EUROGRAPHICS 2017 / L. Barthe and B. Benes
(Guest Editors)

Volume 36 (2017), Number 2

Multiple Vertex Next Event Estimation for Lighting in dense, forward-scattering Media

Pascal Weber, Johannes Hanika and Carsten Dachsbacher

Karlsruhe Institute of Technology, Germany

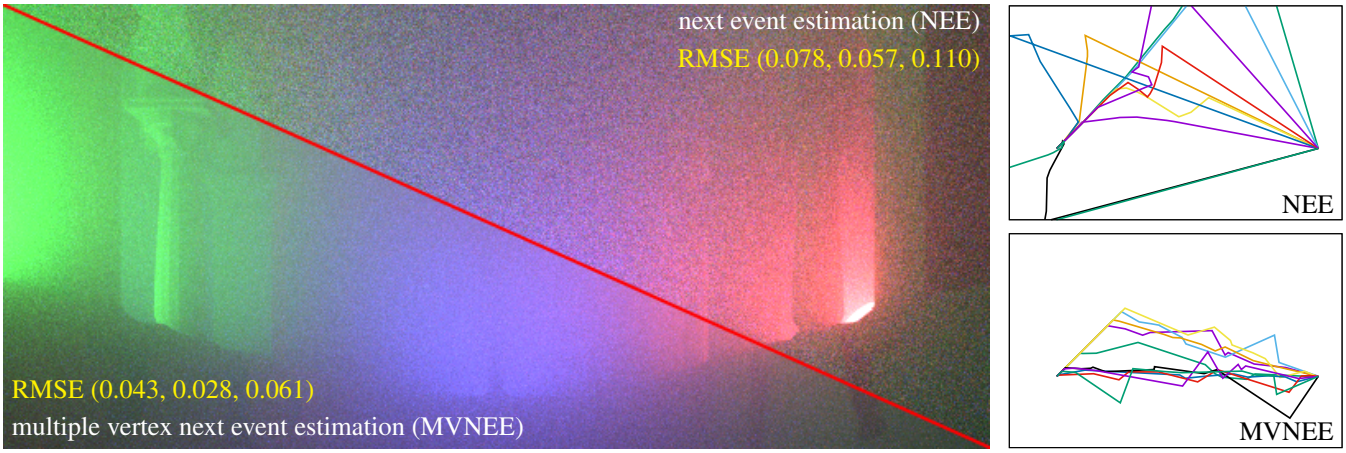

Figure 1: *Equal time comparison (60min) of path tracing with next event estimation (NEE, top) and Multiple Vertex Next Event Estimation (MVNEE, our technique, bottom) in the Sponza scene filled with a homogeneous medium with* $\mu_t = \mu_s = 0.2\frac{1}{m}$ *and an anisotropic Henyey-Greenstein phase function with a mean cosine of* $g = 0.9$*. The three area light sources are positioned at a distance of 33m, 18.4m and 12.2m respectively from the camera. To the right, ten paths are plotted (to scale), created to connect the left point with a 45-degree starting direction to the right point. NEE (top) results in a much worse distribution than MVNEE (bottom).*

Abstract

We present a new technique called Multiple Vertex Next Event Estimation*, which outperforms current direct lighting techniques in forward scattering, optically dense media with the Henyey-Greenstein phase function. Instead of a one-segment connection from a vertex within the medium to the light source, an entire sub path of arbitrary length can be created and we show experimentally that 4-10 segments work best in practice. This is done by perturbing a seed path within the Monte Carlo context. Our technique was integrated in a Monte Carlo renderer, combining random walk path tracing with multiple vertex next event estimation via multiple importance sampling for an unbiased result. We evaluate this new technique against standard next event estimation and show that it significantly reduces noise and increases performance of multiple scattering renderings in highly anisotropic, optically dense media. Additionally, we discuss multiple light sources and performance implications of memory-heavy heterogeneous media.*

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

1. Introduction

Interaction of light with participating media is one of the most important aspects of photorealistic rendering. Light gets scattered in fluids, clouds, fog, and even in air itself. In addition light can be scattered beneath the surface of a variety of materials, most prominently organic material like skin, plants and leafs, but also ceramics to name a few. The scattering effects add to the realism of a rendered image, missing scattering effects on the other hand become apparent.

Rendering participating media, however, is a complex and thus time-consuming task [Jen01], which is why a lot of research went into developing new techniques that aim to speed up the rendering process. One of the most common simplifications for rendering volume scattering is single scattering [PAS03], which assumes that

light gets scattered only once on its way from the light source to the camera. This assumption, however, is not valid for optically dense media. Usually, light is scattered multiple times on its way to the camera. Rendering multiple scattering produces a more realistic result, especially in optically dense media [PAS03], but takes a long time to compute.

We propose *Multiple Vertex Next Event Estimation* (MVNEE) to improve direct lighting calculation in homogeneous participating media by creating multiple-segment sub paths instead of one-segment connections. It is designed as an extension for existing multiple scattering rendering techniques in a Monte Carlo context. It is especially beneficial in optically dense, highly forward scattering media. Our technique creates entire sub paths connecting a vertex in the participating medium to the light source. These sub paths are created by constructing a seed path and performing a one-step perturbation of the vertex locations. MVNEE is developed for infinite, homogeneous participating media using the Henyey-Greenstein phase function [HG41], but we also discuss extensions to spot light sources and heterogeneous media.

2. Background and Related Work

Radiative transfer Lighting effects caused by interaction with the medium in a participating medium can be described by the Radiative Transfer Equation (RTE) [Cha60]:

$$\begin{aligned} L(x,\omega) = {} & \tau(x,\omega,d)L_e(y,\omega) + \int_0^d \tau(x,\omega,s)L_{e,V}(z,\omega)\,\mathrm{d}s \\ & + \tau(x,\omega,d)\int_\Omega f_r(\omega,y,\omega')L(y,\omega')\cos\theta'\,\mathrm{d}\omega' \\ & + \int_0^d \tau(x,\omega,s)\mu_s(z)\int_\Omega \phi(\omega'\bullet\omega)L(z,\omega')\,\mathrm{d}\omega'\,\mathrm{d}s, \end{aligned} \tag{1}$$

where $z = x+\omega\, s$ and $y = intersectScene(x,\omega)$. The scattering and absorption effects at a position x are given by the scattering coefficient $\mu_s(x)$ and absorption coefficient $\mu_a(x)$, which together form the extinction coefficient $\mu_t(x) = \mu_s(x)+\mu_a(x)$ [Raa08]. In a homogeneous medium, their values are independent of position. The distribution of scattering directions is described by the phase function $\phi(\omega\bullet\omega')$.

Absorption and out-scattering effects lead to an exponential attenuation of light over a specific distance [Jen01]. This effect is described by the transmittance:

$$\tau(x,\omega,s) = e^{-\int_0^s \mu_t(x-t\omega)\,\mathrm{d}s}. \tag{2}$$

In a heterogeneous medium, the transmittance depends on a starting position x and direction ω, as well as the distance s. In homogeneous media, the transmittance simplifies to $\tau(s) = e^{-\mu_t s}$ [Cer05].

Volume rendering Rendering of multiple scattering is computationally expensive, especially in optically dense media. In forward path tracing, paths are constructed incrementally. Every path segment is determined by sampling a distance proportional to transmittance and a direction proportional to the phase function.

Kulla and Fajardo [KF12] introduced equiangular sampling in order to sample distances proportionally to the geometry term instead of the transmittance. This can be beneficial, since the variance of the geometry term is often higher than that of the transmittance.

Joint importance sampling [GKH*13] constructs single and double scattering sub paths, while accounting for the product of phase functions and geometry terms along the sub path. For isotropic scattering, a fully analytic formula for sampling phase functions and geometry terms for double scattering at once is derived using marginalization. Using tabulation, a generalized method is provided that can handle anisotropic scattering as well. Since the transmittance is not importance sampled, other methods like random walk path tracing can outperform this technique in dense media. In contrast to our technique, Joint importance sampling is used in media with low-order scattering. Additionally an analytical sampling scheme is only provided for isotropic phase functions. Our method is best suited for dense, anisotropic media. Joint importance sampling samples distances based on geometry terms, while we use transmittance distance sampling.

Jensen and Christensen [Jen98] build *volume photon maps* in a first rendering pass and use the stored data in a second pass in order to render the final image. An adaptive ray marching algorithm is used, which looks up the radiance at points along the ray with a specially designed radiance estimation, which is suited for anisotropic scattering. Moon et al. [Moo06] use photon maps for multiple scattering renderings in hair. During the photon trace pass, the particle data is stored along particle paths instead of scattering positions. Using a 5D position-direction data layout for the photon map, anisotropic scattering can be simulated. Additionally a radiance caching method for fibers is integrated. Novák et al. [Nov12] present a many-light algorithm based on *virtual ray lights* (VRLs), that are placed along each path segment in the medium. Their method is unbiased and is able to render multiple scattering in heterogeneous media.

A variety of approximations are used to simplify multiple scattering renderings. Path integration for light transport in volumes [PAS03, APRN04, Pre04] approximates multiple scattering based on Feynman's path integral approach to solving quantum mechanics problems. The approximations are based on *most probable paths*, which are similar to our seed paths. Once they are found, the contribution of the surrounding neighborhood is approximated by blurring the contribution of those *most probable paths* into the surrounding neighborhood.

For subsurface scattering and BSSRDF models, diffusion approximation is a widely used technique. It is used in optically dense media, where scattering events happen frequently. Stam [Sta95] solved the diffusion equation for heterogeneous media using a multi-grid method. Jensen et al. [Jen01] exploit that the light distribution in homogeneous, highly scattering media tends to become isotropic. They introduced the dipole diffusion model, which they integrated into a BSSRDF. Christensen et al. [CB15] presented parameterizations of BSSRDF models based on empirical data. Using brute force Monte Carlo reference data, they manually fit the reflectance profile. This approximation makes their method multiple times faster than the dipole diffusion model.

Perturbations in the Monte Carlo Context Perturbing a transport path is most often done in a Markov chain Monte Carlo (MCMC) context [VG97], where a random acceptance enforces detailed balance which in turn ensures the correct probability distribution in the equilibrium state. However, one of the problems of

MCMC is temporal inconsistency [HDF15], which makes it unsuitable for production rendering. The concept of perturbation, however, can be applied to a standard Monte Carlo context. Let $\mathbf{X}$ be the perturbed path and $\mathbf{S}$ the seed path. The probability density of $\mathbf{X}$ in Monte Carlo context is dependent on the sampled seed path $\mathbf{S}$ and the perturbation probability [HDF15]:

$$p(\mathbf{X}) = \int_{\mathbf{S}\in\Omega} p(\mathbf{X}|\mathbf{S}) \cdot p(\mathbf{S})\, \mathrm{d}\mathbf{S} \tag{3}$$

This integral is costly to compute, but using further constraints on the seed path and the perturbation, its computation can be avoided [HDF15]. First, the seed path is only used as the input of a single-step perturbation. Finally, $\mathbf{S}$ can be reconstructed given $\mathbf{X}$. For a perturbed path $\mathbf{X}$, there is exactly one seed path $\mathbf{S}$ that could be used in order to get $\mathbf{X}$ with the help of perturbation. These constraints lead to the simplified PDF $p(\mathbf{X}) = p(\mathbf{X}|\mathbf{S}) \cdot p(\mathbf{S})$.

3. Multiple Vertex Next Event Estimation

In this section, our technique is introduced for unbounded homogeneous media using the Henyey-Greenstein [HG41] phase function. We create sub paths from path vertices on surfaces or within the medium to the light source. Using this multi vertex connection, additional scattering events are sampled. The sub paths are constructed by first creating a seed path and then using a specialized perturbation strategy in order to translate the seed vertices. The perturbation is performed within the Monte Carlo context. Our technique is combined with a forward path tracer via multiple importance sampling. Source code is available at `http://github.com/pasiwe/MVNEE-Paper`.

Seed path A seed path can be constructed in several ways. The number of intermediate seed vertices, as well as their positioning, has to be determined. Further, it is important to make sure the seed path can be reconstructed after perturbation [HDF15]. This is needed to compute the PDF of the technique given a path, as required by multiple importance sampling.

Given a path vertex x_p, and a light vertex x_{ls}, the seed vertices s_i are positioned on the straight line connecting x_p with x_{ls} (see Fig. 2). This is done by sampling distances proportional to the transmittance. Starting at x_p, distances $d_i, i = 1,...,n$ are sampled until the remaining distance to the light source is exceeded.

Perturbation The perturbation strategy has to make sure that the seed path can be reconstructed. For our seed path construction, sampling distance by transmittance, the perturbation may not alter the distance between the projected sub path vertices, otherwise this condition is breached. This means the perturbation may only translate vertices in the plane of the tangent space of the seed vertices, perpendicular to the connection from the start vertex x_p to the light vertex x_{ls}. This is shown in Fig. 2.

Different strategies for perturbation in the tangent plane can be employed. We chose to use an independent, isotropic 2D bell shaped PDF centered around each seed vertex. Sampling such a bell shaped PDF makes sure that paths near the seed path are created.

An isotropic 2D Gaussian PDF as well as a 2D GGX [WMLT07]

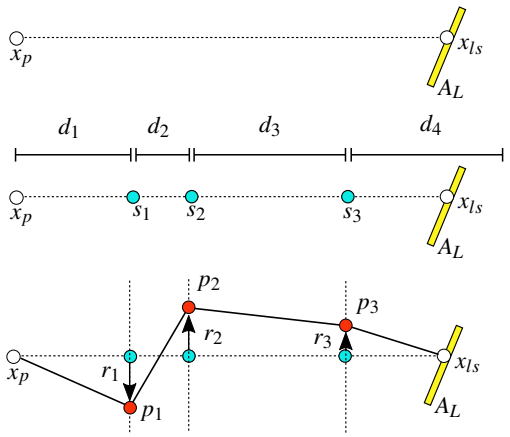

Figure 2: *Creation of the MVNEE Path: At first a vertex x_{ls} on the light source is sampled. Then the seed vertices s_i are determined by sampling distances d_i. During perturbation the vertices p_i are determined by sampling a 2D GGX PDF in the tangent plane through the seed vertex s_i.*

PDF were tested for this purpose. For every PDF, a width parameter has to be set, in the case of the Gaussian PDF, this would be the standard deviation σ. This parameter should be chosen appropriately in order to reduce variance of the perturbation. This width parameter should be adjusted depending on the Henyey-Greenstein mean cosine g, as well as the current seed vertex position. In the next section, our formula for the calculation of the Gauss and GGX width parameters is presented.

Variance reduction In order to calculate the width parameter with a formula dependent on the phase function parameter (the mean cosine), the projection of the Henyey-Greenstein phase function on a 2D plane is approximated by a Gaussian with standard deviation σ_ϕ. For this purpose, we have conducted an experiment:

An image plane is placed in a distance of one mean free path from an initial position x_i as seen in Fig. 3. It has a size of 8×8 mean free paths (MFP) and a pixel resolution of 256×256. The initial direction ω_i is set to be perpendicular to the image plane. For every pixel center x_p on the image plane the phase function value v_p in vertex area measure is calculated and stored at the pixel position. For a pixel center x_p the value v_p is calculated as:

$$v_p = \frac{\phi\left(\omega_i \cdot \frac{x_p - x_i}{\|x_p - x_i\|}\right)}{\|x_p - x_i\|^2}. \tag{4}$$

The resulting distribution of the phase function values is reduced to one dimension by radial averaging. Afterwards a Gaussian PDF fitting is performed.

As a sanity check, the fitting was first performed with $\mu_t = 1.6$ and in a second pass with $\mu_t = 1.0$. The mean free paths are 0.625 meter and 1 meter respectively. The resulting data is listed in Tab. 1. In the last column, the ratio of the fitting values for MFP $= 0.625m$

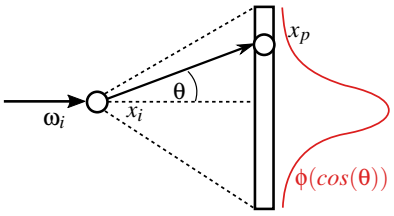

Figure 3: *Henyey-Greenstein phase function approximation using a Gaussian PDF.*

with respect to the fitting results for MFP $= 1m$ are listed. As expected, the ratio is very close to 0.625 in every line.

g	σ_ϕ MFP $= 0.625m$	σ_ϕ MFP $= 1m$	$\frac{\sigma_\phi \text{ for } 0.625m}{\sigma_\phi \text{ for } 1m}$
0.0	0.672021	**1.07526**	0.62498
0.1	0.600581	**0.961113**	0.62488
0.2	0.524889	**0.839835**	0.62499
0.3	0.447689	**0.716209**	0.625
0.4	0.371379	**0.594202**	0.625
0.5	0.297962	**0.476602**	0.62518
0.6	0.228637	**0.365761**	0.625
0.7	0.164514	**0.263209**	0.625
0.8	0.105868	**0.169391**	0.62499
0.9	0.0519185	**0.0830828**	0.6249

Table 1: *Results for fitting a Gaussian PDF to the Henyey-Greenstein values v_p. The fitted standard deviation with respect to mean cosine is provided for a mean free path of 0.625 and 1 meter.*

Gaussian perturbation of multiple scattering vertices To compute the width parameter for multiple scattering vertices along the seed path, we express the product of phase functions as product of Gaussian PDFs. As a result, a final Gaussian standard deviation is calculated for every seed vertex position. It can be directly used for perturbation or multiplied by a constant to emulate the GGX width parameter. The dependency of the width parameter on the Henyey-Greenstein mean cosine g can be satisfied by choosing the σ_ϕ according to g, using linear interpolation between values in Tab. 1.

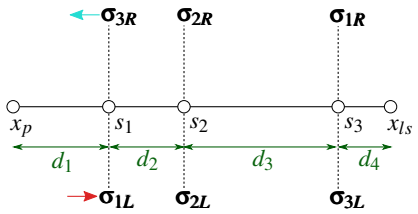

Figure 4: *Input standard deviations σ_{iL} and σ_{iR} for the product of Gaussians at every seed vertex. The standard deviations depend on the mean cosine and the distances d_i to start vertex x_p and end vertex x_{ls}.*

In Fig. 4 an exemplary seed path is depicted. For every seed vertex $s_i, i = 1,2,3$, two Gaussian PDFs with standard deviations σ_{iL} and σ_{iR} are computed. The standard deviation σ_{iL} is calculated using the convolution of Gaussians [Bro03] for all layers up to the current distance:

$$\begin{aligned}\sigma_{iL} &= \sqrt{(\sigma_\phi \cdot d_1)^2 + (\sigma_\phi \cdot d_2)^2 + \ldots + (\sigma_\phi \cdot d_i)^2} \\ &= \sqrt{d_1^2 + d_2^2 + \ldots + d_i^2} \cdot \sigma_\phi.\end{aligned} \tag{5}$$

σ_{iR} is calculated analogously, while starting at the other side of the segment at x_{ls}. For distances $d_i \in \{d_1, \ldots, d_n\}$, σ_{iR} becomes:

$$\sigma_{iR} = \sqrt{d_n^2 + d_{n-1}^2 + \ldots + d_{i+1}^2} \cdot \sigma_\phi. \tag{6}$$

The Gaussian PDF with standard deviation σ_{iL} approximates the distribution of phase function products over i segments with distances d_i, starting from x_p. The same holds true for σ_{iR} starting from vertex x_{ls}. Given both σ_{iL} and σ_{iR} the final standard deviation σ_i for the perturbation of seed vertex s_i is calculated using the product of Gaussians [Bro03]:

$$\sigma_i = \sqrt{\frac{1}{\frac{1}{\sigma_{iL}^2} + \frac{1}{\sigma_{iR}^2}}}. \tag{7}$$

Using Eqs. (5), (6) and (7), the standard deviation σ_i for any seed vertex s_i can be calculated in a deterministic way. The convolution makes sure that the Gaussian PDF gets wider for increasing distances to the end points. The product of Gaussians can be seen as the constraint that makes sure that all paths end in the boundary vertices x_p and x_{ls}.

Handling long tails in the distribution Our numerical experiments show that in fact, a 2D GGX distribution is a better fit for the plane-parameterized phase function than a Gaussian. We thus use the Gaussian approximation only to derive width parameters for every seed vertex. Given a final standard deviation σ_i for a seed vertex s_i, the width parameter for the GGX PDF can be calculated by multiplication by a constant. From our empirical fitting data on various Gaussian and GGX PDFs, it became apparent that the fitted values for GGX α and the standard deviation σ have a linear relation. By averaging the ratio of GGX α and Gauss σ over multiple fitting results, $c_{GGX} = 1.637618734$ was determined as GGX conversion constant. The final GGX width parameter value α_i for seed vertex i is calculated by multiplying it with the calculated standard deviation: $\alpha_i = c_{GGX} \cdot \sigma_i$.

Discussion Rather than constructing a curved seed path, for example with the help of B-Splines [Pie91, BSJS03, JŽ11], this straight line connection has several advantages. First it is easier and faster to create and reconstruct. Secondly it is the shortest connection possible, which is optimal in terms of transmittance. Forward scattering phase functions are satisfied, since multiple scattering events with ideal forward scattering, in this case $\cos\theta = 1$, can outweigh one potentially unfavorable scattering direction at x_p. At last the positions of the seed vertices on the straight line can be easily importance sampled. A downside of this approach is that the light source direction is not considered. We will first use diffuse area lights and discuss spotlights later.

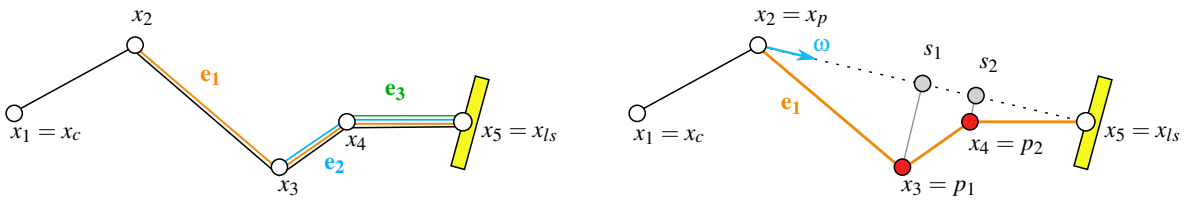

Figure 5: *Visualization of the process of MVNEE PDF calculation for a given path $X = (x_c = x_1, x_2, ..., x_n = x_{ls})$ in the medium, x_c being the start vertex on the image plane, x_{ls} a vertex on the light source. On the left all MVNEE estimators e_i for the path are shown, on the right the seed path reconstruction for estimator e_1 is visualized.*

The perturbation described here ignores any correlation that should exist between consecutive vertices in a path. We found this to only become a problem for very long sub paths, as will be discussed later (Fig. 12).

Computation of MVNEE Estimator PDFs MVNEE can be integrated into path tracing via *multiple importance sampling* (MIS) [VG95] for an unbiased combination. For this, the PDF of the estimator has to be calculated.

An MVNEE sub path can potentially start at every path tracing medium vertex, but the initial vertex x_c and the light vertex x_{ls}. In a path $P = (x_1, ..., x_n = x_{ls})$ containing only vertices in the participating medium with $n-1$ segments, there is one dedicated path tracing estimator and $n-2$ MVNEE estimators $e_i, i = 1, ..., n-2$ starting at vertex $x_2, ..., x_{n-1}$ respectively. The different MVNEE estimators e_i for a given path are shown in Fig. 5 (left).

In order to calculate the PDF of an MVNEE estimator e_i, its seed path has to be reconstructed. This is visualized in Fig. 5 (right) for the MVNEE estimator e_1. The path vertices starting after $x_p = x_2$ have to be treated as the perturbed vertices of the MVNEE path. In a MVNEE path with k segments, $k-1$ seed vertices $s_i, i = 1, ..., k-1$ have to be determined by orthogonal projection onto the dotted line between x_p and x_{ls}.

For certain sub paths created with path tracing, seed path reconstruction is impossible. It is necessary to check that all vertices s_i are positioned between x_p and x_{ls} in consecutive order, i.e. there can be no loops. Additionally no zero-distances due to floating point precision are allowed. Whenever these conditions are breached, the MVNEE estimator PDF is zero.

The PDF of the MVNEE sub path is the product of the PDF for sampling the light vertex $p(x_{ls})$, the PDF for sampling all seed segment lengths and the PDF for perturbation in the tangent plane. For the PDF of the entire estimator, the path tracing PDF up to start vertex x_p has to be multiplied as well. Given all estimator PDFs in vertex area measure, the MIS weight can be calculated.

Performance considerations MVNEE needs more time to compute with increasing number of segments. Naturally the MIS weight computation depends on the number of perturbed vertices to the light source. The expensive step, however, is the sheer amount of estimator PDFs that have to be calculated. In dense media, a path from the camera to the light source can have a high number of segments. After every one of them, a new MVNEE sub path may have been created, spanning the whole path suffix to the light source. Thus, the time needed to evaluate a single path is a lot higher compared to a simple technique like NEE.

In order to reduce the number of created MVNEE paths with high segment count to a reasonable number, we introduce a limit for the *maximum expected segment count* (MESC). Before an MVNEE connection is attempted, the expected segment count is calculated by dividing the current distance to the sampled light vertex by the mean free path. If this count is equal to or below the threshold for the *maximum expected segment count*, the sub path creation is performed, otherwise it is aborted and no further calculation is done. Note that this further reduces the sub domain in path space that can be sampled via MVNEE, but a combined estimator including standard random walk path tracing will still be unbiased.

During the MIS weight calculation, if the expected segment count is equal or below the value of the MESC, the estimator PDF is calculated, otherwise set to zero. Using this a-priori decision, the number of created MVNEE paths, as well as their average segment count can be decreased. Additionally the number of calculated estimator PDFs is significantly reduced, thus shortening the time for rendering. The loss of lighting contributions has to be compensated by path tracing. Note that is is still possible for MVNEE paths to have more segments than the MESC, by sampling multiple distances smaller than the mean free path.

4. Extensions

So far MVNEE has been considered in unbounded, homogeneous participating media without opaque objects and with only one light source. In this section we discuss extensions to surface geometry interaction, multiple light sources, and heterogeneous media.

Surface geometry When using MVNEE in a scene containing surface objects, occlusion checks have to be performed for every MVNEE sub path segment. Once the perturbation is finished, normal culling at the surface point has to be performed as well. If a path is occluded, it has to be rejected. Note that it can happen that the seed path is completely visible, yet the perturbed path intersects scene geometry. On the other hand the seed path can potentially be occluded, while the perturbed path is not. Another important issue is that only the path tracing random walk can create a valid object intersection. This means the number of MVNEE estimators is reduced when a path contains surface intersections: the longest sub path is from the light source to the first surface intersection.

Multiple light sources When using multiple light sources, the light source has to be sampled prior to sampling a vertex on that light source. Usually the light source is chosen uniformly or proportionally to the light source emission intensity. In case of participating media however, the distance to the light source plays an important role. It makes sense to adjust the sampling of light sources to the distance as well. This is why we used an *intensity-distance-based* sampling scheme. Given a current path vertex x, used for direct lighting, and L light sources $l_1,...,l_L$ with center vertices $c_1,...,c_L$ the probability for sampling one of the light sources is calculated using the distance and the light source intensities $I(l_i)$:

$$p(l_j|x) = \frac{I(l_j)/\|c_j - x\|}{\sum_{i=1}^{L} (I(l_i)/\|c_i - x\|)}. \quad (8)$$

The linear falloff is motivated by the distance term in the diffusion approximation [Jen01]; we also tested a quadratic falloff, which did not work equally well.

Heterogeneous volumes Heterogeneous media pose several challenges for MVNEE. Performance analyses become very different, since spatially varying scattering coefficients impose a lot of memory accesses. On top of that there is a mix between vacuum and medium. This makes the presented GGX perturbation strategy less viable, because vertices can be translated into vacuum by perturbation. Additionally, since the perturbation depends on the length of the seed segments, its quality is affected when part of the seed path is inside vacuum, as shown in Fig. 6. Further the MESC can only be approximated by using the extinction coefficient at the start vertex x_p, assuming a homogeneous medium. This can be problematic, especially when starting from a vertex in vacuum.

Keeping these special cases in mind, MVNEE as a technique can be adapted to heterogeneous media with little effort, as long as an interface for transmittance calculation and distance sampling is provided (e.g. by residual ratio tracking [NSJ14] and Woodcock tracking [Raa08]). Whenever the perturbation translates a vertex into vacuum, MVNEE has to be aborted. When reconstructing a seed path, the seed vertices have to be inside the medium, otherwise the seed path could not have been sampled.

The execution time can be further reduced when sampling the seed distances based on the homogeneous transmittance formula. This fast version is realized by sampling the first heterogeneous distance and using the density and extinction coefficient at the first seed vertex to sample all further distances. Again when sampling vertices in the vacuum, the sub path has to be aborted.

Adapting MVNEE to heterogeneous media can be simplified when combined with a homogeneous background medium. This way no problems with vertices in vacuum occur, and fast seed path sampling can be applied as well.

5. Results

In this section, the performance of MVNEE is compared to NEE. Additionally the effect of the MESC, the limit on the segment count of MVNEE paths, is assessed. All techniques were implemented in a custom renderer and ran on an Intel Core i7-4790K CPU with 8 threads. The following techniques are compared: Path tracing combined with MVNEE using MIS with an MESC= n (MVNEE-MIS-n) and Path tracing combined with NEE using MIS (NEE-MIS).

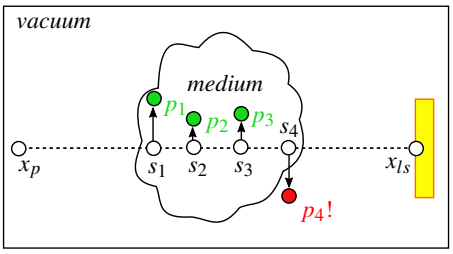

Figure 6: *Perturbation problems in heterogeneous media.*

Correctness In order to show that our technique produces unbiased results, we compare our technique to reference images. The results can be seen in Fig. 7.

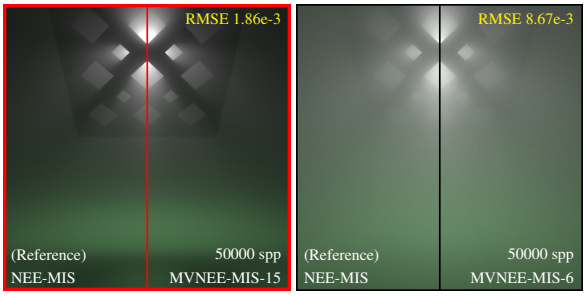

Figure 7: *Comparison of the NEE-MIS reference image to an almost converged MVNEE-MIS result rendered with 50000 spp. Left halves: NEE-MIS reference image (1 Mspp, 16.2h on a 256-thread machine), right: MVNEE-MIS image. The rendering of the MVNEE images took 12.8/7.2 hours on the Intel Core i7 CPU with 8 threads (left/right, corresponding to a dense/thin medium).*

Dense medium Fig. 8 shows equal time comparisons (20min and 60min) on a circular area light with radius 0.2 meters, positioned above the camera facing towards the ground. The distance between camera and light source is around 4.5 meters. A perforated plate mesh is placed between light source and camera. The medium is parameterized with $\mu_s = \mu_t = 2.0$, the Henyey-Greenstein mean cosine is set to $g = 0.9$ for strong forward scattering. For all renderings in this setup, the maximum path segment count is set to 30.

The images are sorted from left to right by the RMSE of the green color channel. MVNEE outperforms NEE-MIS in every version with respect to visible noise and the RMSE, for both equal time comparisons. As expected our techniques generate more samples for lower MESCs. MVNEE-MIS-4 managed to generate almost 3 times more samples than MVNEE-MIS-8 and about twice as many as MVNEE-MIS-6.

However a higher sample count does not automatically result in a lower RMSE. Longer sub paths provide a better quality for direct lighting. MVNEE-MIS-2 generated even more samples than

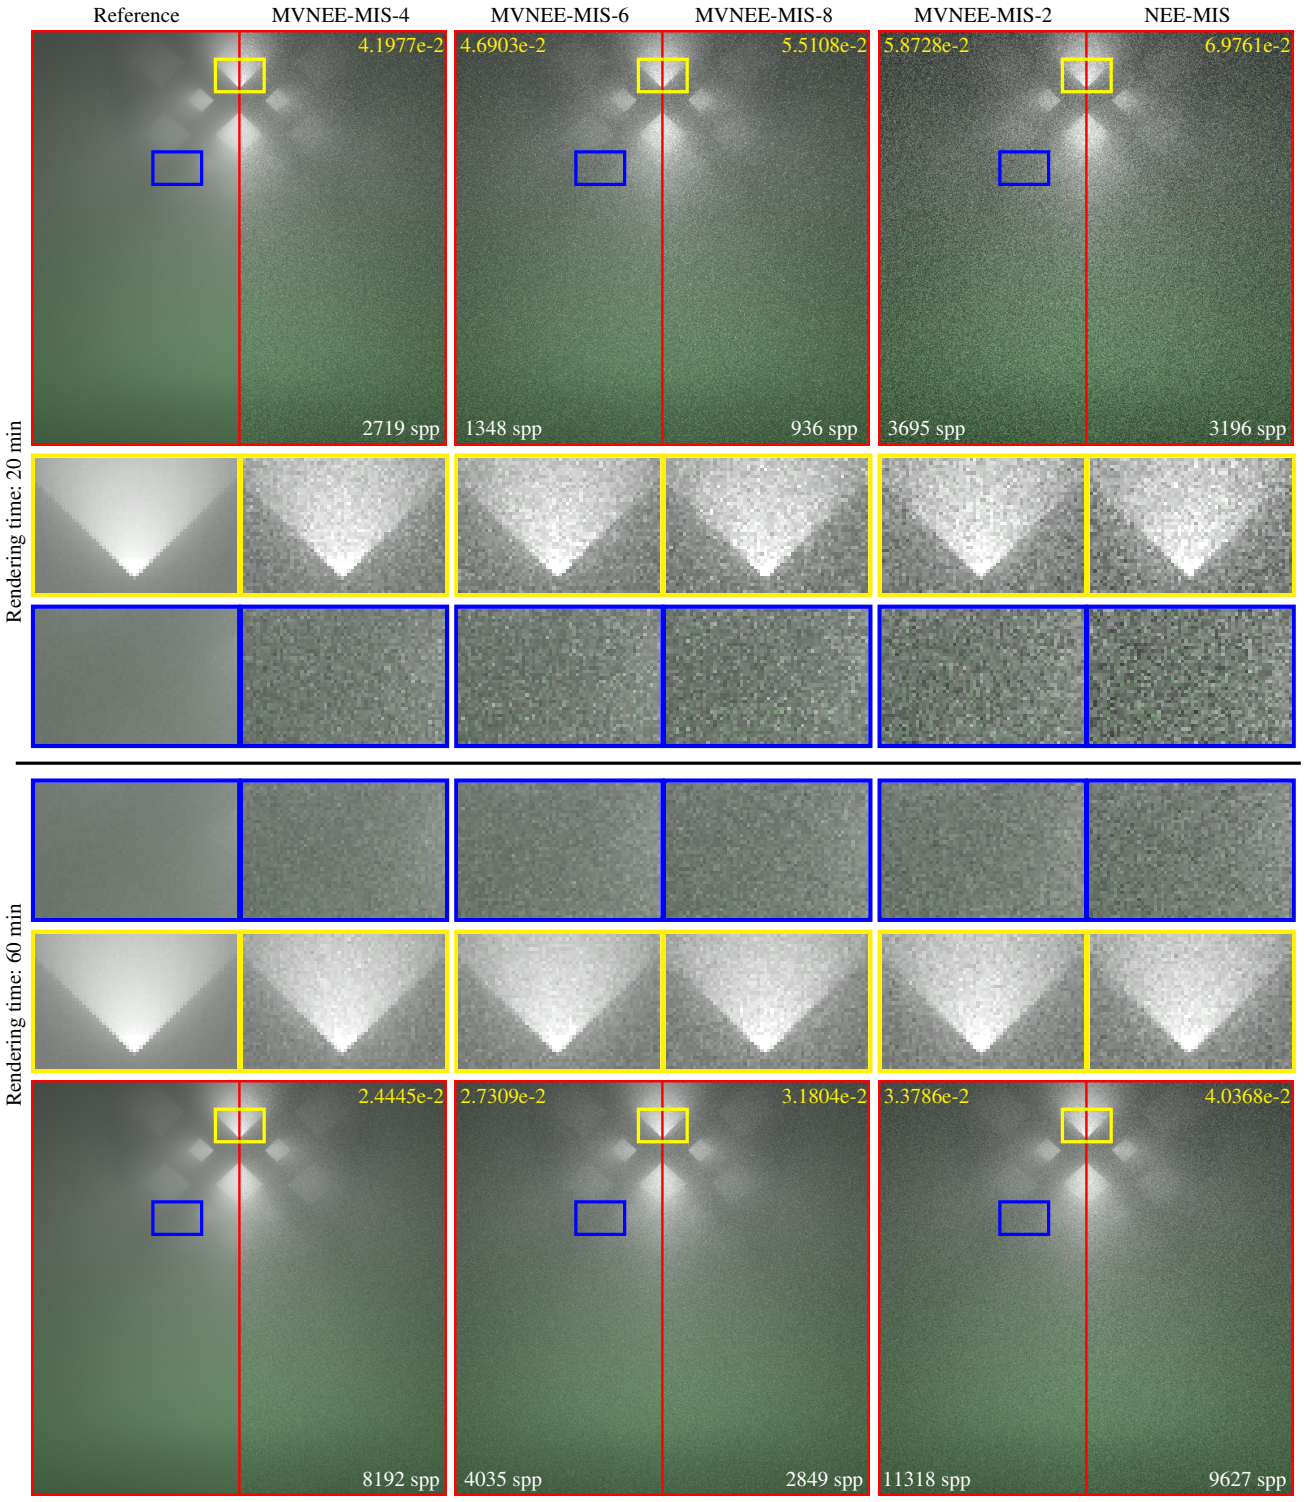

Figure 8: *Equal time comparison in dense medium: For the collection of images above the black line, a rendering time of 20 minutes was used, for the images below the line the rendering time was set to 60 minutes. The Henyey-Greenstein mean cosine is set to 0.9, the extinction coefficient is set to $\mu_t = \mu_s = 2.0$. The output images of the different techniques are sorted by the RMSE in increasing order. The RMSE of the green color channel, which is the dominating channel in this scene, is provided at the top of the output images in yellow. The samples per pixel are shown at the bottom of the images.*

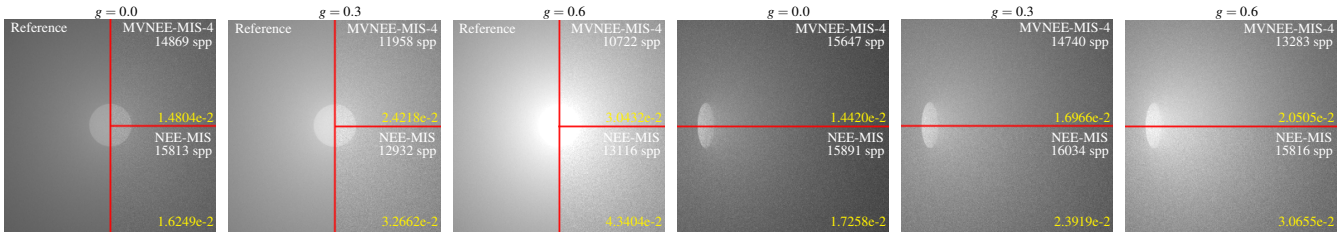

Figure 9: *Equal time comparison (15 min) in an unbounded homogeneous medium using different Henyey-Greenstein mean cosines. The distance from camera to light source is 5m, the light radius is* 0.5*m. Medium parameters:* $\mu_t = \mu_s = 1.7$. *In the left three images, the light source directly faces the camera, in the right three, the light source is perpendicular to the camera direction.*

NEE-MIS, still it is the worst of the MVNEE versions. This shows how crucial multiple importance sampling is. Since a direct lighting approximation is not always possible, the technique relies on random walk path tracing to sample vertices that can be used for direct lighting approximation in the next step. MVNEE without path tracing and multiple importance sampling is not an unbiased technique, because it cannot generate loops.

Different mean cosines Fig. 9 examines the performance of MVNEE for different mean cosines of the Henyey-Greenstein phase function. The medium is parameterized with $\mu_t = \mu_s = 1.7$. Three different mean cosines (0.0, 0.3 and 0.6) are compared. In order to improve the efficiency of MVNEE-MIS, the MESC was reduced to 4 segments (MVNEE-MIS-4).

Even though designed for strongly forward scattering media, MVNEE-MIS-4 outperforms NEE in terms of the RMSE for every presented setting. However the discrepancy between both techniques is increasing with stronger forward scattering. In the isotropic case, our technique is only slightly better than NEE. For $g = 0.6$ our technique is ca. 1.4 times better than NEE after 15 minutes. This holds true for the 90 degrees turned light source, too.

Heterogeneous volumes Fig. 10 shows a 3.5GB voxel dataset rendered as heterogeneous volume. This is a difficult case for MVNEE in two ways. First, the expensive data accesses make it hard to compute many samples per time. The fast approximate transmittance sampling for the seed path helps there (MVNEE-MIS-4-FAST is on par with NEE). Second, the medium is very dense in its core, which means constructing only four segments on the sub path is not enough to reduce variance a lot. It still helps to find indirect light from the cloud shining on the ceiling (blue inset).

Spot light In Fig. 11 a focused spot light shines on a sphere in homogeneous fog. Even though we ignore the emission distribution at the light source, MVNEE performs a lot better than NEE in this example. Some of this gain comes from the a-priori cutoff via MESC. Our technique will only use path tracing to extend the path until the random walk is close enough to the light to make a direct connection viable. This saves on ray tracing costs. In the figure, MVNEE-MIS-4-LIS is a very constrained subset of bidirectional path tracing (BDPT) [LW93], which only casts one segment from the light source. We conducted this experiment to keep the combinations of estimators low for fast PDF evaluation, while still explicitly modelling the light source emission distribution. Even in this case, however, the extra cost of the MIS computation causes a drop from 8061spp to 2565spp in the same time as compared to pure MVNEE-MIS-4.

6. Limitations

Performance MVNEE decreases in performance for long sub paths, due to its sampling process and PDF calculation. The latter is crucial since many estimator PDFs have to be calculated for the MIS weight. We alleviated this issue by a-priori limiting the maximum expected segment count (MESC), which allows us to cut down on the possible number of estimators. In heterogeneous media, however, the expected segment count cannot be calculated reliably without touching a lot of memory, thus impacting performance. The increased sampling overhead compared to standard NEE is especially problematic whenever the created sub path has to be rejected due to occlusion or normal culling.

Bidirectional path tracing So far we ignore the emission distribution function at the light source. This can cause trouble with strongly focused spot lights, and we believe using MVNEE as a generic path vertex connection strategy to extend full bidirectional path tracing [LW93] may lead to better results. Initial experiments with a small subset of connections showed however that careful work will be needed to not impact performance in this case.

Correlation between sub path vertices Our perturbation strategy ignores correlation between sub path vertices, i.e. it perturbs every vertex independently from each other. While this works well for moderate sub path lengths (we found 4–10 is a good range, cf. Fig. 1, right), this becomes a problem for more extreme settings. Fig. 12 shows three plots of paths connecting two vertices in a volume. On the top are the best paths out of one billion path tracing samples (highest measurement contribution in solid angle measure), in the middle are ten random paths from NEE, and at the bottom are ten random MVNEE paths. Fig. 1 uses 10 path vertices, Fig. 12 uses 100 vertices. With that many segments, MVNEE regresses toward the mean and the shape of the path looks more zig-zag than it should, leading to low throughput values. This is another reason why long sub path connections do not work well.

7. Conclusion and Future Work

We presented a new technique called Multiple Vertex Next Event Estimation to calculate lighting from multiple scattering in homogeneous, anisotropic participating media. It creates sub paths to the

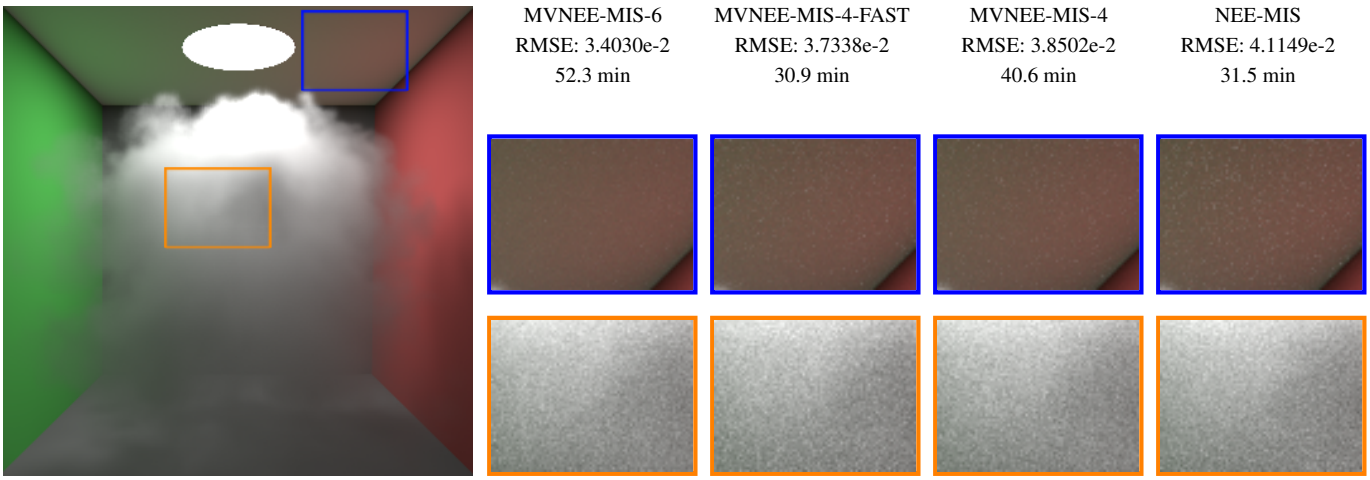

Figure 10: *Equal sample count comparison (1000 spp) with a 3.5GB heterogeneous volume with density multiplied to $\mu_s = 4.0$ and $\mu_a = 1.0$ ($g = 0.9$) in a Cornell box. The maximum path length is 25. The area light source has a radius of 0.8m and is placed in a distance of ca. 8.2 meters from the camera. The reference image (left) was rendered with NEE-MIS using 10^6 spp. MVNEE-MIS-4-FAST uses homogeneous transmittance distance sampling for the seed path. The provided RMSE is from the red color channel, since it had the highest error value.*

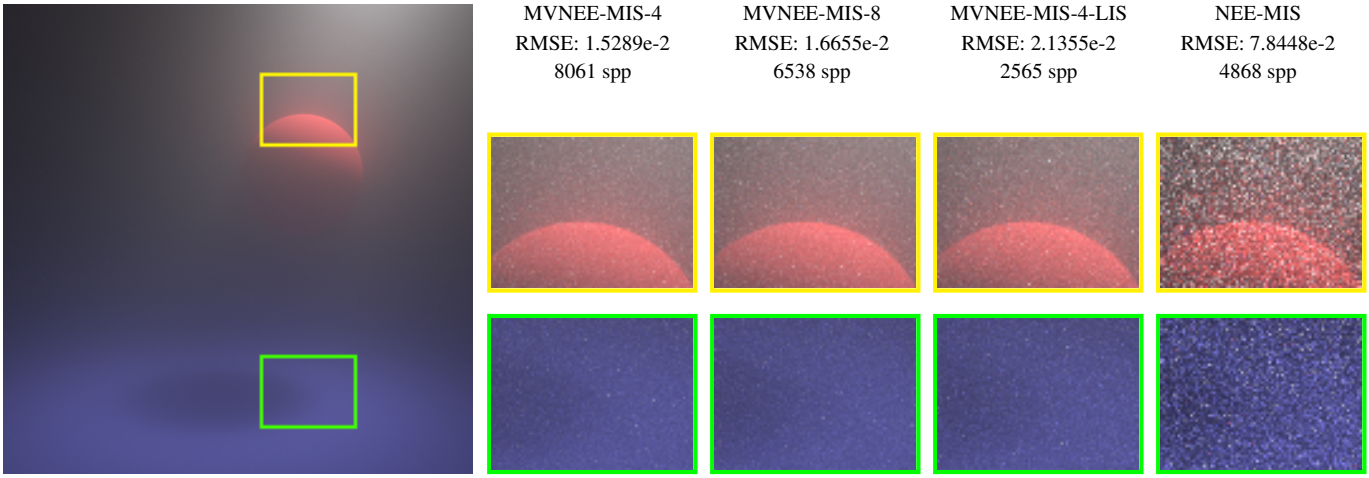

Figure 11: *Equal time comparison (120 min) for rendering a spot light (radius 0.4m, exponent $n = 50$) in a homogeneous volume with $\mu_t = \mu_s = 0.4$ and $g = 0.9$. The maximum path length is 20. The spotlight is placed at a distance of ca. 9.48 meters from the camera. The reference image (left) was rendered with MVNEE-MIS-4 using 10^6 spp. MVNEE-MIS-4-LIS uses the specialized sampling technique where one intermediate vertex is sampled starting from the light source and the MVNEE sub path is connecting to this vertex. Additionally this technique uses NEE whenever the current vertex is closer than 4 mean free paths to the light source. The provided RMSE value is the RMSE of the red color channel, which had the highest error value of the channels.*

light source via perturbation in the Monte Carlo context. The technique can be integrated into any path tracer with multiple importance sampling. It works best in optically dense media with strong forward scattering, using 4–10 sub path segments.

MVNEE suffers from high execution times for long sub paths, which can be decreased by the *maximum expected segment count*. This way it outperforms NEE in a variety of medium parameters.

In the future we would like to be able to render more complex scenes with MVNEE, involving specular boundary conditions, proposed by mirrors or dielectric objects. For this purpose, an attempt of combining MVNEE and MNEE could be made (similar to [KNK*16]). Once MNEE has found a path with multiple specular interactions, the specular vertices could be connected with MVNEE in the participating medium.

In this work MVNEE was used in a Monte Carlo context for forward random walk path tracing. The concept could be applied to BDPT [LW93, Laf96] as well, where MVNEE could be used to connect an eye sub path with a light sub path in a participating medium, though initial experiments show that great care has to be taken not to deteriorate performance.

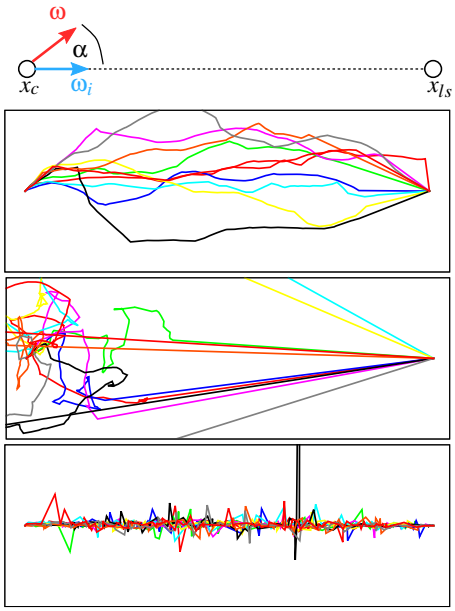

Figure 12: *Fail case for long sub paths. From top to bottom: setup; ten best paths out of 1 billion with highest solid angle measurement contribution; ten random paths from NEE; ten random paths from MVNEE. Since we ignore correlation between vertices, MVNEE starts to clump together near the seed path for long chains.*

References

[APRN04] ASHIKHMIN M., PREMOZE S., RAMAMOORTHI R., NAYAR S. K.: *Blurring of Light due to Multiple Scattering by the Medium: A Path Integral Approach*. Tech. Rep. CUCS-017-04, Columbia University, May 2004. 2

[Bro03] BROMILEY P.: Products and Convolutions of Gaussian Probability Density Functions. *Tina-Vision Memo 3* (2003). 4

[BSJS03] BERGLUND T., STRÖMBERG T., JONSSON H., SÖDERKVIST I.: Epi-Convergence of Minimum Curvature Variation B-Splines. *Dept. Comp. Sci. Elect. Eng., Luleå Univ. Technol., Sweden* (2003), 1402–1536. 4

[CB15] CHRISTENSEN P. H., BURLEY B.: *Approximate Reflectance Profiles for Efficient Subsurface Scattering*. Tech. Rep. 15-04, Pixar, jul 2015. 2

[Cer05] CEREZO, EVA AND PÉREZ, FREDERIC AND PUEYO, XAVIER AND SERON, FRANCISCO J AND SILLION, FRANÇOIS X: A Survey on Participating Media Rendering Techniques. *The Visual Computer 21*, 5 (2005), 303–328. 2

[Cha60] CHANDRASEKHAR, SUBRAHMANYAN: Radiative transfer, 1960. 2

[GKH*13] GEORGIEV I., KŘIVÁNEK J., HACHISUKA T., NOWROUZEZAHRAI D., JAROSZ W.: Joint Importance Sampling of Low-Order Volumetric Scattering. *ACM Transactions on Graphics (TOG) 32*, 6 (2013), 164. 2

[HDF15] HANIKA J., DROSKE M., FASCIONE L.: Manifold Next Event Estimation. *Computer Graphics Forum (Proceedings of Eurographics Symposium on Rendering) 34*, 4 (June 2015), 87–97. 3

[HG41] HENYEY L. G., GREENSTEIN J. L.: Diffuse Radiation in the Galaxy. *The Astrophysical Journal 93* (1941), 70–83. 2, 3

[Jen98] JENSEN, HENRIK WANN AND CHRISTENSEN, PER H.: Efficient Simulation of Light Transport in Scenes with Participating Media Using Photon Maps. In *Proc. SIGGRAPH* (1998), pp. 311–320. 2

[Jen01] JENSEN, HENRIK WANN AND MARSCHNER, STEPHEN R. AND LEVOY, MARC AND HANRAHAN, PAT: A Practical Model for Subsurface Light Transport. In *Proc. SIGGRAPH* (2001), pp. 511–518. 1, 2, 6

[JŽ11] JAKLIČ G., ŽAGAR E.: Curvature Variation Minimizing Cubic Hermite Interpolants. *Applied Mathematics and Computation 218*, 7 (2011), 3918–3924. 4

[KF12] KULLA C., FAJARDO M.: Importance Sampling Techniques for Path Tracing in Participating Media. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering) 31*, 4 (2012), 1519–1528. 2

[KNK*16] KOERNER D., NOVÁK J., KUTZ P., HABEL R., JAROSZ W.: Subdivision Next-Event Estimation for Path-Traced Subsurface Scattering . In *Eurographics Symposium on Rendering - Experimental Ideas & Implementations* (June 2016), The Eurographics Association. 9

[Laf96] LAFORTUNE, ERIC AND WILLEMS, YVES: Rendering Participating Media with Bidirectional Path Tracing. In *Rendering Techniques '96* (1996), Springer, pp. 91–100. 9

[LW93] LAFORTUNE E. P., WILLEMS Y. D.: Bi-Directional Path Tracing. In *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (CompuGraphics)* (1993), vol. 93, pp. 145–153. 8, 9

[Moo06] MOON, JONATHAN T. AND MARSCHNER, STEPHEN R.: Simulating Multiple Scattering in Hair Using a Photon Mapping Approach. *ACM Trans. on Graphics (Proc. SIGGRAPH)* (2006), 1067–1074. 2

[Nov12] NOVÁK, JAN AND NOWROUZEZAHRAI, DEREK AND DACHSBACHER, CARSTEN AND JAROSZ, WOJCIECH: Virtual Ray Lights for Rendering Scenes with Participating Media. *ACM Transactions on Graphics (TOG) 31*, 4 (2012), 60. 2

[NSJ14] NOVÁK J., SELLE A., JAROSZ W.: Residual Ratio Tracking for Estimating Attenuation in Participating Media. *ACM Trans. on Graphics (Proc. SIGGRAPH Asia) 33*, 6 (Nov. 2014). 6

[PAS03] PREMOŽE S., ASHIKHMIN M., SHIRLEY P.: Path Integration for Light Transport in Volumes. In *Proc. Eurographics Workshop on Rendering* (2003), pp. 25–27. 1, 2

[Pie91] PIEGL, LES: On NURBS: A Survey. *IEEE Computer Graphics and Applications*, 1 (1991), 55–71. 4

[Pre04] PREMOŽE, SIMON AND ASHIKHMIN, MICHAEL AND TESSENDORF, JERRY AND RAMAMOORTHI, RAVI AND NAYAR, SHREE: Practical Rendering of Multiple Scattering Effects in Participating Media. In *Proceedings of the Fifteenth Eurographics conference on Rendering Techniques* (2004), pp. 363–374. 2

[Raa08] RAAB, MATTHIAS AND SEIBERT, DANIEL AND KELLER, ALEXANDER: Unbiased Global Illumination with Participating Media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer, 2008, pp. 591–605. 2, 6

[Sta95] STAM J.: Multiple Scattering as a Diffusion Process. In *Rendering Techniques '95*. Springer, 1995, pp. 41–50. 2

[VG95] VEACH E., GUIBAS L. J.: Optimally Combining Sampling Techniques for Monte Carlo Rendering. *Proc. SIGGRAPH* (1995), 419–428. 5

[VG97] VEACH E., GUIBAS L. J.: Metropolis Light Transport. *Proc. SIGGRAPH* (1997), 65–76. 2

[WMLT07] WALTER B., MARSCHNER S. R., LI H., TORRANCE K. E.: Microfacet Models for Refraction through Rough Surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques* (2007), Eurographics Association, pp. 195–206. 3