EUROGRAPHICS 2018 / D. Gutierrez and A. Sheffer
(Guest Editors)

Volume 37 (2018), Number 2

Multiple Scattering in Inhomogeneous Participating Media Using Rao-Blackwellization and Control Variates

László Szirmay-Kalos[1], Milán Magdics[1], and Mateu Sbert[2]

[1]: Budapest University of Technology and Economics, Deptartment of Control Engineering and Information Technology (Hungary)
[2]: Tianjin University, School of Computer Science and Technology (China); University of Girona (Spain)

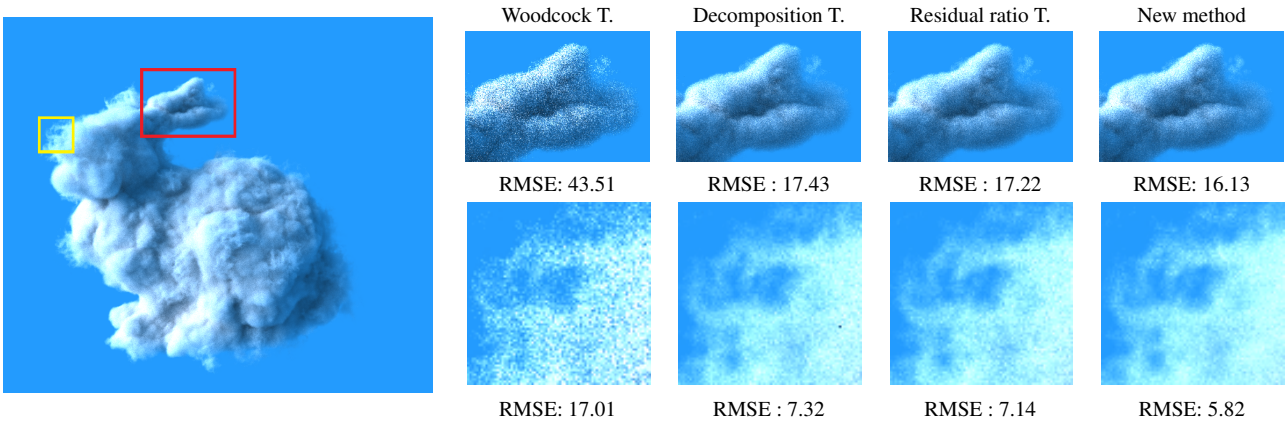

Figure 1: *Bunny shaped dense medium rendered with 50k fetches per pixel.*

Abstract

Rendering inhomogeneous participating media requires a lot of volume samples since the extinction coefficient needs to be integrated along light paths. Ray marching makes small steps, which is time consuming and leads to biased algorithms. Woodcock-like approaches use analytic sampling and a random rejection scheme guaranteeing that the expectations will be the same as in the original model. These models and the application of control variates for the extinction have been successful to compute transmittance and single scattering but were not fully exploited in multiple scattering simulation. Our paper attacks the multiple scattering problem in heterogeneous media and modifies the light–medium interaction model to allow the use of simple analytic formulae while preserving the correct expected values. The model transformation reduces the variance of the estimates with the help of Rao-Blackwellization and control variates applied both for the extinction coefficient and the incident radiance. Based on the transformed model, efficient Monte Carlo rendering algorithms are obtained.

1. Introduction

This paper focuses on the multiple scattering rendering of inhomogeneous participating media [CPP*05, Fat09, KGH*14]. Consistent numerical approaches are usually based on Monte Carlo quadrature and trace photons or importons randomly [JC98, JNT*11, NNDJ12, KF12, BJ17]. Particle tracing has high computational cost in heterogeneous participating media since we cannot efficiently compute the transmittance and analytically sample the free flight. To attack these problems, we modify the underlying model of the participating media to allow the partial analytic solution while preserving the expectation of the Monte Carlo simulation. In its core, we apply two variance reduction techniques, Rao-Blackwellization and control variates.

Rao-Blackwellization means that in a random estimate the variance is usually decreased if some internal random variable is replaced by its mean. For example, in direct physical simulation a photon survives an interaction with the probability of the scattering albedo when its energy is unchanged. The variance of the estimator can be reduced if we never terminate a photon, but its energy is scaled by the survival probability [AK90].

The method of *control variates* separates a control variate $f_{\text{main}}(x)$ from integrand $f(x)$, computes the integral of the control

variate analytically, and adds its result to the Monte Carlo estimate of the integral of difference $f(x) - f_{\text{main}}(x)$. If control variate $f_{\text{main}}(x)$ can capture the variation of the original integrand, the difference will be small, resulting in low variance estimates [LW95, RJN16].

The objective of this paper is to put the modification of the laws of light–material interaction in a unified framework, allowing to obtain low variance estimates for the scattered radiance by simultaneous application of Rao-Blackwellization and control variates. Our main contributions are:

- A transformed particle model for heterogeneous participating media where free flight sampling is analytic and the estimators are unbiased.
- The application of the control variates both for the extinction and for the scattered radiance to reduce the variance.
- The combination of control variates with medium decomposition to reduce the cost of sampling.
- Moderating the exponential explosion of the variance caused by the negative difference extinction in decomposed volumes.
- Natural extension to efficient multi-spectral tracking.
- The proof of unbiasedness and variance analysis.

In Section 2 we review the related previous work. Sections 3 and 4 discuss our new model and rendering algorithm. Finally, Section 5 presents rendering results and we compare the new method to Woodcock tracking [WMHL65], generalized residual ratio tracking [NSJ14, SKGM*17] and weighted decomposition tracking [KHLN17].

2. Previous work

A Monte Carlo solution of the light transport in participating media requires the sampling of the free flight of light particles, the computation of the transmittance between two points, and the simulation of the interaction of the material and light particles. Transmittance calculation needs the integral of the medium extinction, and free flight sampling requires the inverse of this integral. Analytic integration of the extinction and its inversion are possible only in simple cases, like the homogeneous medium. In heterogeneous media numerical methods are needed, including ray marching that approximates the integral by a finite sum or *Woodcock tracking* [WMHL65] that samples the free flight with multiple random steps declaring visited points either real or fictitious. Random steps are sampled using a constant majorant extinction coefficient. Efficient Woodcock tracking needs tight majorants, for which Yue et al. [YIC*10] proposed the application of a kd-tree and Szirmay-Kalos et al. [SKTM11] introduced the concept of virtual particles to deal with arbitrary majorants. Galtier et al. [GBC*13] investigated the integrals showing up in the solution and concluded that with taking the absolute value of the difference, they make sense even if the sampling density is not a majorant. Although they mentioned that using non-majorant sampling increases the variance, no formal variance analysis was given, neither did they propose solution for this problem. In Section 3 and in the Appendix, we provide variance analysis and show that the variance grows exponentially with the length of the path where the sampling extinction is not a majorant, thus such cases must be minimized.

It was shown [SKTM11] that Woodcock tracking and control variates can also be exploited for transmittance calculation, which requires a minorant of the extinction. The ratio tracking method [NSJ14] modified the weights of light particles instead of randomly terminating them, and the *residual ratio tracking* method combined the weighting scheme with control variates. This concept has been further generalized giving physical interpretation of non-majorant sampling and even for negative extinction [SKGM*17] showing up when the control variate is not a minorant. In these methods control variates were used to compute the transmittance between two given points and not during free flight.

If the evaluation of the extinction coefficient is expensive, the cost of deciding whether a collision is real or fictitious is significant. *Decomposition tracking* [KHLN17] reduces this burden by considering the original volume as a mixture of a homogeneous medium and a residual medium. If the homogeneous medium is responsible for most of the interactions, then the majority of the real extinction evaluations can be saved. We should emphasize that although Kutz et al. [KHLN17] call the two media resulting from the decomposition as control and residual media, they do not apply the method of control variates, and their goal is not variance reduction, but to save the computation of the expensive extinction coefficient. As we show it in this paper, the decomposition increases the variance and if the residual volume has negative extinctions, the variance grows exponentially. Thus, decomposition tracking is efficient only if a constant positive minorant can be found.

The method of this paper exploits the data generated during a pre-processing step similarly to [SKTM11, KHLN17], works with non-bounding majorant in sampling and approximate minorant in decomposition as [SKGM*17, KHLN17], and naturally extends to spectral tracking. According to our knowledge, this is the first paper that exploits the statistical method of control variates for the extinction in multiple scattering rendering and non-trivially combines the control variates of the extinction and the scattered radiance. With these control variates, we can attack the problem of the exponentially growing variance caused by the non-majorant sampling extinction values or the negative extinction introduced by decomposition.

3. Model transformation

The particular building blocks to obtain low variance and low cost estimators are adding virtual particles, decomposing the volume to a mixture of particles of different properties, Rao-Blackwellization, and control variates. From these, control variates have been discussed only in the integral formulation framework, and used only for transmittance and single scattering. Decomposition has been proposed also only with the integral formulation. Our goal is to combine, generalize, and improve these techniques in the particle model to allow the exploitation of control variates in multiple scattering simulation and to optimize the model based on rigorous variance analysis. To incorporate control variates of the extinction, the radiance modification cannot be restricted to discrete light–material interaction points, but the weight should change even between such interactions. To handle decomposition, we should allow a variety of particles types.

We talk about light particles and their associated *weight* W that

can mean photon energy or importance [Chr03]. Our practical implementation is based on path tracing, so it traces importons. During Monte Carlo simulation the weight of the light particle is modified randomly and the contribution is the product of its weight and the incident radiance, which are sampled independently. Thus, the statistical properties of the estimator can be described by the properties, i.e. the mean and the variance of the weight.

We first analyze the analog simulation following the physical laws. Then the medium and the simulated laws are changed in parallel with the aim of maintaining the same expected values, but with a simpler computation and lower variance.

3.1. Original medium

The behavior of a light particle of initial weight W_0 and direction $\vec{\omega}$ in the real medium can be described by the following postulates (Figure 2):

1. A light particle collides with a material particle in $[s, s+\mathrm{d}s)$ with probability $\sigma_t(s)\mathrm{d}s$ provided that the light particle has reached distance s, where $\sigma_t(s)$ is the *extinction coefficient* that is proportional to the density of material particles.
2. Direction and weight do not change while the light particle travels in free space, i.e. $W(s) = W_0$.
3. Upon collision at distance s, the particle survives with the probability of *scattering albedo* $a(s)$, its direction $\vec{\omega}$ is modified randomly according to *phase function* $\rho(\vec{\omega}\cdot\vec{\omega}')$, and the weight W^{scat} of the scattered particle is equal to the weight of the incident particle.

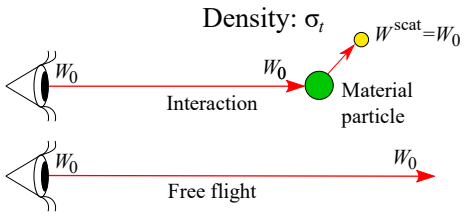

Figure 2: *The original medium. The weight of the particle does not change, but particles may be terminated.*

This physics-driven approach can be analysed by examining expectation $\mathbf{E}[\cdot]$ and variance $\mathbf{V}[\cdot]$ of weights $W(s)$ and W^{scat}. As the indicator whether or not a particle is scattered in $[0,S]$ follows a Bernoulli distribution, we obtain:

$$\mathbf{E}[W(s)] = W_0 T_{\sigma_t}(0,S), \tag{1}$$

$$\mathbf{V}[W(s)] = W_0^2\left(T_{\sigma_t}(0,S) - T_{\sigma_t}^2(0,S)\right), \tag{2}$$

where

$$T_{\sigma_t}(s_1,s_2) = \exp\left(-\int_{s_1}^{s_2}\sigma_t(\tau)\mathrm{d}\tau\right) \tag{3}$$

is the *transmittance*, i.e. the probability of flying through $[s_1,s_2]$ without collision.

The probability density (pdf) of the *free flight* to distance s and scattering at s is

$$\mathrm{pdf}(s) = T_{\sigma_t}(0,s)\sigma_t(s). \tag{4}$$

Let us consider the *scattered weight* W^{scat} that represents the random weight when scattering occurs in $[s, s+\mathrm{d}s)$. Using the probability density of collision from Eq. 4 and the scattering albedo as the probability of survival, we obtain:

$$W^{\mathrm{scat}}(s) = \begin{cases} W_0 & \text{with prob. } a(s)T_{\sigma_t}(0,s)\sigma_t(s)\mathrm{d}s, \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

The expectation and the variance of the scattered weight are

$$\mathbf{E}[W^{\mathrm{scat}}(s)] = W_0 a(s) T_{\sigma_t}(0,s)\sigma_t(s)\mathrm{d}s, \tag{6}$$

$$\mathbf{V}[W^{\mathrm{scat}}(s)] = W_0^2 a(s) T_{\sigma_t}(0,s)\sigma_t(s)\mathrm{d}s + o(\mathrm{d}s) \tag{7}$$

where the little-o notation $o(\mathrm{d}s)$ represents the terms for which

$$\lim_{\mathrm{d}s\to 0}\frac{o(\mathrm{d}s)}{\mathrm{d}s} = 0.$$

The total expected *contribution* of a ray leaving the volume at distance S is:

$$\mathbf{E}[\hat{\Phi}] = W_0 L(S,\vec{\omega}) T_{\sigma_t}(0,S) + \int_0^S W_0 L^{\mathrm{scat}}(s,\vec{\omega}) T_{\sigma_t}(0,s)\sigma_t(s)\mathrm{d}s \tag{8}$$

where $L(s,\vec{\omega})$ is the incident radiance, and we used the following shorthand notation for the expectation of the *scattered radiance* at a point of distance s:

$$L^{\mathrm{scat}}(s,\vec{\omega}) = a(s)\int_\Omega L(s,\vec{\omega}')\rho(\vec{\omega}\cdot\vec{\omega}')\mathrm{d}\omega'.$$

The direct simulation of this physically based model is usually not feasible since it is impossible to analytically compute $T_{\sigma_t}(0,s)$ and to sample the distance with $T_{\sigma_t}(0,s)\sigma_t(s)$. Numerical quadrature methods like *ray marching* are expensive and would make the algorithm biased as it was pointed out in [RSK08], which can be reduced by more samples or higher order integration schemes [Mun14, JLSJ11]. On the other hand, when the particle flies over a region, we ignore what the scattering albedo and the incident radiance are in this region. Should we have at least approximate information about these, it could be built into the estimator reducing its variance.

3.2. Transformed medium

To incorporate control variates of the extinction, the radiance modification cannot be restricted to discrete light–material interaction points, but the weight should change even between such interactions. To reduce the sampling cost, we allow volume decomposition [KHLN17], i.e. assume the medium to be a mixture of different materials where material m is described by real extinction $\sigma_{t,m}$ and by two proxy densities, *sampling extinction* μ_m that is good for free flight sampling, and *main or control extinction* ν_m that allows the analytic evaluation of a reasonable approximation of the transmittance. For the sake of simplicity, the phase functions are supposed to be the same, but the generalization to different phase functions in different materials would be straightforward. The real extinction σ_t of the combined material is the sum of extinctions $\sigma_t = \sum_m \sigma_{t,m}$. Similarly, we introduce the notations of the total sampling density $\mu = \sum_m \mu_m$ and total control variate $\nu = \sum_m \nu_m$, as the sum of the respective parameters of the component materials.

Sampling extinctions μ_m and control extinctions ν_m are free

parameters of our model, which is constructed to make the estimated contribution unbiased for arbitrary control extinction and non-negative sampling extinction. The objective of finding the free parameters is to minimize the variance and the cost of the estimators requiring only approximate or sampled information about the medium. During the development of the practical rendering algorithm, these proxy densities are piece-wise constant functions.

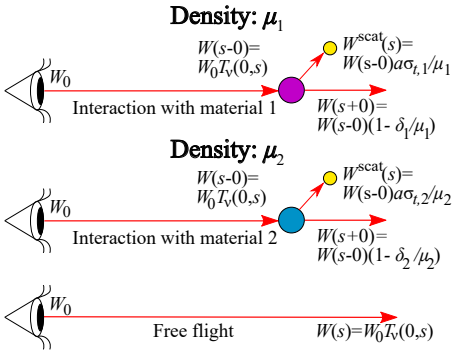

Figure 3: *The transformed model. The weight is continuously decreasing during free flight. When scattering occurs, the particle splits into a transmitted and a scattered particle.*

We express the real extinction $\sigma_{t,m}(s)$ as a sum of control extinction $\nu_m(s)$ and a *difference extinction* in material m:

$$\delta_m(s) = \sigma_{t,m}(s) - \nu_m(s). \tag{9}$$

The postulates of the light particle behavior in the transformed medium are stated to guarantee expected values to remain correct (Figure 3):

1. A light particle collides with a particle of type m in $[s, s+\mathrm{d}s)$ with probability $\mu_m(s)\mathrm{d}s$.
2. While the light particle travels in free space, its direction does not change but its weight $W(s)$ decays in any interval $[s_1, s_2]$ as
$$W(s_2) = W(s_1)\exp\left(-\int_{s_1}^{s_2}\nu(\tau)\mathrm{d}\tau\right) = W(s_1)T_\nu(s_1,s_2) \tag{10}$$
where ν is the sum of the control variates of materials composing the medium.
3. Upon collision at distance s with a material particle of type m, the light particle is broken into a transmitted light particle and a scattered light particle. The transmitted light particle has the same direction as the incident particle. The scattered light particle modifies its direction randomly according to the phase function. The weights of the transmitted and scattered particles are:
$$W(s+0) = W(s-0)\left(1-\frac{\delta_m(s)}{\mu_m(s)}\right), \tag{11}$$
$$W^{\mathrm{scat}}(s) = W(s-0)a(s)\frac{\sigma_{t,m}(s)}{\mu_m(s)} \tag{12}$$
where we used the notation $s-0$ to refer to the weight before scattering and $s+0$ to just after scattering, to resolve the ambiguity caused by the not continuous function $W(s)$.

We show in Appendix 1 that the expectation and the variance of transmitted weight W are

$$\mathbf{E}[W(s)] = W_0 T_{\sigma_t}(0,s), \tag{13}$$
$$\mathbf{V}[W(s)] = W_0^2\left(T_{2\sigma_t-w}(0,s) - T_{2\sigma_t}(0,s)\right) \tag{14}$$
$$= W_0^2\exp\left(-\int_0^s 2\sigma_t(\tau)\mathrm{d}\tau\right)\left(\exp\left(\int_0^s w(\tau)\mathrm{d}\tau\right)-1\right)$$

where $w = \sum_m \delta_m^2/\mu_m$ is the only parameter that depends on the sampling strategy and is called the *variance introduction density* [SKGM*17]. The expectation of the transmitted weight is the same as in Eq. 1, thus it is an unbiased estimator.

The random scattered weight is:

$$W^{\mathrm{scat}}(s) = \begin{cases} W(s-0)a(s)\frac{\sigma_{t,1}(s)}{\mu_1(s)} & \text{with prob. } \mu_1(s)\mathrm{d}s, \\ \dots & \\ W(s-0)a(s)\frac{\sigma_{t,m}(s)}{\mu_m(s)} & \text{with prob. } \mu_m(s)\mathrm{d}s, \\ 0 & \text{otherwise.} \end{cases} \tag{15}$$

The expectation and the variance are (proof in Appendix 2):

$$\mathbf{E}[W^{\mathrm{scat}}(s)] = W_0 a(s) T_{\sigma_t}(0,s)\sigma_t(s)\mathrm{d}s, \tag{16}$$
$$\mathbf{V}[W^{\mathrm{scat}}(s)] = W_0^2 T_{2\sigma_t-w}(0,s)a^2(s)\sum_m \frac{\sigma_{t,m}^2(s)}{\mu_m(s)}\mathrm{d}s + o(\mathrm{d}s).$$

As the expectation of the scattered weight in the transformed model (Eq. 16) is the same as in the original model (Eq. 6), the scattered contribution of the transformed model is also unbiased.

3.2.1. Discussion of the variance

The variance of the transmitted weight W is determined by function w, which should be kept as small as possible. As the square function is convex, Jensen inequality gives us the following bound:

$$w = \sum_m \frac{\delta_m^2}{\mu_m} = \mu\sum_m \frac{\mu_m}{\mu}\left(\frac{\delta_m}{\mu_m}\right)^2 \geq \mu\left(\sum_m \frac{\mu_m}{\mu}\frac{\delta_m}{\mu_m}\right)^2 = \frac{(\sigma_t-\nu)^2}{\mu}, \tag{17}$$

which means that the result is more accurate if the medium is not decomposed.

According to Eq. 14 the minimization of the variance is equivalent to the minimization of the integral of w. If sampling densities μ_m are constant, this means that control variates $\nu_m(s)$ must be the L_2 optimal approximations of real extinction $\sigma_{t,m}(s)$ since

$$\int_0^s w(\tau)\mathrm{d}\tau = \int_0^s \sum_m \frac{\delta_m^2(\tau)}{\mu_m(\tau)}\mathrm{d}\tau = \sum_m \frac{1}{\mu_m}\int_0^s (\sigma_{t,m}(\tau)-\nu_m(\tau))^2\mathrm{d}\tau.$$

Variance $\mathbf{V}[W(s)]$ is surely smaller than that of the original model if the following condition holds:

$$T_{2\sigma_t-w} - T_{2\sigma_t} < T_{\sigma_t} - T_{2\sigma_t} \;\Rightarrow\; 2\sigma_t - w > \sigma_t \;\Rightarrow\; \sigma_t > w = \sum_m \frac{\delta_m^2}{\mu_m}, \tag{18}$$

which is definitely the case, for example, when

$$\sum_m |\delta_m| = \sum_m |\sigma_{t,m} - \nu_m| < \sigma_t \quad \text{and} \quad |\delta_m| < \mu_m, \tag{19}$$

i.e. the control variate does not make the situation worse than not using it at all, and in each material the sampling density is a majorant of the difference extinction.

To examine the price of the variance reduction, let us associate the cost with the expected number of sampled interaction points in a ray, which is $\int_0^s \mu(\tau)d\tau$ for the transformed model and $1-T_{\sigma_t}(0,s)$ for the original model. Using the $\exp(-x) \geq 1-x$ inequality for non-negative $x = \int_0^s \sigma_t(\tau)d\tau$, we can see that the transformed model has higher cost if μ is a majorant of σ_t, and the difference is significant, when the transparency of the whole volume is negligible. Thus, we should also introduce a re-randomization that reduces the cost, which is discussed in Section 3.4.

Let us now examine the scattered weight W^{scat}. Comparing the variance of Eq. 16 to Eq. 7 describing the original model, we can conclude that the transformed model has lower variance if

$$T_{\sigma_t} a\sigma_t > T_{2\sigma_t - w} a^2 \sum_m \frac{\sigma_{t,m}^2}{\mu_m} \implies \sigma_t \geq a T_{\sigma_t - w} \sum_m \frac{\sigma_{t,m}^2}{\mu_m}. \quad (20)$$

The first factor on the right hand side is the scattering albedo, which is always less than 1. Second factor $T_{\sigma_t - w}$ is surely less than 1 if $\sigma_t > w$, which is the same condition that makes the transmission have lower variance (Eq. 18). The third factor is less than σ_t if μ_m is a majorant of $\sigma_{t,m}$. Considering all of these, a sufficient condition for the variance reduction is

$$\sum_m |\delta_m| < \sigma_t \quad \text{and} \quad |\sigma_{t,m}| < \mu_m. \quad (21)$$

Decomposition not only increases variance introduction density w but also makes the third factor of the variance of the scattered weight worse since

$$\sum_m \frac{\sigma_{t,m}^2}{\mu_m} = \mu \sum_m \frac{\mu_m}{\mu} \left(\frac{\sigma_{t,m}}{\mu_m}\right)^2 \geq \mu \left(\sum_m \frac{\mu_m}{\mu} \frac{\sigma_{t,m}}{\mu_m}\right)^2 = \frac{\sigma_t^2}{\mu}. \quad (22)$$

3.3. Control variate for the in-scattered radiance

So far, we have applied the control variate technique for the extinction coefficient and provided formulae to obtain weight W^{scat} that needs to be multiplied by the estimate of the incident radiance. The incident radiance estimate can be obtained by generating a random direction with the pdf of the phase function and by executing the same process for the child ray.

Let us assume that we have a control variate $L_{\text{main}}^{\text{scat}}(s)$ for the scattered radiance for which the line integral problem can be analytically solved in the transformed model. For the sake of simplicity, we use a piece-wise constant control variate along the ray that is direction independent:

$$L_{\text{main}}^{\text{scat}}(s) = L_i^{\text{scat}} \quad \text{if } s_i \leq s < s_{i+1}, \quad i = 0, 1, \ldots, M.$$

We mention that directional dependent approximations would improve the control variate but the computational cost would also be higher [LW95, PWP08].

The estimate of the radiance gathered along a ray is then a sum of two estimates. The first one is the recursive path tracing estimate that is reduced by its control variate $L_{\text{main}}^{\text{scat}}(s)/a(s)$ before adding its contribution to the considered ray. The second one is the line integral taking the pre-defined control variate $L_{\text{main}}^{\text{scat}}(s)$ as the scattered radiance. Now, we examine this calculation and consider the integral of Eq. 8 when $L^{\text{scat}}(s,\vec{\omega})$ is replaced by a piece-wise constant control variate $L_{\text{main}}^{\text{scat}}(s)$:

$$\begin{aligned} \Phi_{\text{main}}^{\text{scat}} &= W_0 \int_0^S L_{\text{main}}^{\text{scat}}(s) T_{\sigma_t}(0,s) \sigma_t(s) ds \\ &= \sum_{i=0}^{M} L_i^{\text{scat}} W_0 \int_{s_i}^{s_{i+1}} T_{\sigma_t}(0,s) \sigma_t(s) ds \\ &= \sum_{i=0}^{M} L_i^{\text{scat}} W_0 \left(T_{\sigma_t}(0,s_i) - T_{\sigma_t}(0,s_{i+1})\right). \end{aligned} \quad (23)$$

In the classical control variate technique the integral of the control variate is analytically computed. However, Eq. 23 and $T_{\sigma_t}(0,s)$ cannot be evaluated analytically in practical cases. Instead of the analytic integral, a low variance, unbiased estimator for $W_0 T_{\sigma_t}(0,s)$ is $W(s)$ according to Eq. 13, thus we obtain

$$\Phi_{\text{main}}^{\text{scat}} \approx \sum_{i=0}^{M} L_i^{\text{scat}} \left(W(s_i) - W(s_{i+1})\right). \quad (24)$$

3.4. Path termination and splitting

Rao-Blackwellization eliminated random particle termination, thus a single ray will generate a binary tree where nodes are scattering points and leafs correspond to points where the weight becomes zero or the ray leaves the medium. It means that the simulation may spend a lot of time processing particles with weight of low absolute values (note that the particle weight can also be negative). To keep the number of particles and the size of the binary tree under control, we introduce random branch cutting.

The random termination affects particles with weights in $(-W_{\min}, W_{\min})$ [SKAS05, VK16]. A light particle in this window survives with probability p and is terminated with probability $1-p$. To compensate the not computed contribution, the weights of surviving particles are divided by p. The particle weight is increased to the window boundary if $p = |W|/W_{\min}$.

There is an even better strategy that considers control variate $L_{\text{main}}^{\text{scat}}(s)$. Instead of assuming that the incident radiance is zero for the terminated path, the incident radiance is approximated by $L_{\text{main}}^{\text{scat}}(s)/a(s)$. If the particle is terminated, we add $W(s)L_{\text{main}}^{\text{scat}}(s)/a(s)$ to its contribution. If the particle survives, the weight is compensated first as $W' = W/p$, and then the contribution is decreased by $(1-p)W'(s)L_{\text{main}}^{\text{scat}}(s)/a(s)$. The expectation will be correct:

$$\begin{aligned} \mathbf{E}[W(s)L(s,\vec{\omega})] = \mathbf{E}[W'(s)L(s,\vec{\omega}) - (1-p)W'(s)L_{\text{main}}^{\text{scat}}(s)/a(s)] \cdot p \\ + W(s)L_{\text{main}}^{\text{scat}}(s)/a(s) \cdot (1-p). \end{aligned} \quad (25)$$

Comparing to classical Russian roulette using the albedo or to the single particle model of [SKGM*17], our windowed termination scheme has advantages and disadvantages as well. Unlike other methods, it also uses the weight accumulated so far and not just the factor of the current interaction, so its variance is smaller. However, it means that if a particle weight should have a high absolute value, then it would create a large binary tree. To handle this, we use the windowed termination scheme before reaching depth 10, and then turn to the single particle model at higher depths.

Termination eliminates small weight particles. High weight particles also pose problems and cause high variance. Such high

weight particles are generated when $2\sigma_t - w$ is negative since in such cases the simulation produces a chain reaction (Eq. 14). To keep the weights under control, these particles can be split.

4. Rendering algorithm

According to the variance analysis, decomposition increases the variance, thus should be applied only if extinction σ_t is expensive to evaluate and the medium can be decomposed to a homogeneous and a residual volume [KHLN17]. The homogeneous volume should have an extinction that approximates the minimum of σ_t, otherwise the residual volume has negative extinction that leads to the exponential growth of the variance. To adapt this idea, we also decompose the medium into two media, the first has extinction $\sigma_{t,1} = \sigma_{\min}$ that approximates the minimum extinction of the original medium (Figure 4). The second medium has the residual extinction $\sigma_{t,2} = \sigma_t - \sigma_{\min}$. Low variance scattering is guaranteed if sampling extinction μ_1 is not less than $\sigma_{t,1}$, thus a reasonable choice is $\mu_1 = \sigma_{\min}$. There are several alternatives to select the control variate ν_1 in the homogeneous material. If the control extinction is set to the real extinction here, i.e. $\nu_1 = \sigma_{\min}$, then the difference extinction will be zero, $\delta_1 = 0$, and the transmitted weight will have a zero variance estimator in the homogeneous medium for the cost of generating two child particles at each interaction. However, this option has high computational cost if the volume is large or dense. Therefore, the better alternative is to set control extinction ν_1 to zero, which means that the weight of the transmitted particle is also zero, thus only the scattered particle needs to be traced in the homogeneous medium.

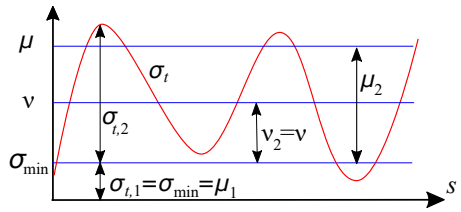

Figure 4: *Decomposition of the material and the setting of proxy extinctions. Parameters μ and $\sigma_{\min}$ approximate the maximum and the minimum of extinction σ_t but they are not necessarily upper and lower bounds. Material 1 is homogeneous with real extinction equal to $\sigma_{\min}$. Material 2 is the residual, in which control variate ν_2 approximates its mean.*

In the residual material, sampling extinction $\mu_2 = \mu - \sigma_{\min}$ is preferred to be close to the maximum extinction in this material. Control extinction of the second medium is the approximation of the mean of $\sigma_{t,2}$.

To represent primary proxy densities $\sigma_{\min}$, ν and μ, as well as the control variate of the scattered radiance, $L^{\text{scat}}_{\text{main}}$, we use a coarse grid structure similarly to [SKTM11, NSJ14, KHLN17]. Each cell contains one value for each of the four parameters, thus proxy densities and the control radiance are piece-wise constant functions along arbitrary rays. Their integration and free flight calculation with the sampling extinctions are straightforward, the algorithm should traverse the grid structure with a DDA algorithm. The way how a ray is traced is shown by Algorithm 1. First termination or splitting is executed, should any of them be needed. We use two running variables I_1 and I_2 for free flight sampling in the homogeneous and the residual materials, that are accumulated while the algorithm steps through the voxels.

Algorithm 1 Path Tracing of a ray defined by start $\vec{e}$, direction $\vec{\omega}$ and initial weight W. The medium is composed of two materials.

procedure TRACE($\vec{e}$,$\vec{\omega}$, W)
  if $|W| < W_{\min}$ **then** ▷ Termination
    $P = |W|/W_{\min}$
    if $P <$ rand() **then**
      return $W \cdot L^{\text{scat}}_{\text{main}}/a(\vec{p})$ ▷ Terminate
    else
      $W = W/P$ ▷ Compensate
      $\hat{\Phi} = -(1-P) \cdot W \cdot L^{\text{scat}}_{\text{main}}/a(\vec{p})$
  if $|W| > W_{\max}$ **then** ▷ Splitting
    $n =$(int)$|W|/W_{\max}$ ▷ number of children
    W /= n
    for $i = 1$ TO $n-1$ **do**
      $\hat{\Phi}$ += Trace($\vec{p}$, $\vec{\omega}$, W) ▷ Child
  $\hat{\Phi} = 0$ ▷ Estimated contribution
  $t_s = 0$ ▷ Start of the ray
  $I_1 = I_2 = 0$ ▷ Running variable of $\int \mu_m(\tau)d\tau$
  $M_1 = -\log(1 - \text{rand}())$ ▷ Threshold
  $M_2 = -\log(1 - \text{rand}())$
  while DDAGet($\vec{e}, \vec{\omega}, t_s \rightarrow t_e, \sigma_{\min}, \mu, \nu, L^{\text{scat}}_{\text{main}}$) **do**
    $\hat{\Phi}$ += $L^{\text{scat}}_{\text{main}} \cdot W$ ▷ Eq. 24
    $\mu_2 = \mu - \sigma_{\min}$
    $t_1 = t_s + (M_1 - I_1)/\sigma_{\min}$, $t_2 = t_s + (M_2 - I_2)/\mu_2$
    if $t_1 > t_e$ AND $t_2 > t_e$ **then** ▷ No scattering in cell
      I_1 += $\sigma_{\min} \cdot (t_e - t_s)$, I_2 += $\mu_2 \cdot (t_e - t_s)$
      $W = W \cdot \exp(-\nu(t_e - t_s))$ ▷ Eq. 10
      $\hat{\Phi}$ -= $L^{\text{scat}}_{\text{main}} \cdot W$ ▷ Eq. 24
      $t_s = t_e$ ▷ Step to next cell
    else ▷ Scattering in cell
      $\vec{\omega}'$ = PhaseFunction($\vec{\omega}$) ▷ Scattering direction
      if $t_1 < t_2$ **then** ▷ Collision with material 1
        $\vec{p} = \vec{e} + \vec{\omega} t_1$ ▷ Location of scattering
        $W = W \cdot \exp(-\nu(t_1 - t_s))$ ▷ Eq. 10
        $W^{\text{scat}} = W \cdot a(\vec{p})$ ▷ Eq. 12
        $\hat{\Phi}$ += Trace($\vec{p}$, $\vec{\omega}'$, W^{scat}) $- L^{\text{scat}}_{\text{main}} \cdot W^{\text{scat}}/a(\vec{p})$
      else ▷ Collision with material 2
        $\vec{p} = \vec{e} + \vec{\omega} t_2$ ▷ Location of scattering
        $W = W \cdot \exp(-\nu(t_2 - t_s))$ ▷ Eq. 10
        $\sigma_{t,2} = \sigma_t(\vec{p}) - \sigma_{\min}$ ▷ Real extinction fetch
        $W^{\text{scat}} = W \cdot a(\vec{p}) \cdot \sigma_{t,2}/\mu_2$ ▷ Eq. 12
        $\delta_2 = \sigma_{t,2} - \nu$
        $W = W \cdot (1 - \delta_2/\mu_2)$ ▷ Eq. 11
        $\hat{\Phi}$ += Trace($\vec{p}$, $\vec{\omega}'$, W^{scat}) $- L^{\text{scat}}_{\text{main}} \cdot W^{\text{scat}}/a(\vec{p})$
        $\hat{\Phi}$ += Trace($\vec{p}$, $\vec{\omega}$, W) ▷ Transmission
      $\hat{\Phi}$ -= $L^{\text{scat}}_{\text{main}} \cdot W$ ▷ Eq. 24
      return $\hat{\Phi}$
  $\hat{\Phi}$ += $W \cdot L^{\text{env}}(\vec{\omega})$ ▷ Environment illumination
  return $\hat{\Phi}$

The DDAGet function is used to query the grid structure, which takes the ray of start $\vec{e}$ and direction $\vec{\omega}$ as well as ray parameter t_s of the current point. The function determines ray parameter t_e where the ray exits the current cell, and the pre-computed proxy densities and the main scattered radiance are returned. Comparing

free flight distances t_1 and t_2 to the distance of the exit point of the cell, t_e, the algorithm determines whether there are potential collisions in the current cell and which material is hit first. If there is no interaction in this cell, running variables are updated and we step to the next cell. If there are potential collisions in both materials, the fist collision of smaller distance is selected, where scattering direction $\vec{\omega}'$ is sampled, and the weights are updated depending on the type of the hit material particle. Note that we choose the control variate of the extinction in the homogeneous material in a way to make the weight of the transmitted particle zero, which needs not to be computed. In the residual material, the same Trace function is called both for the scattered and the transmitted light particles.

The grid structure can be initialized in a pre-processing phase when we take extinction samples in each cell and assign the average, minimum and the maximum to the ν, $\sigma_{\min}$ and μ values of the cell. Setting the average to ν is optimal from the point of view of the transmitted contribution but is not necessarily the best option if scattering is more significant. Control radiance $L^{\text{scat}}_{\text{main}}$ can be estimated by sending pilot rays from random points.

4.1. Multi-spectral tracking

In the transformed model, interaction points are found with sampling extinctions $\sigma_{\min}$ and $\mu_2 = \mu - \sigma_{\min}$, which should be scalars. However, all other parameters, including particle weights, radiance values, real and control extinctions, as well as scattering albedos can be vectors, just the arithmetic operations should be applied on them in element-wise manner, and the absolute value of the weight should be replaced by a vector norm. Thus our model can also be used in a spectral rendering framework where a light particle simultaneously carries energy or importance on several wavelengths. We follow the approach of [KHLN17] and set $\sigma_{\min}$ and μ to the minimum and maximum of the wavelength dependent extinctions, respectively.

5. Results

To demonstrate the proposed approach, we consider three different media, the first is defined by an analytic function and rendered with constant control variates, the second is generated procedurally by multi-octave Perlin noise, and the third is a high resolution voxel grid. In the second and third cases, we also evaluate the benefits of the coarse grid structure. Throughout our tests we use the Henyey-Greenstein phase function with anisotropy parameter g. When not given explicitly, we consider the $g = 0$ isotropic case. We omit the negligible pre-computation costs from the comparison and use the Root Mean Squared Error (RMSE) to quantify the accuracy, which is rounded to the nearest integer for clarity.

5.1. Analytic medium and constant control variates

The first test medium [SKGM*17] is a sphere of radius 10 and center $(0,0,10)$, and its extinction at point (x,y,z) is

$$\sigma_t = \left(\frac{\cos(3(x+y+z)/2)+1}{2}\right)^5 \cdot \frac{\sin(z/2)+2}{3}(1-\sigma_0)+\sigma_0. \tag{26}$$

Parameter σ_0 defines the minimum of the extinction and describes how strongly homogeneous the medium is. The environment illumination is an RGB color box. The scattering albedo is 0.7. For this model, we use constant μ, ν and $L^{\text{scat}}_{\text{main}}$ values.

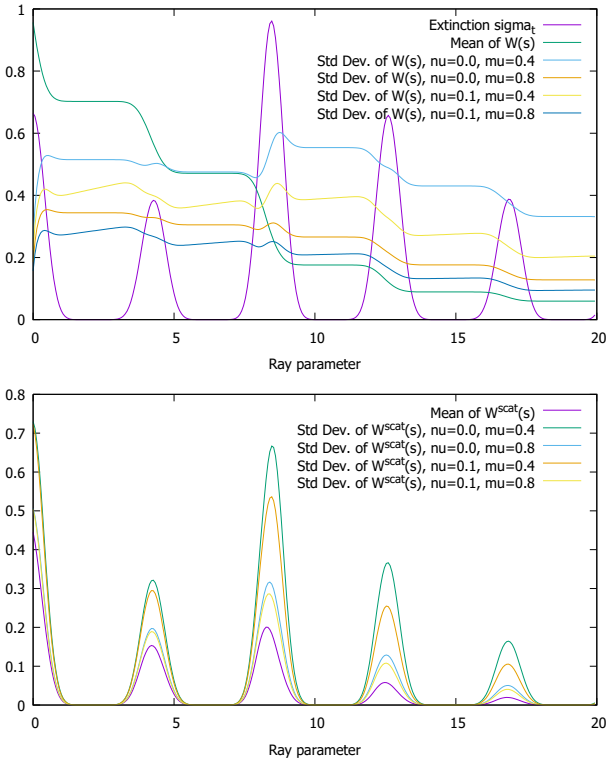

Figure 5: *The extinction, the mean and the standard deviations of transmitted weight $W(s)$ (upper plot) and scattered weight $W^{\text{scat}}(s)$ (lower plot) along the ray going through the center of the analytic medium when $\sigma_0 = 0$.*

First, we do not apply decomposition, i.e. $\mu_1 = \nu_1 = \sigma_{\min} = 0$, $\mu_2 = \mu$, and set minimum extinction σ_0 to zero, which corresponds to the most heterogeneous case. Figure 5 depicts the extinction and the statistical properties of weights $W(s)$ and $W^{\text{scat}}(s)$ along a ray of axis z that goes through the center of the medium. Note that increasing μ makes the variance smaller. We can also observe that control variate ν improves the estimate of $W(s)$ and also of $W^{\text{scat}}(s)$ if μ is far from being a majorant of the extinction. However, control variate ν does not help much to reduce the variance of the scattered contribution if μ is a majorant. Generally, larger sampling extinction also increases the computational cost, since samples are taken more frequently. To provide a fair comparison, the number of cast rays is set to make the number of medium fetches equal in all compared cases (we target 200 fetches per pixel). The corresponding images and their RMSE values are presented by Figure 6. Note that increasing sampling extinction μ makes the algorithm more effective until it reaches 0.8, which is a little less than the maximum extinction of the medium. The optimal control variate ν is about 0.1, which is the average value of σ_t. We set $W_{\max} = 10$ in all tests since the algorithm is not too sensitive to this parameter.

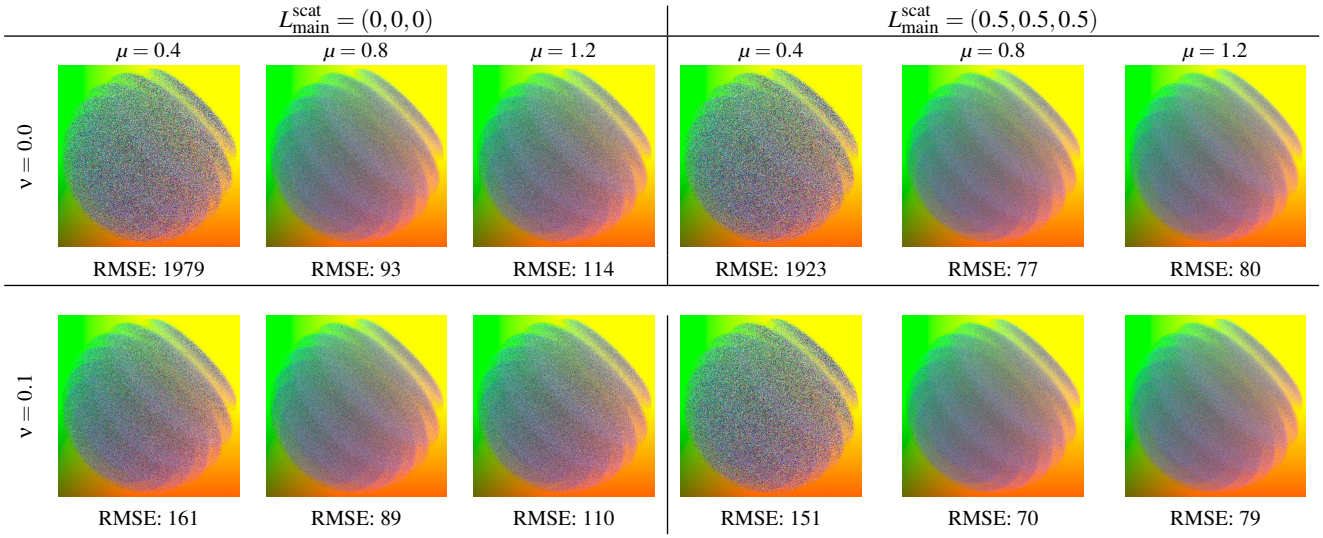

Figure 6: *Particle tracing of the medium of Eq. 26 setting $\sigma_0 = 0$ with 200 medium fetches per pixel. No control variate is used for scattered radiance L^{scat}_{main} (left hand side) and $L^{scat}_{main} = (0.5, 0.5, 0.5)$ (right hand side). Random termination is executed for those light particles where $|W| < W_{min} = 0.1$. No medium decomposition is applied, i.e. $\sigma_{min} = 0$.*

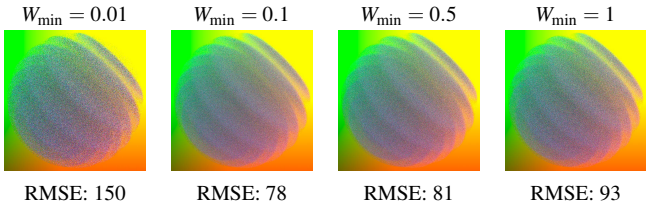

Figure 7: *The effect of termination window W_{min} on the results. We set $\mu = 0.8$ and $\nu = 0.1$.*

Figure 7 shows the results when termination window W_{min} changes. Too large window adds noise similarly to Russian-roulette. On the other hand, if the window is small, then small weight particles are traced, thus in the equal effort comparison, the number of rays should be reduced, which also increases the variance. From now on, $W_{min} = 0.1$ is assumed.

As demonstrated by Figure 8, the direction independent control variate for the scattered radiance can significantly reduce the error even when the medium is not isotropic.

Figure 9 compares the new method to weighted decomposition tracking [KHLN17] and to an improved version of residual ratio tracking [NSJ14] that is generalized to handle non-majorant sampling density and non-minorant control variate [SKGM*17] and to compute not only transmittance but also scattering. To use residual ratio tracking in multiple scattering simulation, we applied the same termination scheme as in the new method. Generalized residual ratio tracking reduces the absolute weight of light particles even during free flight, which is advantageous in variance reduction especially when the sampling density is low. On the other hand, decomposition tracking can save extinction coefficient evaluations, but increases the variance especially when the residual volume has

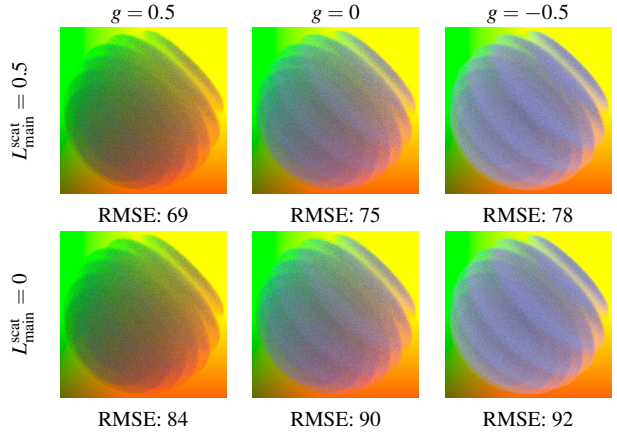

Figure 8: *The effect of the control variate for the radiance in forward scattering ($g = 0.5$), isotropic ($g = 0$), and backward scattering ($g = -0.5$) media, when $\mu = 1$, $\sigma_0 = 0$ and $\nu = 0$.*

negative extinction. Our method combines the advantages of both residual ratio tracking and decomposition tracking, and further improves the variance by the exploitation of the scattered radiance control variate. In this test, we have violated the minorant and majorant conditions since the minimum of the real extinction is 0 and the maximum is 1 in this volume. Note that decomposition tracking is very sensitive to this issue, residual ratio tracking is less so, and the new method is the most robust.

Figure 10 is the repetition of the same experiment for spectral tracking when the environment illumination is white but the extinctions on the wavelengths of r,g,b are proportional to the absolute value of the x, y, z coordinates expressing the direction from the center of the volume. For comparison, we implemented the

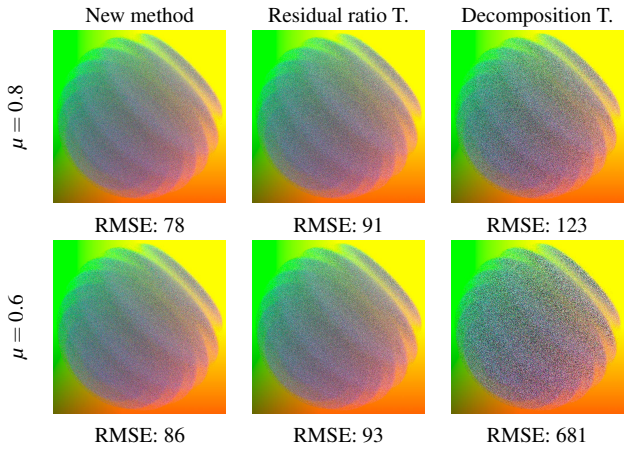

Figure 9: *Rendering the analytic model of* $\sigma_0 = 0$ *when the minorant and majorant conditions are violated, i.e.* μ *is 0.8 and 0.6, respectively, which are less than* $\max(\sigma_t) = 1$*, and* $\sigma_{\min} = 0.01$*, which is larger than* $\min(\sigma_t) = 0$*. We compare the new method to generalized residual ratio tracking and decomposition tracking when all methods compute 200 evaluations of* σ_t *per pixel. We use constant control proxy values* $\nu = 0.1$*,* $L^{\text{scat}}_{\text{main}} = (0.5, 0.5, 0.5)$*.*

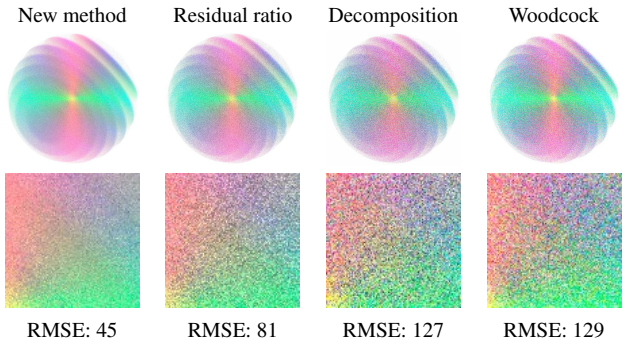

Figure 10: *Spectral tracking with the settings of Figure 9 when* $\mu = 0.8$*,* $g = -0.5$ *and extinction* σ_t *is modulated with the coordinates of the direction vector from the center of the medium, illuminated by a uniform white environment map.*

history-aware-avg version of Decomposition tracking, and Woodcock tracking is executed independently on the three wavelengths.

Figure 11 and Table 1 compares the methods again when all majorant and minorant conditions are met and a non-zero minimum extinction σ_0 is introduced. Decomposition tracking can benefit the most from the positive minimum extinction, but the new method still outperforms all other compared techniques. We also test the methods in this case for spectral tracking in Figure 12, and the conclusions are similar to the scalar case.

5.2. Procedural medium generated with 12-octave Perlin noise

Now, we render a cloud model generated by 12-octave Perlin noise and illuminated by a strongly inhomogeneous environment map

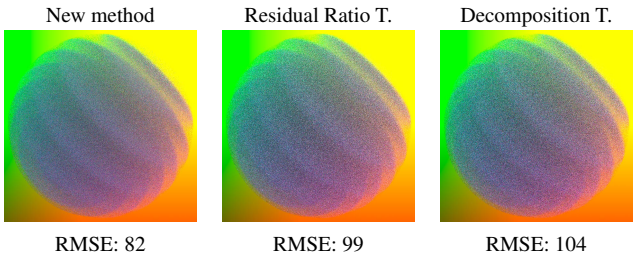

Figure 11: *Rendering the analytic model of* $\sigma_0 = 0.05$ *with the three compared methods when the minorant and majorant conditions are not violated, i.e.* $\mu = 1$*,* $\sigma_{\min} = 0.05$*.*

σ_0	New meth.	Woodcock	Residual ratio	Decomposition
0	71	120	90	110
0.05	82	118	99	104
0.1	85	114	103	96
0.2	66	112	106	88
0.5	50	110	108	51

Table 1: *RMSE values for different degrees of homogeneity* σ_0 *when the minorant and majorant conditions are not violated. Sampling density* μ *is 1 for all models. The new method sets* ν *to 0.1 and makes* $\sigma_{\min}$ *equal to* σ_0*. Generalized residual ratio tracking sets* ν *to* σ_0*, and decomposition tracking* $\sigma_{\min}$ *to* σ_0*.*

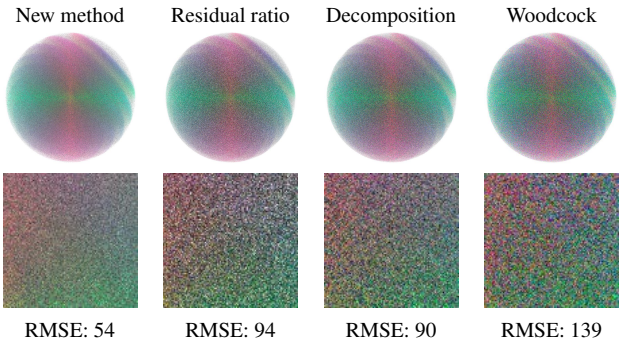

Figure 12: *Spectral tracking with the settings of Figure 11, i.e. when the minorant a majorant conditions are not violated, and* $\sigma_{\min} = \sigma_0$*,* $g = 0.5$*, and above* $\sigma_0 = 0.1$ *extinction* σ_t *is modulated with the coordinates of the direction vector from the center of the medium, illuminated by a uniform white environment map.*

(Figure 13). We also repeat the test with a lower dynamic range environment map (Figure 15). The scattering albedo is 0.9 to emphasize the multiple scattering effects.

Firstly, control variates are computed in a preprocessing phase and stored on a coarse grid of 32^3 resolution. This precomputation takes only 160 seconds on a single-core CPU. Figure 13 shows the comparison of the new method to previous work for 3k volume fetches per pixel on average. Note that decomposition tracking, residual ratio tracking and the new method work on the same grid structure and thus rendering times are similar (10 minutes). If

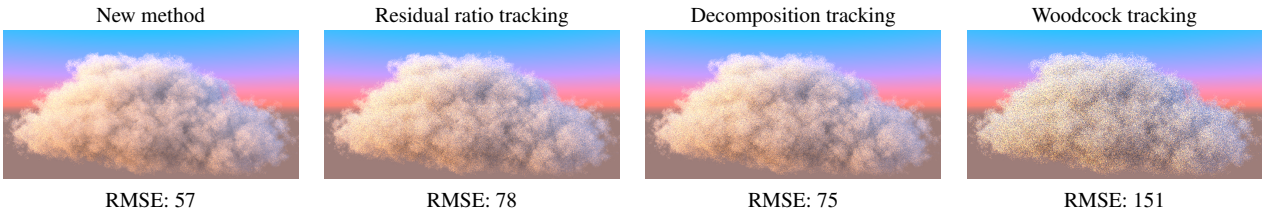

Figure 13: *Procedurally generated cloud with 3k volume evaluations per pixel.*

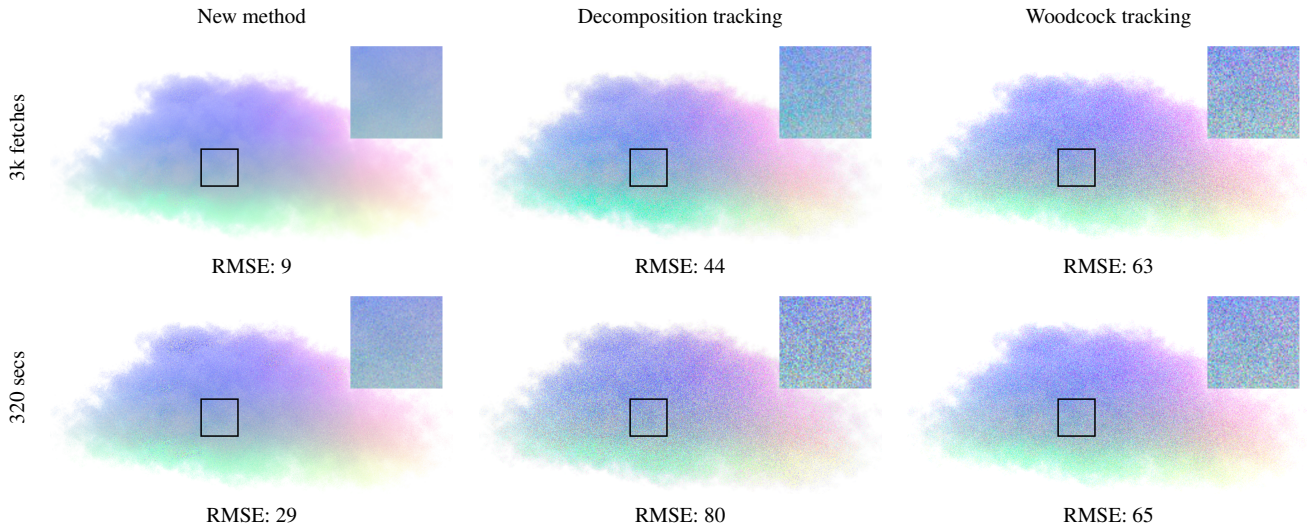

Figure 14: *Spectral tracking with equal number of volume fetches per pixel (3k) in the upper row and with equal time (320 seconds) in the lower row. The environment illumination is constant, the medium is forward scattering ($g = 0.5$) and its extinction σ_t varies in space.*

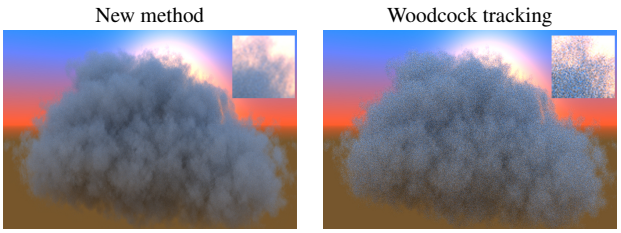

Figure 15: *Comparison to Woodcock tracking when both methods fetch the procedural data by 170 times per pixel.*

the grid resolution is decreased from 32^3 to 16^3 and 8^3, the error increases from 57 to 61 and 67, respectively, i.e. even a relatively low resolution grid may lead to a noticeable improvement over Woodcock tracking.

Figure 14 compares the new method performing spectral tracking to the history-aware-max version of Decomposition tracking and to Woodcock tracking executed independently on the three wavelengths. The anisotropy parameter of the Henyey-Greenstein phase function is $g = 0.5$.

5.3. Participating medium defined by a voxel grid

The final test case is a bunny shaped medium that is defined by a voxel array of $577 \times 572 \times 438$ resolution (Figure 1). The scattering albedo is 0.95 and the model has large physical extent (60 units radius) with an average and maximum extinctions of 0.7 and 5.9, respectively. With such parameters, the transmittance is practically zero and the multiple scattering becomes the dominant phenomenon. The grid resolution is $64 \times 64 \times 64$. First, we apply fairly high dynamic range illumination mimicking a directional light behind the camera. Figure 1 compares the new method with Woodcock tracking, residual ratio tracking and decomposition tracking allowing 50k volume fetches per pixel. Close-ups are also shown by Figure 16 when the rendering is executed with 10k fetches per pixel.

We also render a less dense bunny with a lower dynamic range environment map (Figure 17). As this is a simpler problem, the number of required samples can be significantly lower.

6. Conclusions

This paper presented a transformed model for participating media in which the steps of multiple scattering simulation can be done with analytic formulae, and the result is an unbiased estimate with low variance. The method does not require the a-priori knowledge

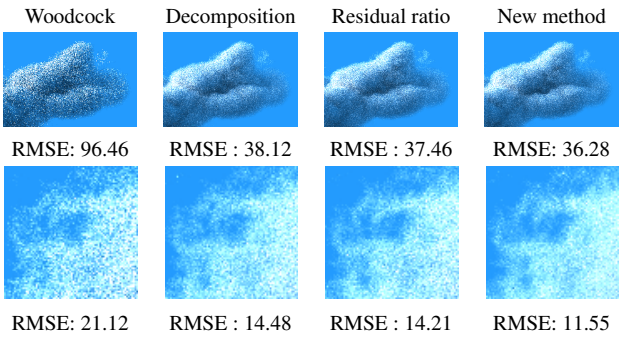

Figure 16: *Dense medium rendered with 10k fetches per pixel.*

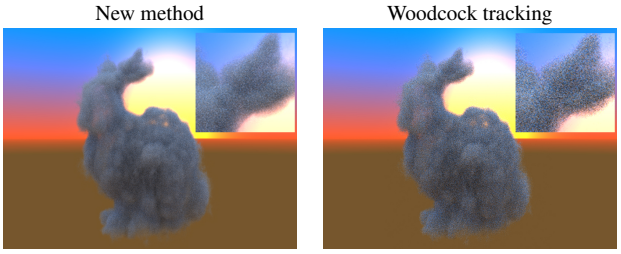

Figure 17: *Sparser medium with 40 fetches per pixel.*

of the extinction function, which is an advantage in case of procedural models. We can use approximate sampled values to set the sampling extinction and the control variates that can reduce the variance. The approximation error does not compromise the unbiasedness of the result, only the variance may increase. Control variates are stored in a coarse grid, and we have shown that even low resolution grids of low overhead can significantly reduce the variance of the Monte Carlo estimation.

Acknowledgement

The bunny model has been downloaded from http://www.openvdb.org. This work has been supported by OTKA K–124124, VKSZ-14 PET/MRI 7T (Hungary), TIN2016-75866-C3-3-R (Spain), Erasmus Mundus Action 2 PANTHER and EFOP-3.6.2-16-2017-00013 (EU).

Appendix 1: Statistical analysis of transmitted weight W

Let us consider how transmitted weight W changes when an infinitesimal step $\mathrm{d}s$ is made along the ray:

$$W(s+\mathrm{d}s) = W(s)\mathcal{T}(s,\mathrm{d}s), \tag{27}$$

where $\mathcal{T}$ is a random variable describing the differential transmission:

$$\mathcal{T}(s,\mathrm{d}s) = \exp(-\nu(s)\mathrm{d}s)\left(1-\frac{\delta_m(s)}{\mu_m(s)}\right) \tag{28}$$

if the light particles collides with a material particle of type m, which happens with probability $\mu_m(s)\mathrm{d}s$. If there is no collision, then

$$\mathcal{T}(s,\mathrm{d}s) = \exp(-\nu(s)\mathrm{d}s), \tag{29}$$

which occurs with probability $1-\mu\mathrm{d}s$. The probability of multiple collisions in $\mathrm{d}s$ is in $o(\mathrm{d}s)$.

The expectation of the differential transmission is

$$\begin{aligned}\mathbf{E}[\mathcal{T}(s,\mathrm{d}s)] &= \exp(-\nu(s)\mathrm{d}s)\left(\sum_m\left(1-\frac{\delta_m(s)}{\mu_m(s)}\right)\mu_m(s)\mathrm{d}s+(1-\mu(s)\mathrm{d}s)\right)\\ &+ o(\mathrm{d}s).\end{aligned} \tag{30}$$

This expectation is rewritten by applying a linear expansion for the exponential $\exp(-\nu(s)\mathrm{d}s)$ and hiding all higher order terms of $\mathrm{d}s$ behind the little o notation:

$$\begin{aligned}\mathbf{E}[\mathcal{T}(s,\mathrm{d}s)] &= 1-(\nu(s)+\mu(s)-\sum_m(\mu_m(s)-\delta_m(s)))\mathrm{d}s+o(\mathrm{d}s)\\ &= 1-\sigma_t(s)\mathrm{d}s+o(\mathrm{d}s).\end{aligned} \tag{31}$$

since $\nu+\mu-\sum_m\mu_m+\sum_m\delta_m=\nu+\delta=\sigma_t$.

The variance of the differential transmittance estimator is:

$$\begin{aligned}\mathbf{V}[\mathcal{T}(s,\mathrm{d}s)] &= \sum_m\left(\exp(-\nu(s)\mathrm{d}s)\left(1-\frac{\delta_m(s)}{\mu_m(s)}\right)\right)^2\mu_m(s)\mathrm{d}s\\ &+ (\exp(-\nu(s)\mathrm{d}s))^2(1-\mu(s)\mathrm{d}s)\\ &- (1-\sigma_t(s)\mathrm{d}s)^2+o(\mathrm{d}s)\\ &= \left(\sum_m\mu_m(s)-2\delta_m(s)+\frac{\delta_m^2(s)}{\mu_m(s)}\right)\mathrm{d}s\\ &+ (1-2\nu(s)\mathrm{d}s-\mu(s)\mathrm{d}s)-(1-2\sigma_t(s)\mathrm{d}s)+o(\mathrm{d}s)\\ &= \sum_m\frac{\delta_m^2(s)}{\mu_m(s)}\mathrm{d}s+o(\mathrm{d}s)=w(s)+o(\mathrm{d}s).\end{aligned} \tag{32}$$

Using the $\mathbf{E}[XY]=\mathbf{E}[X]\mathbf{E}[Y]$ identity for independent random variables $X=W$ and $Y=\mathcal{T}$, the expectation of the transmission is

$$\mathbf{E}[W(s+\mathrm{d}s)] = \mathbf{E}[W(s)]\mathbf{E}[\mathcal{T}(s,\mathrm{d}s)] = \mathbf{E}[W(s)](1-\sigma_t(s)\mathrm{d}s+o(\mathrm{d}s)).$$

Subtracting $\mathbf{E}[W(s)]$ from both sides, dividing the equation by $\mathrm{d}s$ and taking the $\mathrm{d}s\to 0$ limit, we can establish a differential equation for the expectation

$$\frac{\mathrm{d}\mathbf{E}[W(s)]}{\mathrm{d}s} = -\sigma_t(s)\mathrm{d}s.$$

Taking into account the $\mathbf{E}[W(s)]=W_0$ initial condition, this equation can be solved as

$$\mathbf{E}[W(s)] = W_0\mathcal{T}_{\sigma_t}(0,s). \tag{33}$$

For the variance of the weight in the transformed model, we exploit the identity $\mathbf{V}[XY]=\mathbf{V}[X]\mathbf{V}[Y]+\mathbf{V}[X]\mathbf{E}^2[Y]+\mathbf{V}[Y]\mathbf{E}^2[X]$ valid for independent random variables $X=W(s)$ and $Y=\mathcal{T}(s,\mathrm{d}s)$:

$$\begin{aligned}\mathbf{V}[W(s+\mathrm{d}s)] &= \mathbf{V}[W(s)]\left(\mathbf{V}[\mathcal{T}(s,\mathrm{d}s)]+\mathbf{E}^2[\mathcal{T}(s,\mathrm{d}s)]\right)\\ &+ \mathbf{E}^2[W(s)]\mathbf{V}[\mathcal{T}(s,\mathrm{d}s)]\\ &= \mathbf{V}[W(s)](w(s)\mathrm{d}s+1-2\sigma_t(s)\mathrm{d}s)\\ &+ W_0^2T_{\sigma_t}^2(0,s)w(s)\mathrm{d}s+o(\mathrm{d}s),\end{aligned}$$

We again obtain a differential equation

$$\frac{\mathrm{d}\mathbf{V}[W(s)]}{\mathrm{d}s} = (w(s)-2\sigma_t(s))\,\mathbf{V}[W(s)]+w(s)W_0^2T_{\sigma_t}^2(0,s).$$

The solution of this equation is

$$\begin{aligned}\mathbf{V}[W(s)] &= W_0^2\exp\left(-\int_0^s 2\sigma_t(\vec{p}(\tau))\mathrm{d}\tau\right)\left(\exp\left(\int_0^s w(\tau)\mathrm{d}\tau\right)-1\right)\\ &= W_0^2\left(T_{2\sigma_t-w}(0,s)-T_{2\sigma_t}(0,s)\right).\end{aligned} \tag{34}$$

Appendix 2: Statistical analysis of the scattered weight

Using Eq. 12 defining the weight of the scattered light particle and realizing that a scattering on material of type m happens with probability $\mu_m(s)\mathrm{d}s$ in $[s, s+\mathrm{d}s)$, the random scattered weight is

$$W^{\mathrm{scat}}(s) = W(s-0)a(s)\frac{\sigma_{t,m}(s)}{\mu_m(s)} \tag{35}$$

with probability $\mu_m(s)\mathrm{d}s$ and zero with probability $1-\mu\mathrm{d}s$.

Random variable $W^{\mathrm{scat}}(s)$ is the product of two independent random variables, $X = W(s-0)$ and $Y = a(s)\frac{\sigma_{t,m}(s)}{\mu_m(s)}$ occurring with probability $\mu_m(s)\mathrm{d}s$ and is zero otherwise. The expectation is the product of the expectations of the two factors:

$$\begin{aligned}
\mathbf{E}[W^{\mathrm{scat}}(s)] &= \mathbf{E}[W(s-0)]\mathbf{E}\left[a(s)\frac{\sigma_{t,m}(s)}{\mu_m(s)}\right] \\
&= W_0 T_{\sigma_t}(0,s)a(s)\sum_m \sigma_{t,m}(s)\mathrm{d}s \\
&= W_0 a(s) T_{\sigma_t}(0,s)\sigma_t(s)\mathrm{d}s.
\end{aligned} \tag{36}$$

The calculation of the variance is also based on identity $\mathbf{V}[XY] = (\mathbf{V}[X]+\mathbf{E}^2[X])\mathbf{V}[Y] + \mathbf{V}[X]\mathbf{E}^2[Y]$. The expectation and the variance of X is taken from Eqs. 33 and 34, and the expectation of Y is proportional to $\mathrm{d}s$, thus its square is in $o(\mathrm{d}s)$:

$$\begin{aligned}
\mathbf{V}[X]+\mathbf{E}^2[X] &= W_0^2\left(T_{2\sigma_t - w}(0,s) - T_{2\sigma_t}(0,s)\right) + W_0^2 T_{2\sigma_t}(0,s) \\
&= W_0^2 T_{2\sigma_t - w}(0,s), \\
\mathbf{E}^2[Y] &= o(\mathrm{d}s), \\
\mathbf{V}[Y] &= \sum_m \left(a(s)\frac{\sigma_{t,m}(s)}{\mu_m(s)}\right)^2 \mu_m(s)\mathrm{d}s + o(\mathrm{d}s).
\end{aligned} \tag{37}$$

The variance is then:

$$\mathbf{V}[W^{\mathrm{scat}}(s)] = W_0^2 T_{2\sigma_t - w}(0,s)a^2(s)\sum_m \frac{\sigma_{t,m}^2(s)}{\mu_m(s)}\mathrm{d}s + o(\mathrm{d}s). \tag{38}$$

References

[AK90] ARVO J., KIRK D.: Particle transport and image synthesis. In *Computer Graphics (SIGGRAPH '90 Proceedings)* (1990), pp. 63–66. 1

[BJ17] BITTERLI B., JAROSZ W.: Beyond points and beams: Higher-dimensional photon samples for volumetric light transport. *ACM Trans. Graph. 36*, 4 (2017). 1

[Chr03] CHRISTENSEN P. H.: Adjoints and importance in rendering: An overview. *IEEE Trans. on Visualization and Computer Graphics 09*, 3 (2003), 329–340. 3

[CPP*05] CEREZO E., PÉREZ F., PUEYO X., SERON F. J., SILLION F. X.: A survey on participating media rendering techniques. *The Visual Computer 21*, 5 (2005), 303–328. 1

[Fat09] FATTAL R.: Participating media illumination using light propagation maps. *ACM Trans. Graph. 28*, 1 (2009), 1–11. 1

[GBC*13] GALTIER M., BLANCO S., CALIOT C., COUSTET C., DAUCHET J., HAFI M. E., EYMET V., FOURNIER R., GAUTRAIS J., KHUONG A., PIAUD B., TERRÉE G.: Integral formulation of null-collision monte carlo algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer 125* (2013), 57 – 68. 2

[JC98] JENSEN H. W., CHRISTENSEN P. H.: Efficient simulation of light transport in scenes with participating media using photon maps. *SIGGRAPH '98 Proceedings* (1998), 311–320. 1

[JLSJ11] JOHNSON J. M., LACEWELL D., SELLE A., JAROSZ W.: Gaussian quadrature for photon beams in Tangled. In *ACM SIGGRAPH 2011 Talks*, (2011), pp. 54:1–54:1. 3

[JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *SIGGRAPH '01 Proceedings* (2001), pp. 511–518.

[JNT*11] JAROSZ W., NOWROUZEZAHRAI D., THOMAS R., SLOAN P.-P., ZWICKER M.: Progressive photon beams. *ACM Trans. Graph. 30*, 6 (2011). 1

[KF12] KULLA C., FAJARDO M.: Importance sampling techniques for path tracing in participating media. *Comput. Graph. Forum 31*, 4 (2012), 1519–1528. 1

[KGH*14] KŘIVÁNEK J., GEORGIEV I., HACHISUKA T., VÉVODA P., ŠIK M., NOWROUZEZAHRAI D., JAROSZ W.: Unifying points, beams, and paths in volumetric light transport simulation. *ACM Trans. Graph. 33*, 4 (2014), 103:1–103:13. 1

[KHLN17] KUTZ P., HABEL R., LI Y. K., NOVÁK J.: Spectral and decomposition tracking for rendering heterogeneous volumes. *ACM Trans. Graph. 36*, 4 (2017), 111:1–111:16. 2, 3, 6, 7, 8

[LW95] LAFORTUNE E. P., WILLEMS Y. D.: A 5d tree to reduce the variance of Monte Carlo ray tracing. In *Rendering Techniques '95* (1995), pp. 11–20. 2, 5

[Mun14] MUNOZ A.: Higher order ray marching. *Comput. Graph. Forum 33*, 8 (2014), 167–176. 3

[NNDJ12] NOVÁK J., NOWROUZEZAHRAI D., DACHSBACHER C., JAROSZ W.: Virtual ray lights for rendering scenes with participating media. *ACM Trans. Graph. 31*, 4 (2012), 60:1–60:11. 1

[NSJ14] NOVÁK J., SELLE A., JAROSZ W.: Residual ratio tracking for estimating attenuation in participating media. *ACM Trans. Graph. 33*, 6 (2014), 179:1–179:11. 2, 6, 8

[PKK00] PAULY M., KOLLIG T., KELLER A.: Metropolis light transport for participating media. In *Rendering Techniques* (2000), pp. 11–22.

[PWP08] PEGORARO V., WALD I., PARKER S. G.: Sequential Monte Carlo adaptation in low-anisotropy participating media. *Comput. Graph. Forum 27*, 4 (2008), 1097–1104. 5

[QXFN07] QIU F., XU F., FAN Z., NEOPHYTOS N.: Lattice-based volumetric global illumination. *IEEE Trans. on Visualization and Computer Graphics 13*, 6 (2007), 1576–1583.

[RJN16] ROUSSELLE F., JAROSZ W., NOVÁK J.: Image-space control variates for rendering. *ACM Trans. Graph. 35*, 6 (2016), 169:1–169:12. 2

[RSK08] RAAB M., SEIBERT D., KELLER A.: Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer, 2008, pp. 591–606. 3

[SKAS05] SZIRMAY-KALOS L., ANTAL G., SBERT M.: Go with the winners strategy in path tracing. In *WSCG* (2005), pp. 49–56. 5

[SKGM*17] SZIRMAY-KALOS L., GEORGIEV I., MAGDICS M., MOLNÁR B., LÉGRÁDY D.: Unbiased light transport estimators for inhomogeneous participating media. *Comput. Graph. Forum 36*, 2 (2017), 9–19. 2, 4, 5, 7, 8

[SKSU05] SZIRMAY-KALOS L., SBERT M., UMENHOFFER T.: Real-time multiple scattering in participating media with illumination networks. In *Eurographics Symposium on Rendering* (2005), pp. 277–282.

[SKTM11] SZIRMAY-KALOS L., TÓTH B., MAGDICS M.: Free Path Sampling in High Resolution Inhomogeneous Participating Media. *Comput. Graph. Forum 30*, 1 (2011), 85–97. 2, 6

[VK16] VORBA J., KŘIVÁNEK J.: Adjoint-driven russian roulette and splitting in light transport simulation. *ACM Trans. Graph. 35*, 4 (2016), 42:1–42:11. 5

[WMHL65] WOODCOCK E., MURPHY T., HEMMINGS P., LONGWORTH S.: Techniques used in the GEM code for Monte Carlo neutronics calculation. In *Proc. Conf. Applications of Computing Methods to Reactors, ANL-7050* (1965). 2

[YIC*10] YUE Y., IWASAKI K., CHEN B.-Y., DOBASHI Y., NISHITA T.: Unbiased, adaptive stochastic sampling for rendering inhomogeneous participating media. *ACM Trans. Graph. 29*, 6 (2010), 177:1–177:8. 2