions. Closed-form solutions are rare and normally not general enough to be useful for computer graphics. For our specific case, there exist approaches [Wil77; dEo16] but do not generalise to arbitrary phase functions and are valid for point light sources only. Furthermore, they do not allow for Monte Carlo sampling of other sources of variation such as occlusion and general emitter distribution functions.

A closely related direction of research is concerned with the air light integral, collecting all the flux along a fixed viewing direction. This has been explored for real time applications in homogeneous volumes with isotropic phase function [SRNN05] as well as for the more general case with arbitrary phase function and analytic integration based on series expansion of the phase function and the light source [PSP09; PSP10; PSS11].

Product sampling. Product sampling will importance sample the product of several terms. Georgiev et al. [GKH*13] introduced a method to importance sample volumetric paths of two bounces. The sampling of the geometric terms is related to our approach, but the phase function has to be tabulated and the method is best suited for isotropic scattering. Villeneuve et al. [VGGN21] sample the product of transmittance, phase function and point normal. Phase function and transmittance have to be approximated with a polynomial equation, which becomes worse the more forward scattering the phase function becomes. In contrast, we can use any phase function sampling while still keeping the sampling of both geometry terms analytic. Table 1 compares which terms are sampled by different techniques.

Volume path guiding. Volume path guiding uses a data-driven approach that comes with a significant implementation overhead and their memory consumption often limit their usefulness when rendering big scenes. Reibold et al. [SHJD18] proposed to use 3D truncated Gaussians for volumes and a mix of path guiding for paths with high variance and otherwise regular path tracing. Herholz et al. [HZE*19] use von Mises-Fisher mixtures to sample scattering directions according to the product of phase function and incident radiance. Furthermore, the distance sampling accounts for learned distributions of inscattering. Another approach [DWWH20] is to use a spatio-directional tree to represent the incoming radiance distribution, though this approach does not include distance sampling. Anisotropic volumetric scattering over more than one scattering event is very challenging to learn in a guiding framework. The spiky nature of the multidimensional integrand makes it difficult to represent well in the guiding data-structure and the efficiency of the training suffers from very noisy input samples. In contrast to these memory-intensive methods, our approach does not account for occlusion but is very lightweight and easy to implement.

Next event estimation (NEE). NEE simulates a path up to a current end vertex $\boldsymbol{x}_0$ that is then in turn deterministically connected to a point $\boldsymbol{x}_2$ sampled either on an emitter (path tracing) or the camera aperture (light tracing). Distance sampling [LW96] samples a distance and connects to the light source or camera. Kulla and Fajardo [KF12] introduced equiangular sampling that samples the inverse squared distance fall-off of a given point. NEE can be further improved by selecting additional vertices along a path.

Manifold next event estimation [HDF15] (MNEE) allows sampling through a refractive interface by finding a point on the surfaces that satisfies specular constraints. Similarly, [WZHB09] address the case of single scattering through a specular interface. Koerner et al. [KNK*16] introduce an additional vertex to allow for better sampling across a dielectric boundary after subsurface scattering. Specular next event estimation [LZHJ20] builds upon an effective parametrisation to construct a sampling data structure to handle specular and glossy light transport over one vertex. Weber et al. [WHD17] create sub-paths with multiple vertices. Our method enforces a geometric constraint on the path by first sampling the phase function at the vertex location $\boldsymbol{x}_1$ and then filling in the rest of the path, and may be seen as a volumetric analogue of MNEE.

Rays, beams, photon volumes. Novák et al. [NNDJ12] compute global illumination in volumes using virtual ray lights, i.e. 1D line primitives that act as light sources to connect transport paths from the eye to. They also propose a scheme to importance sample the phase function of the extra scattering vertex $\boldsymbol{x}_1$ using discretisation and tabulation along fixed ray light and viewing directions. In contrast, we first sample the phase function in the middle and adjust the directions at both other vertices accordingly. The idea of using larger, non-point geometric primitives as loci of geometric connections was later extended to planes and volumes [DJBJ19].

Once-more collided flux. The neutron transport literature contains a few interesting and highly related works. In particular, Kalos [Kal63] proposed a framework to extend next event estimation by inserting an additional scattering vertex in between the current end of the path history and the sensor, to estimate the *once-more collided flux*. Our technique is similar in spirit, but they sampled a fixed outgoing direction first and only considered a non-analog PDF for the distance to the once-more collided vertex. Later works seem to indicate that this approach does not work well for anisotropic scattering [ABN01]. Kalli and Cashwell [KC77] combine Kalos' once-more collided flux estimator with a directional resampling technique, which, as we do, changes the outgoing direction at a scattering vertex to a non-analog random walk [SK71]. This biasing

scheme, however, does not explicitly sample the phase function but aims to gently guide the random walk closer to a sensor.

3. Method

In this paper we will focus on improving next event estimation within a homogeneous thin volume. We want to efficiently render a single extra scattering step in a volume before the light reaches the camera. This causes some characteristic blur, which, similar to depth of field but with a larger radius, changes over depth. For this, we extend next event estimation (NEE) for light tracing, connecting a path vertex to the camera aperture. It can be applied completely analogously, to path tracing, constructing connections to the light source. Specifically, we introduce *once-more scattered next event estimation* (OMNEE), augmenting the NEE segment with an additional scattering vertex.

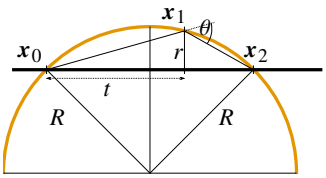

Figure 5: *The geometry of the three path vertices x_0, x_1, x_2 and the angle θ between the two segments, sampled according to the phase function $f_s(x_1)$. The arc radius R, the polar coordinate r and the fractional distance t assume the path geometry has been normalised such that $\overline{x_0 x_2} = 1$.*

Overview. We call the last path vertex of the light path x_0. This can be for instance a vertex on the caustic on the back of the whale in Figure 1 or a vertex on a sunbeam in the medium. The NEE vertex on the camera aperture is called x_2. We will insert a third vertex x_1 between these.

First, we sample the phase function $f_s(x_1)$ to obtain an angle θ between the segments $x_0 x_1$ and $x_1 x_2$ (see Figure 5). This factors out the material and enables us to use standard sampling routines for any forward scattering phase function. This ensures that one of the dominant terms can be handled effectively.

Next, we sample the centre vertex x_1 somewhere on a plane orthogonal to $x_0 x_2$, sliding the plane between the two end vertices. We know where to look for x_1 given θ: the inscribed angle theorem tells us it will be on a circular arc from x_0 to x_2 with centre somewhere off the direct connection. The radius R of this arc can be inferred from the angle θ. In homogeneous media, the flux through the vertex x_1 will also be rotationally symmetric around the axis x_0, x_2 if the phase function depends on θ only.

Analysis of the integrand. To compute this flux transported from x_0 to x_2 we need to integrate over all possible x_1 in (3D) vertex area measure:

$$I(x_0 \leftrightarrow x_2) = \int_{x_1} f(x_0 \leftrightarrow x_1 \leftrightarrow x_2)\,\mathrm{d}x_1, \tag{2}$$

where the integrand is the measurement contribution function

$$f(x_0 \leftrightarrow x_1 \leftrightarrow x_2) = f_r(x_0)\cdot G(x_0,x_1)\cdot G(x_1,x_2)\cdot \mu_s \cdot f_s(x_1)\cdot T(x_0,x_1)\cdot T(x_1,x_2)\cdot W(x_2), \tag{3}$$

where G and T denote the geometry terms and transmittance between two vertices. We observe that in the setting of classic, forward scattering scattering media, the following terms do not usually change rapidly i.e. are less important to sample:

$$\mu_s \cdot T(x_0,x_1)\cdot T(x_1,x_2)\cdot W(x_2). \tag{4}$$

The scattering coefficient μ_s will only change in heterogeneous media. The transmittances are more or less constant because they depend on the total distance travelled, i.e. $d_1 + d_2$ if we call $d_1 = |x_1 - x_0|, d_2 = |x_2 - x_1|$. This is not varying significantly for very forward scattering media. In this case, for typical values of θ, the triangle x_0, x_1, x_2 is very flat. For instance, for $\cos\theta = 0.98$ we have a maximum of $d_1 + d_2 \approx 1.005 \cdot |x_2 - x_0|$. The camera responsivity W is essentially the sensor's vignetting.

This leaves us with the following terms to sample:

$$G(x_0,x_1)\cdot G(x_1,x_2)\cdot f_s(x_1) = \frac{\cos\theta_0}{d_1^2}\cdot\frac{\cos\theta_2}{d_2^2}\cdot f_s(x_1) \tag{5}$$

$$= \underbrace{\cos\theta_0\cdot\cos\theta_2}_{\text{not varying much}}\cdot\frac{f_s(x_1)}{d_2^2\cdot d_1^2} \tag{6}$$

The cosine part in the G terms remains roughly invariant, again because the angular spread is very focused in highly forward scattering media. Finally, we split the measurement contribution into an *almost constant* part f_c and an *interesting* part f_i:

$$f_c(x_0 \leftrightarrow x_1 \leftrightarrow x_2) = \cos\theta_0\cdot\cos\theta_2 \cdot \mu_s \cdot T(x_0,x_1)\cdot T(x_1,x_2)\cdot W(x_2), \tag{7}$$

$$f_i(x_0 \leftrightarrow x_1 \leftrightarrow x_2) = \frac{f_s(x_1)}{d_2^2\cdot d_1^2}. \tag{8}$$

Note that f_i has singularities on both ends (at the vertices x_0 and x_2), where either of the distances d approaches zero. In contrast to the sum of both distances, this term is important to sample. To derive a sampling strategy, we will need to sample three dimensions to arrive at all possible x_1:

$$I(x_0 \leftrightarrow x_2) = \int_w \int_v \int_u f(x_0 \leftrightarrow x_1 \leftrightarrow x_2)\,\mathrm{d}u\,\mathrm{d}v\,\mathrm{d}w \tag{9}$$

where u, v, w are defined by an orthonormal coordinate system that aligns w with the direct connection between x_0 and x_2.

3.1. A change of variables

We intend to sample θ according to the phase function first, and then fill in the required two dimensions to yield a 3D vertex x_1. That is, we are looking for a change of variables (u, v, w) to $(\theta, ?, ?)$. We chose to sample a normalised, fractional distance $t \in (0, 1)$ along w and – as an intermediate step – express the other two dimensions in polar coordinates, since the scattering is rotationally symmetric around the axis $x_0 x_2$. Thus, we introduce a new local coordinate system to parameterise x_1: $t \in (0, 1)$ is the fractional distance between x_0 and x_2, r is the radius (i.e. related to the distance from x_1 to the

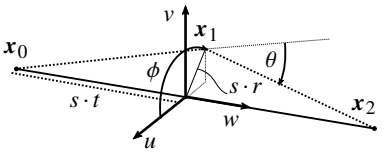

Figure 6: *The local vertex area measure coordinate frame (u,v,w) and the rotationally symmetric scattering along ϕ.*

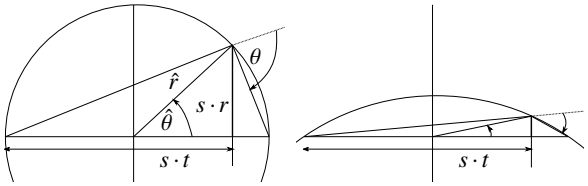

Figure 7: *The spherical coordinate system $(\hat\theta,\hat r,\hat\phi)$, illustrated for two values of θ. Left: close to $\pi/2$, right: more forward scattering. The longitudinal $\hat\phi=\phi$ rotates this 2D illustration around the axes of $s\cdot t$.*

line $\boldsymbol{x}_0\boldsymbol{x}_2$), and ϕ is the polar angle (see Figure 6). This radius r can be expressed in terms of fractional distance t as well as phase function angle θ to complete our change of variables from vertex area measure (u,v,w) to (θ,t,ϕ).

Also, we introduce a scaling variable $s=\overline{\boldsymbol{x}_0\boldsymbol{x}_2}$, the distance from the shading point $\boldsymbol{x}_0$ to the camera vertex $\boldsymbol{x}_2$. We use this to normalise the geometry of the scattering events by dividing out this scale factor. This affects R,r and t (see Figure 5).

We proceed by expressing the normalised polar coordinate radius r in terms of θ (i.e. we assume the geometry has been scaled by s for unit distance between $\boldsymbol{x}_0$ and $\boldsymbol{x}_2$). The radius R of the circumcircle of the triangle $\boldsymbol{x}_0\boldsymbol{x}_1\boldsymbol{x}_2$ (see Figure 5) can be determined as

$$R(\theta)=\frac{\overline{\boldsymbol{x}_0\boldsymbol{x}_2}}{2s\sin(\pi-\theta)}=\frac{1}{2\sin(\theta)}. \tag{10}$$

From this we can compute the normalised polar radius r as

$$r(t,R)=\sqrt{R^2-(1/2-t)^2}-\sqrt{R^2-1/4}, \tag{11}$$

$$r(t,\theta)=\sqrt{\frac{1}{4\sin^2\theta}-(1/2-t)^2}-\sqrt{\frac{1}{4\sin^2\theta}-1/4}. \tag{12}$$

As another intermediate step, we will next go to regular spherical coordinates $(\hat r,\hat\theta,\hat\phi=\phi)$ (see Figure 7):

$$\int_S f\,\mathrm{d}\boldsymbol{x}=\int_{\hat\phi}\int_{\hat\theta}\int_{\hat r} f\cdot\left|\frac{\partial\boldsymbol{x}}{\partial(\hat\theta,\hat r,\hat\phi)}\right|\mathrm{d}\hat r\,\mathrm{d}\hat\theta\,\mathrm{d}\hat\phi \tag{13}$$

$$=\int_{\hat\phi}\int_{\hat\theta}\int_{\hat r} f\cdot|\hat r^2\sin\hat\theta|\,\mathrm{d}\hat r\,\mathrm{d}\hat\theta\,\mathrm{d}\hat\phi. \tag{14}$$

This change of variables separates out the rotation angle ϕ and can naturally limit the domain of integration to a sphere with centre in the middle between $\boldsymbol{x}_0$ and $\boldsymbol{x}_2$ and radius $s/2$. This is the locus of all forward scattering events, i.e. with $\theta\le\pi/2$.

To complete our change of variables from $\boldsymbol{x}_1$ to (θ,t,ϕ), we need to express $\hat\theta$ and $\hat r$ in terms of θ and t. We use the following identities:

$$\phi=\hat\phi \tag{15}$$

$$s(t-1/2)=\hat r\cos\hat\theta \tag{16}$$

$$\sin\hat\theta=s\cdot r/\hat r \tag{17}$$

$$\hat r^2=s^2(t-1/2)^2+s^2r^2. \tag{18}$$

To compute the Jacobian determinant, we are now looking for a few

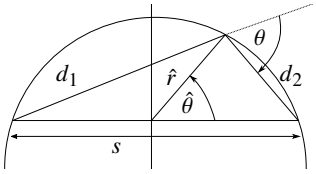

Figure 8: *To compute θ, we can use the law of cosines: $s^2=d_1^2+d_2^2-2d_1d_2\cos\gamma$ where $\gamma=\pi-\theta$ is opposite of the edge s.*

partial derivatives:

$$|J|=\left|\frac{\partial(\theta,t)}{\partial(\hat r,\hat\theta)}\right|=\left|\frac{\partial\theta}{\partial\hat r}\cdot\frac{\partial t}{\partial\hat\theta}-\frac{\partial\theta}{\partial\hat\theta}\cdot\frac{\partial t}{\partial\hat r}\right|. \tag{19}$$

The fractional distance t and its derivatives are simple:

$$t(\hat r,\hat\theta)=\frac{\hat r}{s}\cos\hat\theta+1/2 \tag{20}$$

$$\partial t/\partial\hat\theta=-\frac{\hat r}{s}\sin\hat\theta \tag{21}$$

$$\partial t/\partial\hat r=\frac{1}{s}\cos\hat\theta. \tag{22}$$

We compute θ from d_1,d_2 and s using the law of cosines (Figure 8). It remains to express the distances d_1 and d_2 in terms of $\hat\theta$ and $\hat r$:

$$d_2^2=\hat r^2+\frac{s^2}{4}-\hat r s\cos\hat\theta \tag{23}$$

$$d_1^2=\hat r^2+\frac{s^2}{4}-\hat r s\cos(\pi-\hat\theta)=\hat r^2+\frac{s^2}{4}+\hat r s\cos\hat\theta, \tag{24}$$

and we also simplify the expression for their product

$$d_1^2\cdot d_2^2=-\frac{16\hat r^2s^2\cos^2\hat\theta-s^4-8\hat r^2s^2-16\hat r^4}{16} \tag{25}$$

$$=-\frac{1}{16}\left(8\hat r^2s^2(2\cos^2\hat\theta-1)-16\hat r^4-s^4\right) \tag{26}$$

$$=\frac{1}{16}\left(16\hat r^4+s^4-8\hat r^2s^2\cos2\hat\theta\right). \tag{27}$$

Now we are ready to compute θ from $\hat{\theta}$ and $\hat{r}$ as

$$\theta(\hat{r},\hat{\theta}) = \pi - \arccos\left(\frac{d_1^2+d_2^2-s^2}{2d_1d_2}\right) \tag{28}$$

$$= \pi - \arccos\left(\frac{\hat{r}^2-s^2/4}{d_1d_2}\right) \tag{29}$$

$$= \pi - \arccos\left(\frac{4\hat{r}^2-s^2}{\sqrt{16\hat{r}^4+s^4-8\hat{r}^2s^2\cos 2\hat{\theta}}}\right), \tag{30}$$

as well as the partial derivatives required for the Jacobian in Equation (19)

$$\partial\theta/\partial\hat{r} = \frac{4s(4\hat{r}^2+s^2)\sin\hat{\theta}}{16\hat{r}^4+s^4-8\hat{r}^2s^2\cos 2\hat{\theta}}, \tag{31}$$

$$\partial\theta/\partial\hat{\theta} = \frac{4\hat{r}s(s^2-4\hat{r}^2)\cos\hat{\theta}}{16\hat{r}^4+s^4-8\hat{r}^2s^2\cos 2\hat{\theta}}. \tag{32}$$

3.2. A preliminary estimator

With the identities we collected so far, we construct the following Monte Carlo estimator, by sampling the angle $\phi \sim U(0,2\pi)$ uniformly on the circle, $t \sim U(0,1)$ uniformly between $\boldsymbol{x}_0$ and $\boldsymbol{x}_2$, and $\theta \sim f_s(\theta)\sin\theta$ by the phase function f_s. For the latter, we use the standard solid angle sampling procedures as are presumably already present in a rendering system for analog Monte Carlo random walks. Since this samples a direction ω on the sphere, it also accounts for the Jacobian determinant $\sin\theta$, which will thus be part of our PDF. The estimator for the interesting part f_i becomes

$$\hat{I} = \frac{f_i(\boldsymbol{x}_1)}{p(\boldsymbol{x}_1)} = \frac{1}{d_1^2\cdot d_2^2\cdot|J|}\cdot\frac{|\hat{r}^2\sin\hat{\theta}|}{\sin\theta}, \tag{33}$$

and we furthermore insert the Jacobian determinant $|J|$ and analytically cancel with the other terms. A bit of rearranging yields

$$d_1^2\cdot d_2^2\cdot|J| = \left|\hat{r}^3\cos 2\hat{\theta} - \frac{\hat{r}s^2}{4}\right|. \tag{34}$$

and inserting this into Equation (33) yields an estimator that no longer requires evaluating the reciprocal squared distances

$$\hat{I} = \left|\hat{r}^3\cos 2\hat{\theta} - \frac{\hat{r}s^2}{4}\right|^{-1}\cdot\frac{\hat{r}^2|\sin\hat{\theta}|}{\sin\theta} \tag{35}$$

$$= \left|\frac{4\hat{r}\sin\hat{\theta}}{4\hat{r}^2\cos 2\hat{\theta}-s^2}\right|\cdot\frac{1}{\sin\theta}. \tag{36}$$

We can compute all the variables in this estimator from the sampled (θ,t,ϕ) random variables as defined above. They range in the following intervals $\hat{r}\in[0,s/2]$, $\hat{\theta}\in[0,\pi]$, and $\theta\in[0,\pi/2]$. This works, but Equation (36) has an unpleasant singularity at $\theta=0$, the important forward scattering case.

3.3. Sampling the fractional distance

We have one more degree of freedom to sample the remaining weight in the estimator, by changing the PDF of the fractional distance t. We have established that sampling θ according to the phase function, and uniformly sampling t results in Equation (36). We can express this in terms of (θ,t), and for fixed θ sample t conditional on θ.

Inserting Equations (12) and (15) to (18) and simplifying using a computer algebra program (see `math.pdf` in the supplemental material), these are the results:

$$A := \cot\theta - \sqrt{\sin^{-2}\theta-(2t-1)^2} \tag{37}$$

$$\hat{p}(t|\theta) = \frac{A}{s\sin\theta(4(t-1)t-\cot\theta\cdot A)} \tag{38}$$

$$\int \hat{p}(t|\theta)\,\mathrm{d}t = \frac{\theta}{s\sin\theta} =: c \tag{39}$$

$$p(t|\theta) = \hat{p}(t|\theta)/c = \frac{A}{\theta(4(t-1)t-\cot\theta\cdot A)} \tag{40}$$

$$P^{-1}(\xi|\theta) = \cos(\theta-\xi\theta)\sin(\xi\theta)/\sin\theta. \tag{41}$$

In summary, the full estimator is constructed in the following steps:

- sample an outgoing direction at $\boldsymbol{x}_1$ via the regular solid angle phase function sampling routine, store θ,
- sample the fractional distance $t\sim p(t|\theta)$ via Equation (41)
- sample $\phi\sim\frac{1}{2\pi}$ (or reuse from phase function sampling)
- reconstruct location of $\boldsymbol{x}_1$ using t and r from Equation (12).

The vertex area measure PDF and the estimator are

$$p(\boldsymbol{x}_1) = f_s(\boldsymbol{x}_1)\frac{s}{d_1^2\cdot d_2^2}\frac{\sin\theta}{\theta} \tag{42}$$

$$\hat{I} = f_c(\boldsymbol{x}_0\leftrightarrow\boldsymbol{x}_2)\cdot\frac{\theta}{s\sin\theta}. \tag{43}$$

Note how the weight of the Monte Carlo estimator contains $\theta/\sin\theta$. This suggests that even better importance sampling can be achieved by not using a generic phase function sampler that perfectly samples the phase function in solid angle (incurring the spherical Jacobian determinant $|\sin\theta|$) but one that samples by θ instead. However, $\lim_{\theta\to0}\sin\theta/\theta = 1$ and for highly forward scattering media $\sin\theta\approx\theta$, so we expect only moderate advantages for a significant implementation effort here. In a similar note, when evaluating the inverse CDF $P^{-1}(\xi|\theta)$ for $\theta=0$, we consider that $\lim_{\theta\to0}\sin(\xi\theta)/\sin\theta = \xi$.

3.4. Forward scattering phase functions

Phase functions in thin, homogeneous media like water or fog tend to be strongly forward scattering with only a small backscattering component. Dedicated phase functions are used to describe such media; a Mie phase function can simulate scattering of small dielectric spheres like water droplets and a Fournier-Forand [FF94] phase function can describe ocean water. Figure 9 illustrates their shape and shows a fit to the more commonly used Henyey-Greenstein phase function. While the mean cosine of ocean water or water vapour is normally around $g=0.85$, it should be noted that the forward scattering, when separated from the backward scattering component, often has a mean cosine of $g=0.95$ and higher, which makes classic NEE techniques very inefficient.

Our technique works best for phase functions that are purely forward-scattering phase functions. However, one advantage of our method is that we can use it with *any* phase function that has at least

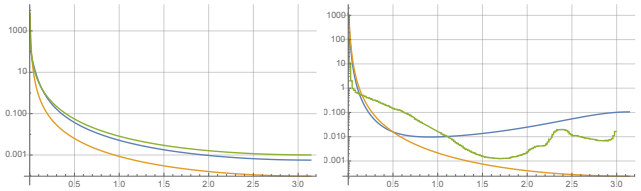

Figure 9: *Fournier-Forand (left) and Mie (right) phase functions fitted to a single Henyey-Greenstein (orange) and a two-term Henyey-Greenstein (blue) phase function [GSMA08]. For the single, $g = 0.99527$ resp. $g = 0.988264$; the double has $g1 = 0.997229$ and $g2 = 0.959952$ with a weight of 0.339242 and $g1 = 0.990344$ and $g2 = -0.439579$ with a weight of 0.712146. The green line is the ground truth phase function. A good summary of existing analytical phase functions can be found in [Sha15].*

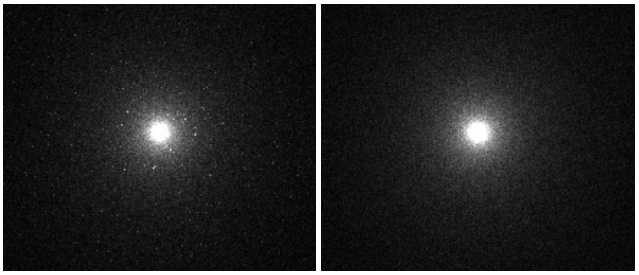

Figure 10: *A point emitter in an infinite homogeneous volume with $g = 0.5$. Left: Sampling the full phase function and rejecting samples from the backward hemisphere. Right: Samples are drawn from the forward hemisphere only.*

some forward-scattering. For phase functions that have a backscattering component like the classic Henyey-Greenstein (HG) phase function

$$f_{HG}(\mu) = \frac{1}{2}\frac{1-g^2}{(1+g^2-2g\mu)^{3/2}} \tag{44}$$

this means that we need to use rejection sampling for our NEE technique. See Section 5 for a more detailed discussion on how to handle this case.

However, for some phase functions like HG it is easy to derive a routine that only considers its forward hemisphere, which avoid wasting samples. Note that this is a pure performance optimisation; a special phase function is not needed for our technique, and we can always fall back to rejection sampling. By introducing the constraint $\mu \geq 0$, we will get *HG forward* (HGF):

$$f_{HGF}(\mu) = -\frac{(-1+g)g\sqrt{1+g^2}}{(-1+g+\sqrt{1+g^2})(1+g^2-2g\mu)^{3/2}} \tag{45}$$

which integrates to unity over $[0,1]$. The CDF of Equation (45) is invertible and allows us perfect importance sampling:

$$\mu = \frac{\xi(1+g^2)(1+y+g(-2+\xi+g-y))}{1+y+g(z-2\xi(-1+g)(1-g+y))} \tag{46}$$

$$y = \sqrt{1+g^2} \tag{47}$$

$$z = 4g+2\xi^2 g-3(1+y)-(-3+g)g(-g+y) \tag{48}$$

where ξ is a random number. As it can be seen in Figure 10, partitioning the phase function into its hemispheres will further decrease noise. For sake of completeness, a similar formula can be found for the backward component:

$$\mu = \frac{(-1+\xi)(1+g^2)(-\xi+g^2+\xi y+gy)}{g(2\xi-2\xi^2+g+2\xi g^2+g^3+y-2\xi y+2\xi^2 y+2\xi gy+g^2 y)} \tag{49}$$

4. Results

We evaluate our technique in two settings: a spectral CPU rendering framework and an isolated unit test simulating a point light source in an infinite thin homogeneous medium and paths with just three vertices (sensor, scattering, light).

Figure 13 shows a final frame comparison of a render of a whale under water. Both the distance in the water and the brightness of the object influence the characteristic blur. The bottom row shows filtered images, where firefly noise has been removed. This shows how the blur effect disappears for other methods, i.e. is unrenderable.

The beauty renders in Figures 1, 2 and 13 to 15 are computed via light tracing, i.e. starting transport paths at the light sources and connecting to the camera via next event estimation. On top of this, equiangular sampling introduces a two-vertex connection to the camera, and our technique does the same but chooses the centre vertex $\boldsymbol{x}_1$ according to our once-more scattering equations. We use the forward-only HG phase function from Equation (44). Render times were about 310ms for one progression with one sample per pixel (plain light tracing), 430ms for equiangular sampling, and 410ms for our technique. The difference between the last two is not significant, the extra work to construct a two-vertex path suffix and combine it using MIS is mostly equivalent.

Figure 11 presents a unit test comparison of our work to Villeneuve et al. [VGGN21]. Here we use the classic full HG phase function form Equation (44). We compare against two different approaches, one better suited for a low g and the other one better for higher g. In this setting, we combined our method (forward scattering) with equiangular sampling [KF12] (backward scattering) to yield an unbiased estimator for the complete domain. All images use 4spp. The render time of our method is split roughly equally between equiangular and once-more scattering NEE. Note that we do not perform MIS in this test, our goal is solely to obtain an unbiased estimator. The figure shows results for different mean cosines between isotropic and highly forward scattering. Our method is simpler and faster, and a lot more efficient at sampling for large mean cosines.

A comparison of just our technique without backscattering can be seen in Figure 12. Here we evaluate only the forward scattering hemisphere (i.e. only the domain of our sampling technique, the

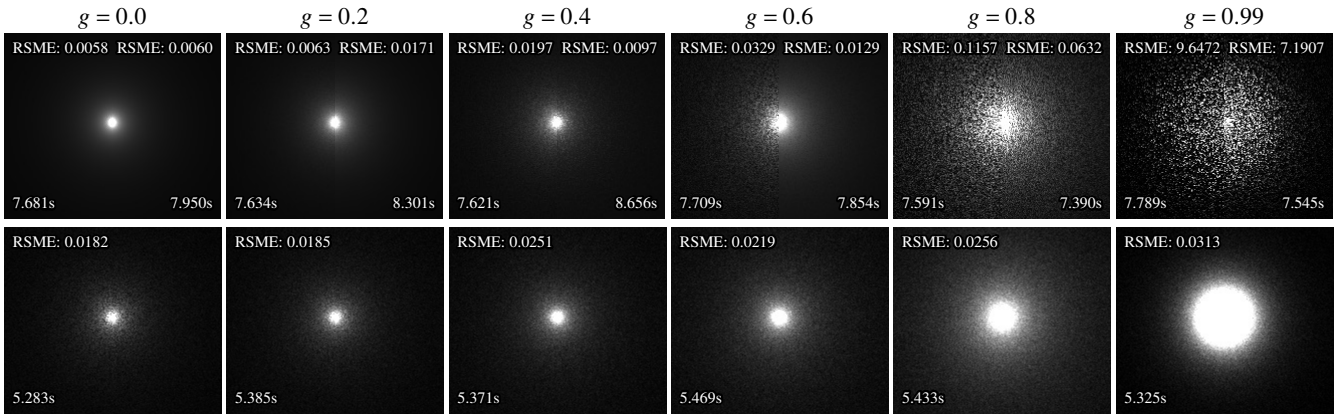

Figure 11: *Comparison of Villeneuve et al. [VGGN21] with (top) their $\mathcal{B}_P(\mathcal{T}_T)$ (Bézier warp for phase function and Taylor series for transmittance) and their $\mathcal{B}_T(\mathcal{T}_P)$ (Bézier warp for transmittance and Taylor series for phase function) and ours (bottom). g increases from 0.0 to 0.99 in 0.2 steps from left to right, all images use 4spp. While our method is less suited for the isotropic case where backward scattering is significant and we cannot use our technique, it starts to overtake [VGGN21] once the forward scattering becomes dominant. Note that our method has a better performance, even when combined with the overhead from equiangular sampling for the backward component.*

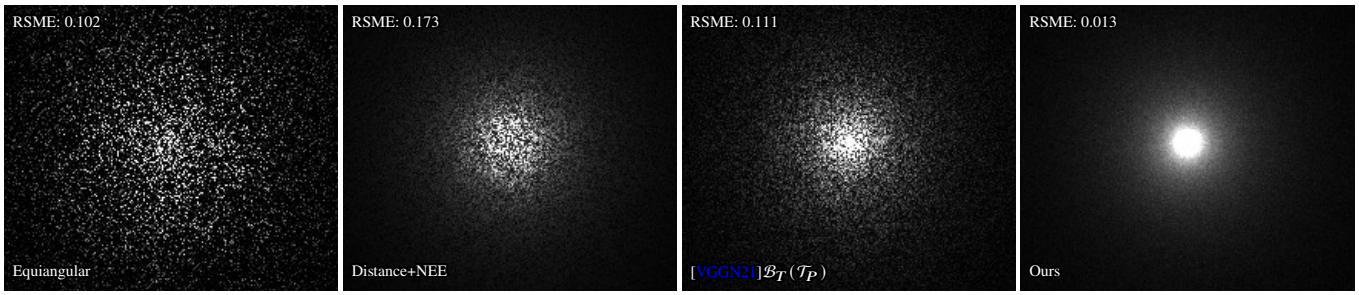

Figure 12: *A unit test with an infinite homogeneous medium and a point emitter, mean cosine $g = 0.9$. Here, only forward scattering is enabled and 4 samples per pixel are simulated. The images show only three path vertices, i.e. single scattering in the medium. The reference for the RMSE numbers was distance sampling with next event estimation at 10k samples per pixel.*

backward scattering hemisphere has been set to zero) with $g = 0.9$. The images show single scattering, they are limited to max. 3 path vertices. This test shows the performance of our estimator in isolation, without diluting the differences by MIS combination.

When scattering off an object, our technique works best on diffuse surfaces. Figure 14 shows the same scene but the whale has now a glossy gold material with Beckmann roughness 0.01. We consider this as a failure case for our method, since we do not sample the BSDF at the vertex $\boldsymbol{x}_0$ on the whale. Indeed variance increases, but our method will still produce less noise compared to others. This can be seen when removing firefly noise, which shows that we can still capture more of the glow than equiangular sampling or light tracing.

Figure 15 shows low order highly forward scattering in atmospheric haze. The glow around specular highlights as well as the more subtle glow above the desert surface in the background can only be reproduced by importance sampling the phase function.

As an additional result, Figure 16 shows our technique integrated in a production rendering system. These images showcase the adjoint connection, next event estimation applied to path tracing, connecting to the emitters. In this scene with torchlights under water, the visual features and the noise reduction of our technique are comparable to what we demonstrated with light tracing before.

5. Discussion and Limitations

We presented a technique to efficiently render a specialised effect: the point-spread-function-like glow around objects in thin, homogeneous, forward scattering media.

We disregard transmittance in our sampling scheme. As discussed briefly earlier, the total distance travelled from the shading point $\boldsymbol{x}_0$ to the next event $\boldsymbol{x}_2$ via a once-more scattered vertex is not much longer than the direct connection in the case of highly forward scattering. For isotropic or backward scattering, this is not the case and thus our method will likely not present an optimal estimator in this setting. Likewise, phase functions that do not pose a strong

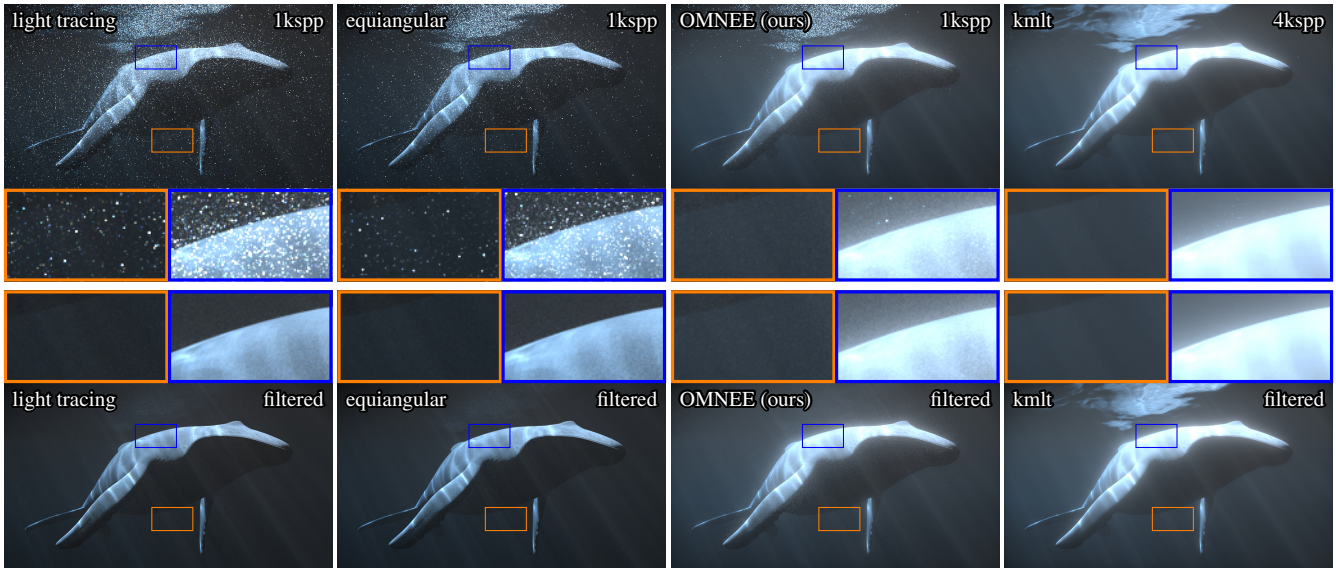

Figure 13: *A whale rendered with max. five path vertices. While equiangular sampling [KF12] improves the quality of the sunbeam towards camera a bit, it does not sample the phase function explicitly. The characteristic glow around objects under water is thus extremely noisy and has to be filtered out for final frames. The bottom row is filtered via density based outlier rejection [ZHD18].*

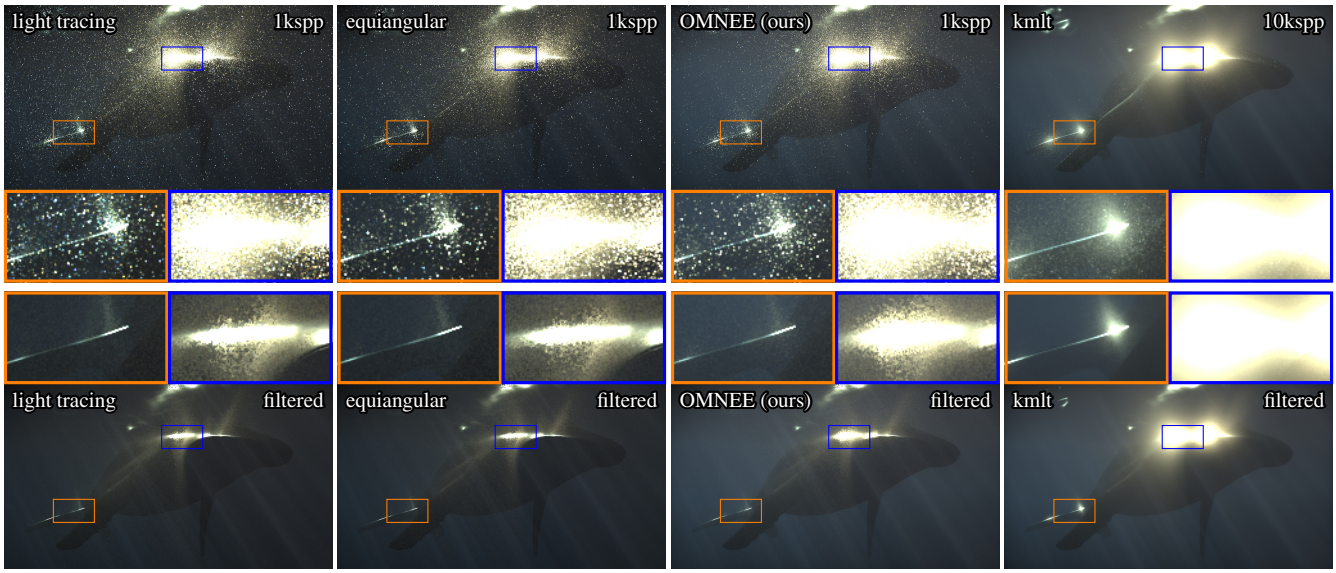

Figure 14: *A failure case for our method. The glossy whale presents another important directional constraint at the shading vertex on its back. Since we do not importance sample this, variance is increased. Again, as in Figure 13, the bottom row shows the images with fireflies filtered out (i.e. no spatial denoising has been applied). The last column shows primary sample space Metropolis light transport (kmlt) used on top of our sampling strategy.*

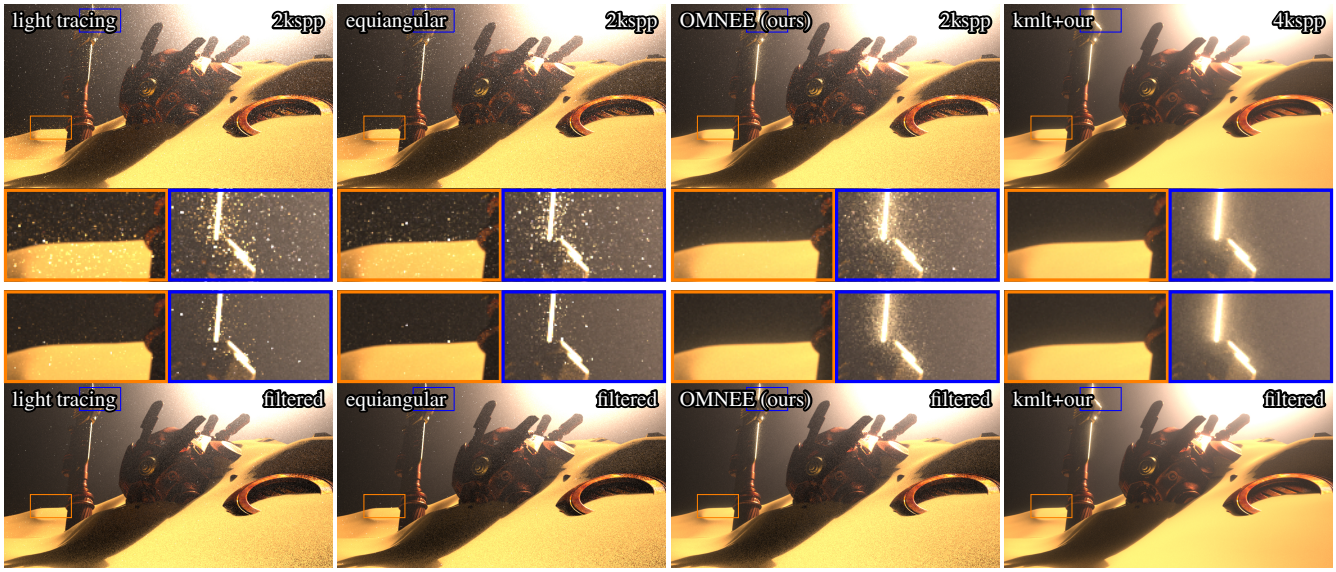

Figure 15: *A desert scene with glossy metal objects, enclosed in an infinite homogeneous forward scattering medium ($g = 0.98$). Even though the additional glossy constraint on the metal surfaces (g around 0.1) is not importance sampled, we can reproduce the glow in the medium here. The last column shows primary sample space Metropolis light transport (kmlt) used on top of our sampling strategy.*

Figure 16: *Once-more scattered next event estimation evaluated towards a light source, i.e. in path tracing from the camera. The images show the importance of the extra scattering vertex on top of single scattering only (right), and also demonstrate how out technique can robustly reduce the noise in this setting. All images are 1024 samples per pixel.*

angular constraint on the path are not crucial to sample and thus transmittance and our sampling are complementary.

We are not explicitly addressing heterogeneous media but treat it by Monte Carlo sampling. The improved importance sampling can still be expected to yield faster convergence over the base line, similar to what has been evaluated before [GKH*13, Fig. 10].

The adjoint effect blurs out shadows for pathtracing, but in our evaluation we focused on the more prominent visual features towards the camera. Thus, light tracing or pathtracing with camera reconnections [BSH02; WGGH20; THD17] are a prerequisite for the application of our technique.

We made some assumptions about the phase functions. In particular, our technique only handles the forward scattering part of a phase function and only the circularly symmetric case (i.e. the phase function only depends on θ, there cannot be an anisotropic medium in the sense of Jakob et al. [JAM*10]). For now we need to either renormalise the importance sampling to only the forward scattering hemisphere or discard samples that are generated in the backward facing hemisphere. For highly forward scattering media, such as we are interested in analysing in here, the amount of backward scattering is usually very small.

The sampling method is unbiased in the sense that NEE is unbiased too: it is an estimator only for the domain it samples. NEE will only sample the light sources and requires another technique to sample indirect contributions. Similarly, once-more scattering NEE requires another technique to sample the backward scattering contributions. Only the combined estimator via MIS will be an unbiased estimate for the whole transport domain (see Figure 11).

We do not consider the BSDF at the path vertex $\boldsymbol{x}_0$. In case of glossy BSDFs this may be a dominating term. Equiangular sampling will consider this BSDF, however, this does not on its own present a good estimator (see Figure 14).

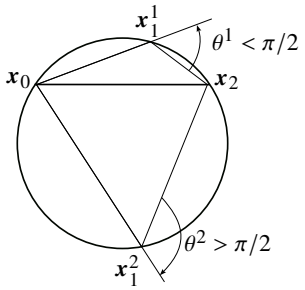

Figure 17: *Illustration of the extended geometry required to include backward scattering ($\theta > \pi/2$) in our framework. The part of the drawing above the x_0x_2 line is the case we consider: the centre of the circle is on the other side of the line as x_1^1. If the centre is on the line, the resulting θ will be exactly $\pi/2$ everywhere on the circle. In the bottom part, x_1^2 results in $\theta^2 > \pi/2$, but the angle θ will again be constant for all x_1 on the lower arc. This arc unfortunately bulges outside the reach of our parameter t on both sides, so an angular parameterisation has to be chosen.*

6. Conclusion and Future work

We demonstrated the once-more scattered next event estimation technique for light tracing. To efficiently simulate subtle blurs on shadows it can also be used for next event estimation in path tracing, to connect to the light sources (not to the camera as in light tracing).

For the special case of connecting to the environment map, the sampling routine would simplify substantially because the vertex x_2 on an environment map is really a directional constraint then. It is thus possible to directly sample the outgoing direction at x_0 following the BSDF.

Our choice of the parameterisation of $t \in (0,1)$ may not be the optimum. It is plausible that an angular parameter on the circumcircle around x_0, x_1, x_2 could serve the same purpose, and at the same time generalise to backscattering contributions, i.e. allow us to sample vertices x_1 outside the sphere defined by $\theta \in [0, \pi/2]$, see Figure 17. This circumcircle would grow larger than radius $s/2$ for $\theta > \pi/2$, and it would be necessary to re-derive the formulas in Section 3.3, which may be possible.

Furthermore, it would be interesting to extend the sampling to include the product with the BSDF at x_0, to improve glow around glossy objects. In fact this could be a starting point to extend the scattering geometry to longer paths, using cascades of intertwined circles. This would be potentially useful for highly scattering and dense media such as skin.

Acknowledgements

We would like to thank Eugene d'Eon, Benedikt Bitterli and Jiří Vorba for their valuable feedback. Also we thank the anonymous reviewers for input and the new title.

References

[ABN01] Authier, N., Both, J. P., and Nimal, J. C. "New photon biasing schemes for the "once-more-collided-flux-estimator" method". *Nuclear Science and Engineering* 137.2 (2001), 146–155. doi: 10.13182/NSE01-A2181 3.

[BSH02] Bekaert, Philippe, Sbert, Mateu, and Halton, John. "Accelerating path tracing by re-using paths". *Eurographics Workshop on Rendering*. Ed. by Debevec, P. and Gibson, S. The Eurographics Association, 2002. doi: 10.2312/EGWR/EGWR02/125-134 10.

[dEo16] d'Eon, Eugene. *A Hitchhiker's Guide to Multiple Scattering*. 2016. url: http://www.eugenedeon.com/hitchhikers 3.

[DJBJ19] Deng, Xi, Jiao, Shaojie, Bitterli, Benedikt, and Jarosz, Wojciech. "Photon surfaces for robust, unbiased volumetric density estimation". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 38.4 (July 2019). doi: 10/gf6rx9 3.

[DWWH20] Deng, Hong, Wang, Beibei, Wang, Rui, and Holzschuch, Nicolas. "A Practical Path Guiding Method for Participating Media". *Computational Visual Media* 6 (Mar. 2020), 37–51. doi: 10.1007/s41095-020-0160-1. url: https://hal.inria.fr/hal-02486014 3.

[FF94] Fournier, Georges R. and Forand, J. Luc. "Analytic phase function for ocean water". *Ocean Optics XII*. Ed. by Jaffe, Jules S. Vol. 2258. International Society for Optics and Photonics. SPIE, 1994, 194–201. url: https://doi.org/10.1117/12.190063 6.

[GKH*13] Georgiev, Iliyan, Křivánek, Jaroslav, Hachisuka, Toshiya, Nowrouzezahrai, Derek, and Jarosz, Wojciech. "Joint importance sampling of low-order volumetric scattering". *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia)* 32.6 (Nov. 2013), 164:1–164:14. doi: 10.1145/2508363.2508411 3, 10.

[GSMA08] Gutierrez, Diego, Seron, Francisco J., Munoz, Adolfo, and Anson, Oscar. "Visualizing underwater ocean optics". *Computer Graphics Forum* 27.2 (2008), 547–556. doi: 10.1111/j.1467-8659.2008.01152.x 7.

[HDF15] Hanika, Johannes, Droske, Marc, and Fascione, Luca. "Manifold next event estimation". *Computer Graphics Forum (Proceedings of Eurographics Symposium on Rendering)* 34.4 (June 2015), 87–97. doi: 10.1111/cgf.12681 3.

[HZE*19] Herholz, Sebastian, Zhao, Yangyang, Elek, Oskar, Nowrouzezahrai, Derek, Lensch, Hendrik P. A., and Křivánek, Jaroslav. "Volume path guiding based on zero-variance random walk theory". *ACM Transactions on Graphics* 38.3 (June 2019), 25:1–25:19. doi: 10.1145/3230635 3.

[JAM*10] Jakob, Wenzel, Arbree, Adam, Moon, Jonathan T., Bala, Kavita, and Marschner, Steve. "A radiative transfer framework for rendering materials with anisotropic structure". *ACM Transactions on Graphics* 29.4 (July 2010), 53:1–53:13 10.

[Kal63] Kalos, Malvin H. "On the estimation of flux at a point by Monte Carlo". *Nuclear Science and Engineering* 16 (1963), 111–117 3.

[KC77] Kalli, H. J. and Cashwell, E. D. *Evaluation of three Monte Carlo estimation schemes for flux at a point*. Tech. rep. LA–6865-MS. Los Alamos Scientific Lab., 1977 3.

[KF12] Kulla, Christopher and Fajardo, Marcos. "Importance sampling techniques for path tracing in participating media". *Computer Graphics Forum (Proc. of Eurographics Symposium on Rendering)* 31.4 (June 2012), 1519–1528. doi: 10.1111/j.1467-8659.2012.03148.x 3, 7, 9.

[KNK*16] Koerner, David, Novák, Jan, Kutz, Peter, Habel, Ralf, and Jarosz, Wojciech. "Subdivision next-event estimation for path-traced subsurface scattering". *Proceedings of EGSR (Experimental Ideas & Implementations)*. The Eurographics Association, June 2016. doi: 10/gf6rzj 3.

[LW96] Lafortune, Eric P. and Willems, Yves D. "Rendering participating media with bidirectional path tracing". *Proc. of Eurographics Workshop on Rendering Techniques*. Porto, Portugal: Springer-Verlag, Aug. 1996, 91–100 3.

[LZHJ20] Loubet, Guillaume, Zeltner, Tizian, Holzschuch, Nicolas, and Jakob, Wenzel. "Slope-space integrals for specular next event estimation". *ACM Transactions on Graphics* 39.6 (Nov. 2020). doi: 10.1145/3414685.3417811 3.

[NGHJ18] Novák, Jan, Georgiev, Iliyan, Hanika, Johannes, and Jarosz, Wojciech. "Monte Carlo methods for volumetric light transport simulation". *Computer Graphics Forum (Eurographics State of the Art Reports)* 37.2 (2018), 1–26. doi: 10.1111/cgf.13383 3.

[NNDJ12] Novák, Jan, Nowrouzezahrai, Derek, Dachsbacher, Carsten, and Jarosz, Wojciech. "Virtual ray lights for rendering scenes with participating media". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 31.4 (July 2012). doi: 10/gbbwk2 3.

[PSP09] Pegoraro, Vincent, Schott, Mathias, and Parker, Steven G. "An analytical approach to single scattering for anisotropic media and light distributions". *Graphics Interface*. Canadian Information Processing Society, 2009, 71–77. isbn: 978-1-56881-470-4 3.

[PSP10] Pegoraro, Vincent, Schott, Mathias, and Parker, Steven G. "A closed-form solution to single scattering for general phase functions and light distributions". *Computer Graphics Forum (Proc. of Eurographics Symposium on Rendering)* 29.4 (2010), 1365–1374. doi: 10/cdm26d 3.

[PSS11] Pegoraro, Vincent, Schott, Mathias, and Slusallek, Philipp. "A mathematical framework for efficient closed-form single scattering". *Graphics Interface*. Canadian Information Processing Society, 2011, 151–158. isbn: 978-1-4503-0693-5 3.

[Sha15] Sharma, Subodh Kumar. "A review of approximate analytic light-scattering phase functions". *Light Scattering Reviews 9* (2015), 53–100 7.

[SHJD18] Simon, Florian, Hanika, Johannes, Jung, Alisa, and Dachsbacher, Carsten. "Selective guided sampling with complete light transport paths". *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37.6 (Dec. 2018). doi: 10.1145/3272127.3275030 3.

[SK71] Steinberg, H. A. and Kalos, Malvin H. "Bounded estimators for flux at a point in Monte Carlo". *Nuclear Science and Engineering* 44 (1971), 406–412 3.

[SRNN05] Sun, Bo, Ramamoorthi, Ravi, Narasimhan, Srinivasa G., and Nayar, Shree K. "A practical analytic single scattering model for real time rendering". *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 24.3 (July 2005), 1040–1049. doi: 10/fgnbqt 3.

[THD17] Tessari, Lorenzo, Hanika, Johannes, and Dachsbacher, Carsten. "Local quasi-Monte Carlo exploration". *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations*. June 2017. doi: 10.2312/sre.20171196 10.

[VGGN21] Villeneuve, Keven, Gruson, Adrien, Georgiev, Iliyan, and Nowrouzezahrai, Derek. "Practical Product Sampling for Single Scattering in Media". *Eurographics Symposium on Rendering - DL-only Track*. Ed. by Bousseau, Adrien and McGuire, Morgan. The Eurographics Association, 2021. isbn: 978-3-03868-157-1. doi: 10.2312/sr.20211290 3, 7, 8.

[WGGH20] West, Rex, Georgiev, Iliyan, Gruson, Adrien, and Hachisuka, Toshiya. "Continuous multiple importance sampling". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 39.4 (July 2020). doi: 10.1145/3386569.3392436 10.

[WHD17] Weber, Pascal, Hanika, Johannes, and Dachsbacher, Carsten. "Multiple vertex next event estimation for lighting in dense, forward-scattering media". *Computer Graphics Forum (Proceedings of Eurographics)* 36.2 (Apr. 2017), 21–30. doi: 10.1111/cgf.13103 3.

[Wil77] Williams, M. M. R. "On the role of the adjoint Boltzmann equation in the calculation of energy deposition". *Journal of Physics D: Applied Physics* 10.17 (Dec. 1977), 2343–2346. doi: 10.1088/0022-3727/10/17/006 3.

[WZHB09] Walter, Bruce, Zhao, Shuang, Holzschuch, Nicolas, and Bala, Kavita. "Single scattering in refractive media with triangle mesh boundaries". *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 28.3 (July 2009). doi: 10.1145/1531326.1531398 3.

[ZHD18] Zirr, Tobias, Hanika, Johannes, and Dachsbacher, Carsten. "Reweighting firefly samples for improved finite-sample Monte Carlo estimates". *Computer Graphics Forum* 37.6 (2018), 410–421. doi: 10.1111/cgf.13335 9.