Residual Ratio Tracking for Estimating Attenuation in Participating Media

Jan Novák 1,2 Andrew Selle¹

Wojciech Jarosz²

¹Walt Disney Animation Studios

²Disney Research Zürich



(a) Path traced light transport in clouds

(b) Delta tracking

(d) Residual ratio tracking

Figure 1: A cloudy sky rendered with our residual ratio tracking estimator for computing transmittance in heterogeneous volumes. Our technique is unbiased, outperforms the delta tracking-based estimator (b), and fits well into path-tracing, production frameworks. The insets show renderings of absorptive-only (top) and scattering (bottom) clouds; the transmittance was estimated using delta tracking (b), ratio tracking (c), and residual ratio tracking (d) with a roughly equal cost reported as the number of extinction coefficient evaluations. Images ©Disney.

Abstract

Evaluating transmittance within participating media is a fundamental operation required by many light transport algorithms. We present ratio tracking and residual tracking, two complementary techniques that can be combined into an efficient, unbiased estimator for evaluating transmittance in complex heterogeneous media. In comparison to current approaches, our new estimator is unbiased, yields high efficiency, gracefully handles media with wavelength dependent extinction, and bridges the gap between closed form solutions and purely numerical, unbiased approaches. A key feature of ratio tracking is its ability to handle negative densities. This in turn enables us to separate the main part of the transmittance function, handle it analytically, and numerically estimate only the residual transmittance. In addition to proving the unbiasedness of our estimators, we perform an extensive empirical analysis to reveal parameters that lead to high efficiency. Finally, we describe how to integrate the new techniques into a production path tracer and demonstrate their benefits over traditional unbiased estimators.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism-Raytracing;

Keywords: participating media, transmittance, delta tracking

Links: DL PDF WEB

1 Introduction

The world around us is filled with participating media that attenuates and scatters light as it travels from light sources, to surfaces, and finally to our eyes. Simulating this transport in heterogeneous participating media-such as smoke, clouds, nuclear reactor housings, biological tissue, or other volumetric datasets-is important in many fields, ranging from neutron transport, to medical physics, scientific visualization, and film and visual effects production.

Monte Carlo (MC) path sampling approaches, including variants of path tracing [Kajiya 1986], bidirectional path tracing [Lafortune and Willems 1993; Veach and Guibas 1994; Pauly et al. 2000], or many-light methods [Keller 1997; Dachsbacher et al. 2013], have proven to be practical approaches for accurately approximating this light transport. All of these rely on generating random paths between the light(s) and the sensor, and there has been extensive research on importance sampling such paths to obtain low-noise images [Raab et al. 2008; Kulla and Fajardo 2012; Georgiev et al. 2013].

Central to all these approaches, however, is the need to evaluate transmittance-or fractional visibility-between two points in the scene. This is needed for shadow connections between light- and camera-subpaths, for rendering colored media, or simply for evaluating the fractional visibility to solid surfaces. In homogeneous media, computing transmittance is trivial since it accepts a simple exponential analytic form. Unfortunately, in heterogeneous media an expensive numerical approximation is necessary and relatively little research has been done on performing this critical operation efficiently. Traditional ray marching techniques result in unpredictable, systematic bias and require many fine steps in high-resolution data. On the other hand, while unbiased free-flight sampling techniques like delta tracking [Woodcock et al. 1965] can be adapted to estimate transmittance, they result in coarse, binary estimators with high variance. These options lead to either substantially increased render times or artifacts in the form of bias or noise. In this paper we are interested in an efficient, unbiased evaluation of transmittance in highly-complex heterogeneous media-the common case in visual effects and film production, and other graphics applications.

[©]ACM, 2014. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in ACM Transactions on Graphics, 33, 6, November 2014. doi.acm.org/10.1145/2661229.2661292



Figure 2: *Homogeneous absorbing (top) and scattering (bottom) media rendered with path tracing that evaluates transmittance analytically* (**a**), *using delta tracking* (**b**), *and using our new method* (**c**).

Unfortunately, due to the aforementioned limitations, rendering heterogeneous media often takes orders of magnitude longer than homogeneous media, in otherwise identical scenes. In fact, simply rendering a homogeneous medium using only the techniques available for heterogeneous media (i.e. disallowing analytic transmittance calculation) immediately leads to significantly higher render times, higher bias, or higher variance—even though the scene is homogeneous, see Figure 2. While this is a sad state of affairs, it also indicates that the problem is not inherent to the heterogeneity of the medium itself, but rather stems from the lack of appropriate tools to compute transmittance accurately and efficiently in this more general case.

In this paper, we tackle this fundamental problem that is shared by many participating media rendering algorithms. We first survey existing approaches (Section 2) for computing transmittance-both from graphics and other related fields like neutron transport-to identify their relative weaknesses. We then propose two novel and complementary approaches to compute unbiased, low-variance transmittance estimates in general heterogeneous media (Section 3), which we call ratio tracking and residual tracking-the latter of which gracefully and automatically simplifies to the standard analytic result when the medium is homogeneous. We prove the correctness and unbiasedness of the proposed techniques (please see the supplementary material), and perform an extensive empirical variance analysis (Section 4) evaluating our approaches against the state-ofthe-art. Our techniques can be used as simple drop-in replacements for evaluating the transmittance between two points in heterogeneous media. We demonstrate the practicality of these improvements by incorporating our methods into a production rendering system, demonstrating substantial noise reduction in a number of scenes with complex heterogeneous media, such as the one in Figure 1.

2 Problem Statement & Previous Work

The propagation of light through participating media is subject to interactions with particles. These interactions lead to absorption and scattering of photons and their rates are described by the absorption $\mu_a(x)$ and scattering $\mu_s(x)$ coefficients. The extinction coefficient $\mu(x) = \mu_a(x) + \mu_s(x)$ then denotes the total loss of light due to absorption and out-scattering per unit distance traveled. Table 1 summarizes the various coefficients that we use throughout the paper.

Consider a differential beam of light propagating through a participating medium with extinction coefficient $\mu(x)$. The fraction of



Figure 3: Illustration of different techniques for estimating transmittance through slabs of different materials. Regular tracking finds intersections with interfaces between individual materials. Ray marching proceeds with a constant step that needs to be small enough to avoid excessive bias. Delta tracking fills optically thinner regions with fictitious particles (red), analytically samples tentative free paths, and probabilistically decides whether collisions occur with real or fictitious particles.

Table 1: List of interaction coefficients used throughout the paper.

μ_a	absorption coefficient
μ_s	scattering coefficient
μ	extinction coefficient
$\bar{\mu}$	majorant extinction coefficient
μ_c	control extinction coefficient
μ_r	residual extinction coefficient
$\bar{\mu}_r$	majorant residual extinction coefficient

light T(d) that "survives" the transport up to distance d is:

$$T(d) = \frac{L_o}{L_i} = \exp\left(-\int_0^d \mu(x) \,\mathrm{d}x\right),\tag{1}$$

where L_i is the radiant energy at the beginning of the beam and L_o the amount that reaches distance d. The fraction is commonly referred to as the *transmittance*. In homogeneous media—where μ is spatially constant—Equation (1) simplifies to a simple exponential: $T(d) = \exp(-d\mu)$. In heterogeneous media, however, evaluating Equation (1) represents one of the major challenges of simulating radiative transport [Chandrasekhar 1960]. This is the problem we focus on in this paper.

In the following, we describe traditional approaches that are used to evaluate T(d). For brevity, we omit an extensive treatment of previous work on light transport simulation algorithms—and refer interested readers to standard literature [Cerezo et al. 2005; Pharr and Humphreys 2010]—as specific simulation approaches are largely orthogonal to our contributions.

2.1 Regular Tracking

The transmittance can be computed analytically when the *optical thickness*, i.e. the integral in Equation (1), has a known closed form anti-derivative. If the scene contains materials that fulfill this criterion (e.g. layered homogeneous slabs, discrete homogeneous voxels), we can use regular tracking [Amanatides and Woo 1987]—also known as the ray tracing method, substepping, or surface tracking in other fields—and split the beam at material interfaces. With homogeneous materials, the transmittance simplifies to:

$$T(d) = \prod_{i=1}^{N} \exp(-d_i \mu_i),$$
 (2)

where d_i is the length of *i*-th beam segment, which goes through medium with extinction coefficient μ_i (see Figure 3 for an illus-

Algorithm 1: *Pseudocode of the delta tracking estimator of transmittance along a ray with origin o, direction* ω *, and length d.*

DeltaTracking (o, ω, d)		
1	t = 0	
2	do:	
3	$\zeta = rand()$	
4	$t = t - \frac{\log(1-\zeta)}{\bar{a}}$	
5	if $t \geq d$: break	
6	$\xi = rand()$	
7	while $\xi > \frac{\mu(o+t*\omega)}{\bar{\mu}}$	
8	return t	
$\textbf{DeltaTrackingEstimator}(o, \omega, d)$		
9	$t = \text{DeltaTracking}(o, \omega, d)$	
10	return $t > d$	

tration). Since the properties of μ_i enable evaluating the integral analytically, the main cost of regular tracking resides in finding the interfaces, which is generally done by tracing rays. The drawback of the technique is that it cannot handle media with arbitrarily varying extinction coefficient that are fairly common in rendering.

2.2 Quadrature Methods

The integral in Equation (1) can also be evaluated numerically, e.g. using quadrature rules that are in this context commonly referred to as ray marching [Perlin and Hoffert 1989]. While this approach is sufficiently general, it unfortunately produces biased estimates of the transmittance. This holds even if the optical thickness is estimated in an unbiased manner (e.g. using MC integration [Pauly et al. 2000]) since $E[\exp(X)] \neq \exp(E[X])$; the exponentiation step "skews" the normal distribution of the error making its mean non-zero. The bias can be reduced by small marching steps, however, this is often too expensive and does not fit well into path tracing-based frameworks that rely on averaging many, relatively low quality samples.

2.3 Free-flight Sampling & Delta Tracking

The estimation of transmittance is highly related to random sampling of so-called free-flight distances between consecutive interactions with the medium. This requires solving for a distance d in Equation (1), which results in a given transmittance value, and can be done analytically in homogeneous media. Specialized approaches exist for certain continuously varying extinction functions [Carter et al. 1972; Brown and Martin 2003].

The general case, however, requires a technique like *delta track-ing*—also known as Woodcock tracking, pseudo scattering, hole tracking, or null-collision algorithms—which is based on von Neumann's [1951] rejection technique for generating numbers with arbitrary distributions. The technique was independently developed in neutron transport [Woodcock et al. 1965] and plasma physics [Skullerud 1968] for unbiased sampling of neutron and ion free paths, respectively, in environments with many materials. It has been later formalized by Coleman [1968] and recently presented in an integral form by Galtier et al. [2013]. Raab et al. [2008] introduced delta tracking to graphics for rendering participating media. As this technique forms the basis for our new estimators, we introduce it in greater detail.

The idea of delta tracking is to "homogenize" the heterogeneous medium by adding *fictitious* particles. The local concentration of fictitious particles is set so that the combined extinction coefficient $\bar{\mu}$,

often referred to as the *majorant*, is spatially homogeneous.¹ The albedo and phase function of fictitious particles are set to 1 and $\delta(\omega)$, respectively. As such, photons interacting with fictitious particles continue unaltered along the original direction. These interactions are in some physics literature referred to as "null collisions".

In practice, delta tracking models the interactions with real and fictitious particles probabilistically. The technique essentially constructs a random walk along a line, whose Euclidean length represents a sample of the free flight distance. As shown in Algorithm 1, the tracking employs the majorant extinction $\bar{\mu}$ to sample a distance tto a "tentative" interaction point (line 4). Then, if t is not beyond a given maximum distance d (e.g. to the nearest surface), the algorithm probabilistically classifies the interaction as either real or fictitious: if a random number ξ is greater than the relative concentration of real particles, i.e. $\xi > \mu(x)/\bar{\mu}$ (line 7), the interaction is said to involve a fictitious particle. The random walk then continues by repeating the process of sampling free flight distances and probabilistically classifying the tentative interactions until a real collision occurs. Coleman [1968] proved that free flight distances generated in this manner have the desired distribution; the algorithm is thus unbiased.

The cost of delta tracking is highly dependent on how tightly the majorant $\bar{\mu}$ bounds the true extinction coefficient $\mu(x)$ as this directly impacts the number of rejected, tentative interactions. Yue et al. [2010] and Szirmay-Kalos et al. [2011] both suggested strategies to optimize this process by localizing the majorant calculation to only bound the extinction coefficient locally. As we show later, these optimizations are both orthogonal and complementary to our contributions for computing transmittance. It is also worth noting that certain variants of delta tracking can handle non-bounding "majorants" [Carter et al. 1972; Galtier et al. 2013].

Free-flight Sampling as a Transmittance Estimate. Unbiased free-flight sampling routines can also be used for estimating transmittance [Raab et al. 2008; Szirmay-Kalos et al. 2011; Jarosz et al. 2011]: if the sampled free flight distance is greater than *d*, the transmittance is estimated as 1, and 0 otherwise. The binary estimate can be further refined by averaging multiple instances, i.e. by counting the relative number of free flight samples that exceed *d*.

Figures 4b (top) and 4c (top) show two instances of delta tracking in the same medium but with two different majorants (see Figure 4a). The trackings differ in the number of tentative steps (represented by blue circles and tics on the x-axis) they generate. With a tight majorant, the collision with a real particles occurs much faster than with a loose majorant. In either case, the transmittance function is approximated by a step function. Figures 4d and 4e show the effect of averaging multiple free flight samples yielding finer approximations (black curves) of the transmittance function (red curves).

3 The New Estimators of Transmittance

In this section, we introduce two new complementary techniques for estimating transmittance. We strive for an intuitive description here; precise definitions and proofs of unbiasedness are included in the supplementary material. The goal for the first estimator, referred to as *ratio tracking*, is to leverage the information discovered during the tracking more efficiently instead of deducing "just" a binary answer. The resulting estimator provides a *piecewise constant* approximation to the transmittance function. The second technique, called *residual tracking*, is complementary to delta tracking and ratio tracking and combines numerical estimation with an analytic approximation, yielding a *piecewise exponential* solution.

¹The majorant can in fact vary spatially as long as it enables tractable, closed-form inversions of optical thickness for sampling free flight distances.



Figure 4: A comparison of delta and ratio tracking-based estimators in a periodic heterogeneous medium bound by a loose and tight majorant extinction coefficient (**a**). In (**b**, **c**), we show single instances (blue) of delta (top) and ratio (bottom) trackings for the two different majorants. In (**d**, **e**), we show approximations of the transmittance function (black) obtained by averaging 100 instances (thin blue) of the corresponding tracker. Red curves represent the ground truth transmittance function. The ticks on horizontal axes mark all sampled (tentative) collision points.

3.1 Ratio Tracking

The delta tracking-based estimator of transmittance can be interpreted as a random walk terminated by Russian roulette. The termination probability at each tentative collision point is set to the local fraction of real particles. In Monte Carlo simulations, Russian roulette often serves as an alternative to discrete or continuous weighting [Hayakawa et al. 2014], where, instead of probabilistically terminating the path, we continue constructing it until a boundary condition is met while weighting down its contribution.

Our ratio tracking approach follows the same rationale: we replace the Russian roulette by a weight that is equal to the probability of colliding with a fictitious particle. Instead of probabilistically terminating the random walk at one of tentative collision points, we thus continue it until we reach the end-point of the ray. During the construction we keep track of the joint probability of colliding with fictitious particles at all the preceding tentative collisions points. This becomes the "weight" that the estimator scores when reaching d:

$$\langle T(d) \rangle_{\mathrm{R}} = \prod_{i=1}^{K} \left(1 - \frac{\mu(x_i)}{\bar{\mu}} \right).$$
(3)

The multiplicand in the product represents the ratio of fictitious to all particles at collision point x_i . Note that in addition to these ratios, the value of the estimator depends also on the number of tentative collision points K that it generates before reaching d. This is a random variable that, in addition to d, depends on $\overline{\mu}$. Its expected value $E[K] = \overline{\mu}d$ determines the cost of the algorithm.

Algorithm 2: *Pseudocode of the ratio tracking estimator of transmittance along a ray with origin o, direction* ω *, and length d.*

 $\textbf{RatioTrackingEstimator}(o, \omega, d)$

t = 01 T = 12 3 do: $\zeta = rand()$ 4 $t = t - \frac{\log(1-\zeta)}{\zeta}$ 5 $t = t - \frac{1}{\bar{\mu}}$ if $t \ge d$: break $T = T * \left(1 - \frac{\mu(o + t * \omega)}{\bar{\mu}}\right)$ 6 7 while true 8 return T9



(a) Delta tracking (b) Ratio tracking

(c) Reference

Figure 5: *Transmittance estimated by averaging 2 instances of delta* (a) *and ratio* (b) *tracking. While the non-analog, ratio tracking estimator reduces variance more efficiently, it is also more expensive. See Section 4 for a comparison of the cost and effective variance.*

Algorithm 2 provides pseudocode for the ratio tracking estimator. The correctness of the method should intuitively follow from the fact that we simply replace Russian roulette by another unbiased technique: the weighting that occurs at (discrete) tentative collision points. To support this intuition in a mathematically rigorous way, we include a formal proof of correctness in the supplementary material. The bottom row of Figure 4 shows the approximations of the transmittance produced by single instances of ratio tracking (blue curves), as well as averages over 100 trackings (black curves). Note how the approximations drop at each tentative collision point proportionally to the relative concentration of fictitious particles.

Discussion. Delta tracking belongs to the category of *analog* estimators whose scores strictly adhere to real physical processes being simulated. The estimator simulates the process of attenuation by considering interactions with real particles only: it scores 1 if only fictitious interactions occurred along the ray, and 0 otherwise. In contrast, our ratio tracking estimator adjusts the score at every tentative collision, including fictitious ones. It thus belongs to the class of *non-analog* estimators that in certain situations yield lower variance [Hykes and Densmore 2009; Hayakawa et al. 2014]; we demonstrate this in Figure 5. However, as ratio tracking needs to unconditionally reach the endpoint of the ray, the cost of the estimation—in comparison to delta tracking—may increase. To better understand these trade-offs, we analyze the variance, cost, and the net efficiency of both estimators in Section 4.



Figure 6: A comparison of residual tracking with different control extinction coefficients. In (**b**) and (**c**), we analytically compute the control transmittance (purple) based on the minimum $\mu(x)$ along the ray and then apply delta tracking (**b**) and ratio tracking (**c**) to numerically evaluate the transmittance through the residual medium (blue curves: individual trackings, black curves: averages). The product of the control and the residual transmittance is represented by the green curve. Ratio tracking can be used with arbitrary control extinctions: in (**d**) and (**e**), we show examples with the average and the maximum $\mu(x)$ used as the control extinction.

3.2 Residual Tracking

The previously introduced ratio tracking becomes expensive when E[K], i.e. the expected number of steps required to reach the endpoint of the ray, is high. This directly depends on the value of the majorant extinction coefficient. Our goal in this section is to reduce the value of the majorant. We first evaluate part of the transmittance analytically, and then use the numerical tracking only on the remaining (residual) extinction. As long as the residual extinction has lower majorant than μ , the ratio tracking—and under some constraints also delta tracking—will advance using longer tentative free-flight distances, which in turn lowers E[K], thereby reducing the cost of the estimation.

We start by introducing the *control* extinction coefficient $\mu_c(x)$, which is a simplified version of the original $\mu(x)$ that enables expressing the optical thickness up to distance d in closed form $\tau_c(d)$. The transmittance up to d can then be written as:

$$T(d) = \exp\left(-\int_{0}^{d} \mu(x) \, \mathrm{d}x\right)$$

= $\exp\left(-\int_{0}^{d} \mu_{c}(x) + \mu(x) - \mu_{c}(x) \, \mathrm{d}x\right)$
= $\exp\left(-\tau_{c}(d)\right)$
Control transmittance $\exp\left(-\int_{0}^{d} \mu(x) - \mu_{c}(x) \, \mathrm{d}x\right)$. (4)
Residual transmittance

The first exponential term, referred to as the control transmittance $T_c(d)$, represents a coarse approximation of T(d), which is computed using the simplified extinction coefficient μ_c . The second exponential then serves as a correction that accounts for the difference between the control and the actual transmittance. We denote this exponential as the *residual* transmittance T_r and refer to the integrand as the residual extinction coefficient $\mu_r(x) = \mu(x) - \mu_c(x)$. Please note that in certain situations—depending on the value of $\mu_c(x)$ —the residual extinction may be negative.



Figure 7: We apply the concept of control variates and split the extinction function (\mathbf{a}) to control (\mathbf{b}) and residual (\mathbf{c}) extinction. The transmittance of (\mathbf{a}) can be then computed as the product of analytically evaluated control transmittance and numerically estimated residual transmittance.

Figure 7 illustrates the concept of decomposing the extinction function into the control and the residual part. The transmittance through the residual part can be computed using any of the techniques described up to now; however, we will restrict ourselves to the general and unbiased estimation via ratio tracking and, for some values of $\mu_c(x)$, also delta tracking. For brevity, we drop the positional parameter and simply write μ_c assuming a constant control extinction. It should be noted, however, that using spatially varying control extinction may further decrease the variance (see Section 6).

The idea of residual tracking resembles the concept of control variates. The differences are subtle, but it is worth noting that since we are evaluating exponentiated integrals, even a *constant* control variate (i.e. constant control extinction in our case) can reduce variance if it decreases the rate of change of the residual exponentiated integral. The residual tracking is also related to the "separation of main part" applied by Szirmay-Kalos et al. [2011] to the evaluation of optical thickness. However, their approach will be effective only if the main part is represented by a non-constant function that matches the extinction function well. In contrast, our residual tracking benefits even from constant control extinctions, which are easy to compute. Furthermore, Szirmay-Kalos et al. proposed to integrate the residual optical thickness using numerical quadratures, thereby biasing the result. Our techniques, that we describe next, remain unbiased.



Figure 8: Residual ratio tracking in a simple procedural volume with the control transmittance (\mathbf{a}) computed using the minimum (top), the average (middle), and the maximum (bottom) extinction coefficient in the volume. The residual transmittance (\mathbf{b}) used to compute the product (\mathbf{c}) was estimated using 4 trackings only to

emphasize the noise typical for each control extinction.

3.2.1 Residual Delta Tracking

In order to estimate the residual transmittance with delta tracking, we have to ensure that the residual extinction coefficient $\mu_r(x)$ is always non-negative, i.e. μ_c must never exceed $\mu(x)$. Otherwise the probabilities of colliding with real particles become negative. Since delta tracking has no means to handle such situations properly, the estimation would not converge to the correct result. Note that this is normally not a problem as media with negative densities do not exist in the real world. In the case of residual tracking, however, this limitation significantly restricts possible values of μ_c , thereby reducing the effectiveness of using the control variate.

Figure 6b shows a log-plot featuring the control transmittance (purple curve), an estimate of residual transmittance obtained by averaging 100 instances of residual delta tracking, and their product (green curve) that approximates the true transmittance function. To prevent negative probability densities of colliding with real particles, we set the control extinction μ_c to the *minimum* extinction coefficient μ_{min} along the ray. For sampling the free-flight distances, we used the majorant of the residual extinction coefficient $\bar{\mu}_r = \mu_{max} - \mu_{min}$.

3.2.2 Residual Ratio Tracking

By adjusting Equation (3) to estimate only the residual transmittance, $T_r(d)$, we obtain the residual ratio tracking estimator:

$$\langle T_r(d) \rangle_{\rm RR} = \prod_{i=1}^K \left(1 - \frac{\mu_r(x_i)}{\bar{\mu}_r} \right) = \prod_{i=1}^K \left(1 - \frac{\mu(x_i) - \mu_c}{\bar{\mu}_r} \right).$$
 (5)

The value of $\bar{\mu}_r$ is in this case defined slightly differently. As stated at the beginning of Section 3.2, our goal is to minimize the tracking cost by prolonging the tentative free-flight distances. This amounts to finding the smallest $\bar{\mu}_r$ that still avoids *negative multiplicands*, which would prevent convergence of the estimator in Equation (5). To ensure that $\mu_r(x)/\bar{\mu}_r \leq 1$, we use $\bar{\mu}_r = \max(|\mu_r(x)|; 0 \leq x \leq d)$, i.e. the maximum absolute difference between $\mu(x)$ and μ_c along the ray.

Algorithm 3: *Pseudocode of the residual ratio tracking estimator for sampling transmittance along a ray with origin o, direction* ω *, and length d.*

ResidualRatioTrackingEstimator (o, ω, d)

1 t = 0 $T_c = \exp\left(-\mu_c * d\right)$ 2 $T_r = 1$ 3 do 4 $\zeta = rand()$ 5 $\begin{aligned} & \zeta = \operatorname{Iund}() \\ & t = t - \frac{\log(1-\zeta)}{\bar{\mu}_r} \\ & \text{if } t \ge d: \text{ break} \\ & T_r = T_r * \left(1 - \frac{\mu(o+t*\omega) - \mu_c}{\bar{\mu}_r}\right) \end{aligned}$ 6 7 8 while true 9 return $T_c * T_r$ 10

An important advantage of ratio tracking over delta tracking is that it can handle *negative* extinction coefficients. To see how, we need to refrain from the previously introduced physical meaning of the multiplicand—the ratio of fictitious to all particles—and simply consider it as a local weight. By inspecting Equation (5), we see that the weights at collisions in regions of negative residual extinction are greater than 1. As such, these regions can be interpreted as light amplifiers (as opposed to attenuators) that boost the residual transmittance to correct the underestimating control transmittance.

In the following, we examine three different values of μ_c :

- if μ_c = μ_{min} = min(μ(x); 0 ≤ x ≤ d), i.e. the control extinction is underestimating, the control transmittance systematically overestimates the real transmittance (i.e. the purple curve in Figure 6c is higher than the red curve) and the residual transmittance thus needs to scale it down.
- if µ_c = µ_{avg} = avg(µ(x); 0 ≤ x ≤ d), the control extinction along the ray matches the real extinction on average. The residual tracking thus corrects only for the local over- and underestimation of the control transmittance w.r.t. the real transmittance along the ray. This can be seen as a wiggle of the black curve around value 1 in Figure 6d.
- if μ_c = μ_{max} = max(μ(x); 0 ≤ x ≤ d), the control transmittance systematically underestimates the real transmittance (see Figure 6e) and the residual tracking thus needs to produce values greater than 1 to scale the control up.

Figure 8 visualizes the control and the residual transmittance where the control extinction was set to the minimum, average, and the maximum extinction inside the volume. The bright regions correspond to cases when the residual transmittance takes on values higher than 1 to correct for the overly low control transmittance. Algorithm 3 provides the pseudocode of the estimator. Please refer to the supplementary material for a proof of unbiasedness.

Discussion. Similarly to ratio tracking, the residual tracking is a non-analog estimator. Furthermore, it combines the properties of *continuous* and *discrete* weighting estimators [Hayakawa et al. 2014]: the control transmittance provides a continuous, exponential approximation, which is then adjusted using a piecewise-constant correction term obtained from the residual (ratio) tracking. The resulting approximation of transmittance is thus piecewise exponential. Note that when the medium is homogeneous (as in Figure 2), the control transmittance is already exact and the result thus noise-free. Residual tracking can also be used for creating "medium-length beams" proposed by Křivánek et al. [2014].



Figure 9: An analysis of a canonical scene (**a**) rendered in (**b**) with delta tracking (top) and ratio tracking (bottom) estimators using roughly the same number of μ evaluations. Visualizations in each subsequent column use the same false-color scale. The vertical and the horizontal axes represent the collision sampling efficiency η and the extinction μ along the ray, respectively. The variance of single instances of delta tracking (**c**, top) is agnostic to η , however, the tracker becomes more expensive with low values of η (**d**, top). The increase in cost in such situations is even more apparent with ratio tracking (**d**, bottom), however, the estimator is able to compensate for this by the lower variance (**c**, bottom). We show the products of the variance and the cost (i.e. the effective variance) in (**e**). The absolute difference between the two products is shown in (**f**), where the white line represents zero difference and the labels mark the estimator that performs better in the region.

4 Analysis of Efficiency

In this section, we analyze the variance, the cost, and their product for the aforementioned tracking estimators.

4.1 Delta vs. Ratio Tracking

In Figure 4, we visualized the delta and ratio trackings for a tight and a loose value of the majorant extinction coefficient. The loose majorant generally leads to a higher number of tentative collision points, thus increasing the cost of both trackings. However, the resulting ratio-tracking approximation consists of more steps that match the true transmittance better, thereby reducing the variance.

The goal of this section is to analyze how the variance and the cost of the two trackers depend on the value of the majorant and the optical thickness of the medium. We use the collision sampling *efficiency* [Leppänen 2010], defined as $\eta(x) = \mu(x)/\bar{\mu}$, to express how "tightly" the majorant bounds the extinction function at x; high values correspond to tight bounding and high probabilities of real interactions. For the comparison we consider a canonical scene consisting of an axis-aligned unit cube filled with an absorbing medium, an orthographic camera, and an area light source, placed behind the cube; see Figure 9a for an illustration. To study the dependency on the optical thickness, we exponentially increase the extinction coefficient of the medium along the x axis so that the transmittance (and thus the intensity of pixels) decreases linearly as we move from left to right across the rendered image (see Figure 9b). The density of the medium is kept constant along y and z axes. To incorporate the dependency on the majorant, we modulate the collision sampling efficiency vertically: pixels at the bottom are rendered with low η (loose majorants) while pixels on top are computed with η approaching 1 (i.e. almost perfectly tight majorants).

Figure 9c visualizes the variance of single instances of the delta tracking (top) and the ratio tracking (bottom) estimators. In the case of delta tracking, the variance does not depend on η . Since the estimator always scores 0 or 1 independent of the number of tentative collisions, its variance is agnostic to the value of the majorant. In contrast, the ratio tracking-based estimator always scores a real

number leading to a lower overall variance. More importantly, its variance reduces with decreasing collision sampling efficiency. This is because ratio tracking—in contrast to delta tracking—can leverage the higher number of tentative collision points to produce a finer approximation of the transmittance function.

Figure 9d depicts the average number of generated tentative collision points per single tracking; this we consider to be the the cost of the estimation. The cost of both estimators increases with lower values of the collision sampling efficiency. Unsurprisingly, ratio tracking is also more expensive than delta tracking as it always constructs the random walk all the way to the back side of the unit cube. This becomes most apparent in the right part of the visualization, where the medium is optically thick and mean free paths relatively short.

In order to compare the net efficiency, we plot the effective variance, i.e. the products of the per-tracker variance and cost, in Figure 9e. High values in these plots correspond to configurations where either the cost, the variance, or both are high; i.e. the tracker has difficulties reducing the noise in these configurations. In Figure 9f, we visualize the absolute difference of the effective variance plots. The white curve represents configurations where both estimators perform equally well. For high values of η , the delta tracking estimator performs marginally better. However, the difference between the two estimators becomes much more significant for low values of η , where the piecewise constant approximation produced by ratio tracking outperforms the binary estimation, despite the higher cost per single instance. This can be seen in Figure 9b that shows equal-cost renderings of the volumetric cube. While the noise in the top part of the renderings is comparable, the bottom part looks significantly better when using ratio tracking.

For best performance, one would want to choose the estimator based on the transmittance and collision sampling efficiency along the ray. However, as neither of these is known a-priori, the selection can be done only heuristically. We also experimented with different on-line switching schemes, but since the relative improvements of delta tracking with high η are rather low, the returns were diminishing. We thus always start with ratio tracking and switch to delta tracking only if the transmittance drops below 0.1%. It is also worth noting that low values of η are fairly common in practice.



Figure 10: The per-tracker variance, cost, and their corresponding product for different extinction functions. Note that independently of the optical thickness, which changes along the horizontal axis of the false-color plots, the product of the variance and the cost is minimized quite well by using the average extinction coefficient along the ray as μ_c . Please refer to the supplementary material for more examples.

4.2 Residual Ratio Tracking with Various Values of μ_c

In this section, we compare how different values of the control extinction coefficient impact the efficiency of the residual ratio tracking. We use the same geometric setup as in Section 4.1, however, this time we also vary the density of the medium along the z-axis. This enables studying how the shape of the extinction function impacts the efficiency of the estimators. Columns (a) and (e) of Figure 10 show four extinction functions that are scaled to produce the desired transmittance along the horizontal axis (as in Figure 9). The vertical axes of all false-color plots are identical and represent the control extinction coefficient (precisely its relative value w.r.t. the maximum, i.e. μ_c/μ_{max}). The three dashed lines correspond to μ_{min} , μ_{avg} , and μ_{\max} , i.e. the minimum, average, and maximum $\mu(x)$ along the ray, respectively. While it is generally hard to classify how the variance and cost depend on μ_c , the minimum of their product (columns (d) and (h)) is in all cases (and for all values of transmittance) very close to μ_{avg} . As such, using the average extinction coefficient as μ_c will be effective in reducing the stochastic error of estimating the residual transport, which in that case averages to value 1. Note that μ_{avg} represents the average *along the ray* and computing it for every ray independently is not practical. We address this in Section 5.1.

5 Practical Rendering Algorithm

The primary motivation for developing the new transmittance estimators was to meet the criteria of production path tracers. These create a large number paths for every pixel whose (weighted) average contribution estimates the pixel color. The transmittance thus needs to be evaluated relatively cheaply along each path, but without systematic errors to ensure correct and predictable results on average.

In cases when the medium is highly heterogeneous and relatively large w.r.t. the mean free path, maintaining a single μ_c for the entire volume is sub-optimal. Fortunately, our technique can be combined with spatial indexing structures, e.g. super-voxels [Szirmay-Kalos et al. 2011] or kd-trees [Yue et al. 2010], that provide localized majorants. Our implementation uses the super-voxels idea by Szirmay-Kalos et al. [2011] but stores localized control and majorant residual extinction coefficients; we detail their computation in Section 5.1. The other difference to the original super-voxels is that our representation is sparse and hierarchical, leveraging the benefits of the VDB data structure [Museth 2013] for compact storage and fast traversal. The size of super-voxels is set to $10 \times$ the size of the smallest representable detail. During rendering, we split each ray into segments—one for each intersected super-voxel—and estimate the transmittance in each segment using residual ratio tracking.



(a) Supergrid

(b) Average $(\mu_c = \mu_{avg})$ **(c)** Heuristic $(\mu_c = \mu'_c)$

Figure 11: When a ray passes just next to a medium, using the average extinction within super-voxels as μ_c may heavily underestimate the control transmittance and make the estimation of residual transmittance prone to high variance (**b**). Our heuristic (**c**) detects such situations and conservatively decreases the control extinction.

5.1 Extinction Parameters of Super-Voxels

The super-voxels aggregate extinction parameters of all heterogeneous volumes. The analysis in Section 4.2 suggests that setting μ_c to the average extinction along a ray (segment) yields good performance. Since finding the exact average along the ray segment is impractical, we could instead approximate it with the average extinction inside the super-voxel. This works well if the mean gradient of the extinction function in the super-voxel is rather low. Figure 11 illustrates a failure case where a ray travels through super-voxels that only partially overlap dense regions, and the average extinction of those voxels heavily overestimates the actual extinction along the ray. Consequently, the control transmittance severely underestimates the true value and the residual tracking is prone to high variance.

To address this issue, we decrease the control extinction coefficient in troublesome super-voxels to a heuristically derived value μ'_c , so that the expected value of the residual transmittance, approximated as:

$$\langle T_r(d) \rangle \approx \left[1 - \frac{\mu_{\text{avg}} - \mu_c}{\bar{\mu}_r} \right]^{d\bar{\mu}_r},$$
 (6)

is always below a user-defined threshold γ . Here, μ_{avg} represents the average extinction along a ray segment of length d. To derive μ'_c , we assume a hypothetical worst case with the highest possible residual transmittance: i.e. the ray travels along the super-voxel's diagonal of length D, the extinction along the ray equals the minimum extinction μ_{\min} found in the super-voxel, and $\bar{\mu}_r = \mu_{\max} - \mu_{\min}$. To compute μ'_c , we use γ as the expectation of the residual transmittance, solve Equation (6) for μ_c : $\mu_c = \mu_{\min} + \bar{\mu}_r (\gamma \frac{\bar{D} \bar{\mu}_r}{\bar{D} \bar{\mu}_r} - 1)$, and evaluate it for the user-defined threshold; we always use $\gamma = 2$. We further restrict μ'_c to be within $\langle \mu_{\min}, \mu_{avg} \rangle$, compute the majorant of the corresponding residual extinction coefficient $\bar{\mu}_r$ (see Section 3.2.2), and store these two coefficients in each super voxel.

5.2 Rendering

In order to estimate the transmittance along a given ray, we perform regular tracking through the super-voxel grid (i.e. 3D DDA [Amanatides and Woo 1987]). For each intersected super-voxel, we identify the overlapping ray segment and query the local μ_c and $\bar{\mu}_r$. The transmittance along the segment is computed by the product of control transmittance and residual transmittance estimated using ratio tracking. The transmittance along the entire ray is then calculated by multiplying the transmittances of all ray segments.

While performing the residual ratio tracking, we temporarily store a functional representation of the transmittance along the entire ray. This is then used for constructing a numerical PDF for importance sampling the in-scattered light along the ray (our PDFs are proportional to the product of transmittance, scattering, and fluence), and for estimating the transmittance to locations where we sample the in-scattered and/or emitted light.

6 Results

We implemented both estimators in the Mitsuba renderer [Jakob 2010] and an in-house production renderer. Figure 12 shows unbiased renderings of a rising smoke plume—we purposely made it absorptive to avoid variance due to scattering—computed by the traditional delta tracking estimator, residual delta tracking estimator, and four variants of our residual ratio tracking estimator. The bottom row visualizes the variance of each estimator. We use super-voxels to store localized control and majorant extinction coefficients. The residual transmittance corrects only in regions where the control transmittance underestimates or overestimates the true value. To visualize the noise, we use rather low per-pixel sample counts that are adjusted to produce a roughly equal number of $\mu(x)$ evaluations; this in practice determines the performance of the algorithm.

While setting $\mu_c = \mu_{avg}$ yields residual transport that is on average closest to 1 among all the tested variants, the product may suffer from high variance in configurations outlined in Section 5.1. Using the heuristically derived μ'_c avoids these issues while preserving the RMSE obtained with μ_{avg} , which is about $2\times$ lower than with the traditional delta tracking-based estimator. A path tracer leveraging the residual ratio tracking thus requires $4\times$ fewer samples to resolve transmittance at the quality obtained with delta tracking. It is worth noting that not all values of μ_c reduce the error equally well, e.g. using $\mu_c = \mu_{max}$ yields significantly higher variance than $\mu_c = \mu_{min}$. It is thus important to estimate the average extinction along the ray accurately, or make μ_c rather underestimate the optimal value.

Figure 13 demonstrates the benefits of using trilinearly interpolated control extinction; we use the same polynomial approximation as Szirmay-Kalos et al. [2011] but apply it to μ_c instead of the majorant. Since the trilinearly interpolated control matches $\mu(x)$ better, the residual transport does not deviate from 1 as much and can be estimated with lower variance. If we make the control extinction continuous, e.g. by using the same μ_c for adjacent super-voxel corners, the noise will change gradually over the image plane; this may be in certain scenarios advantageous.

6.1 Media with Colored Extinction

Figure 1 shows a rendering obtained with our in-house production renderer. The clouds use "colored" extinction coefficients. This poses a problem for delta tracking, which, in order to preserve high collision sampling efficiency, needs to handle the transmittance computation for each color channel independently. The estimation becomes easier with our ratio tracking that maintains good effective performance even with low values of η ; all color channels can thus be efficiently handled at once. For the control transmittance, we computed μ_c coefficients separately for each color and then estimated the residual transmittance by a single instance of residual ratio tracking, which uses the maximum $\bar{\mu}_{\tau}$ across all color channels.

The insets emphasize the noise obtained with delta, ratio, and residual ratio tracking with roughly the same number of $\mu(x)$ evaluations. In the top insets we disabled scattering to visualize the noise due to the transmittance estimation only. As shown in the bottom insets, the different performance of the two estimators is clearly visible even when adding noisy estimates of multiple scattering; all renderings use the same PDFs to sample in-scattering, the only difference resides in the transmittance estimation. The residual ratio tracking yields lower RMSE producing results with roughly $6 \times$ lower effective variance (i.e. the product of MSE and cost) when considering just the transmittance, and about $2.3 \times$ lower effective variance when simulating multiple scattering in the medium. Figure 14 shows two more examples with scattering media.

7 Discussion and Future Work

Overall Impact of Transmittance Estimation. Transmittance estimation is only one of the many components of evaluating radiative-transport integrals. Here, we primarily focused on studying the variance of transmittance estimators in isolation using metrics that are renderer independent. In general, the reported improvements should lead to a corresponding decrease of noise in memory bound scenes where transmittance significantly impacts the quality.

Non-strict Majorants. The residual delta tracking and residual ratio tracking estimators require strictly bounding majorants to avoid negative probabilities or negative multiplicands, respectively. Carter et al. [1972] and Galtier et al. [2013] proposed weighted variants of delta tracking that can deal with non-bounding majorants (i.e. negative densities of fictitious particles). These methods are quite practical as they remove the burden of finding bounding majorants, be it at the cost of increased variance. While employing such weighting in residual delta tracking is trivial, applications to residual ratio tracking require careful variance analysis; we leave it as future work.

Free-flight Sampling. One could also leverage the concept of control variates and residual tracking for free-path sampling. Assuming that the control extinction represents a good approximation of $\mu(x)$ and the control optical thickness is invertible, free paths could be sampled using the inversion method. The sample would then be further weighted by the residual transmittance, which can be efficiently estimated for all channels via the residual ratio tracking.

Integral Formulation. We would like to note that ratio tracking can be seen as a special case of the recently presented integral formulation of delta tracking [Galtier et al. 2013]. While we did not leverage this framework here, we believe it could offer a foundation for deriving and proving new, weighted trackings, which may under certain constraints be more efficient than the original algorithm.

8 Conclusion

We presented two complementary concepts that yield unbiased estimators for efficient evaluation of transmittance in heterogeneous volumes. In comparison to delta tracking, ratio tracking decreases the need for having tight majorants via improved efficiency in those cases. When combined with the control transmittance, the residual ratio tracking estimator gracefully handles media with low degree of heterogeneity, simplifying to an analytic solution for homogeneous (sub)volumes. Furthermore, the estimators do not require any manual tuning (such as the step size in ray marching), and can be easily



Figure 12: Equal-cost renderings of a rising smoke with delta tracking (**a**), residual delta tracking (**b**), and residual ratio tracking with different control extinctions coefficients μ_c (**c-f**). All estimators were accelerated using super-voxels with constant μ_c . The residual transmittance (second row) was estimated using a roughly equal number of $\mu(x)$ evaluations, see the label for the exact value. The estimators with $\mu_c \in {\mu'_c, \mu_{avg}}$ in columns (**d**) and (**e**) yield the lowest error, which is visualized in the bottom row as equally scaled, false-colored variance.



Figure 13: A comparison of different control extinction functions: control transmittance ($\mathbf{a}, \mathbf{b}, \mathbf{c}$) and residual transmittance ($\mathbf{d}, \mathbf{e}, \mathbf{f}$) for constant ($\mu_c = \mu'_c$) and trilinearly interpolated continuous and discontinuous control extinction, respectively. The plot in (\mathbf{g}) shows the relative effective variance of trilinearly interpolated controls w.r.t. the constant control; all plotted as functions of the resolution of the super-grid.



Figure 14: Two scenes from Big Hero 6 rendered using residual ratio tracking for estimating transmittance. The equal-cost insets show renders of the medium only and were not color graded (hence the difference to the final images above) to allow for a better comparison of the noise obtained with traditional delta tracking estimator and our residual ratio tracking estimator. Images ©Disney.

combined with existing spatial indexing structures. Given that the code changes w.r.t. delta tracking are minimal, integrating them into renderers that already use delta tracking is trivial and well justified by the more effective handling of media with colored extinction.

9 Acknowledgments

We thank reviewers for their helpful comments, Simon Kallweit for implementing the interpolated control extinction, and Oliver Klehm and Jaroslav Křivánek for proofreading. The bunny (Figure 2) and the smoke (Figures 5, 13) are courtesy of the Stanford Computer Graphics Laboratory and the OpenVDB library, respectively. Figures 1 and 14 were rendered using Disney's Hyperion Renderer.

References

- AMANATIDES, J., AND WOO, A. 1987. A fast voxel traversal algorithm for ray tracing. In *Eurographics* '87, 3–10.
- BROWN, F. B., AND MARTIN, W. R. 2003. Direct sampling of Monte Carlo flight paths in media with continuously varying cross-sections. In *Proc. of ANS Mathematics & Computation Topical Meeting*, 6–11.
- CARTER, L. L., CASHWELL, E. D., AND TAYLOR, W. M. 1972. Monte Carlo sampling with continuously varying cross sections along flight paths. *Nuclear Science and Engineering* 48, 4, 403– 411.
- CEREZO, E., PÉREZ, F., PUEYO, X., SERON, F. J., AND SILLION, F. X. 2005. A survey on participating media rendering techniques. *The Visual Computer 21*, 5, 303–328.
- CHANDRASEKHAR, S. 1960. *Radiative Transfer*. Dover Publications.
- COLEMAN, W. A. 1968. Mathematical verification of a certain Monte Carlo sampling technique and applications of the technique to radiation transport problems. *Nuclear Science and Engineering* 32, 1 (Apr.), 76–81.
- DACHSBACHER, C., KŘIVÁNEK, J., HAŠAN, M., ARBREE, A., WALTER, B., AND NOVÁK, J. 2013. Scalable realistic rendering with many-light methods. *Computer Graphics Forum 33*, 1, 88– 104.
- GALTIER, M., BLANCO, S., CALIOT, C., COUSTET, C., DAUCHET, J., HAFI, M. E., EYMET, V., FOURNIER, R., GAU-TRAIS, J., KHUONG, A., PIAUD, B., AND TERRE, G. 2013. Integral formulation of null-collision Monte Carlo algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer 125* (Apr.), 57–68.
- GEORGIEV, I., KŘIVÁNEK, J., HACHISUKA, T., NOWROUZEZAHRAI, D., AND JAROSZ, W. 2013. Joint importance sampling of low-order volumetric scattering. *ACM TOG (Proc. of SIGGRAPH Asia) 32*, 6 (Nov.), 164:1–164:14.
- HAYAKAWA, C. K., SPANIER, J., AND VENUGOPALAN, V. 2014. Comparative analysis of discrete and continuous absorption weighting estimators used in Monte Carlo simulations of radiative transport in turbid media. J. Opt. Soc. Am. A 31, 2, 301–311.
- HYKES, J. M., AND DENSMORE, J. D. 2009. Non-analog Monte Carlo estimators for radiation momentum deposition. *Journal* of Quantitative Spectroscopy and Radiative Transfer 110, 13, 1097–1110.
- JAKOB, W., 2010. Mitsuba renderer. http://www.mitsubarenderer.org.

- JAROSZ, W., NOWROUZEZAHRAI, D., THOMAS, R., SLOAN, P.-P., AND ZWICKER, M. 2011. Progressive photon beams. ACM TOG (Proc. of SIGGRAPH Asia) 30, 6 (Dec.), 181:1–181:12.
- KAJIYA, J. T. 1986. The rendering equation. *Computer Graphics* (*Proc. of SIGGRAPH*), 143–150.
- KELLER, A. 1997. Instant radiosity. In *Proc. of SIGGRAPH* 97, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, Annual Conference Series, 49–56.
- KULLA, C., AND FAJARDO, M. 2012. Importance sampling techniques for path tracing in participating media. CGF (Proc. of Eurographics Symposium on Rendering) 31, 4 (June), 1519–1528.
- KŘIVÁNEK, J., GEORGIEV, I., HACHISUKA, T., VÉVODA, P., ŠIK, M., NOWROUZEZAHRAI, D., AND JAROSZ, W. 2014. Unifying points, beams, and paths in volumetric light transport simulation. ACM TOG (Proc. of SIGGRAPH) 33, 4, 103:1–103:13.
- LAFORTUNE, E. P., AND WILLEMS, Y. D. 1993. Bi-directional path tracing. In *Compugraphics* '93, 145–153.
- LEPPÄNEN, J. 2010. Performance of Woodcock delta-tracking in lattice physics applications using the Serpent Monte Carlo reactor physics burnup calculation code. *Annals of Nuclear Energy* 37, 5, 715 – 722.
- MUSETH, K. 2013. VDB: High-resolution sparse volumes with dynamic topology. *ACM TOG 32*, 3 (July), 27:1–27:22.
- PAULY, M., KOLLIG, T., AND KELLER, A. 2000. Metropolis light transport for participating media. In *Proc. of Eurographics Workshop on Rendering Techniques*, Springer-Verlag, London, UK, 11–22.
- PERLIN, K. H., AND HOFFERT, E. M. 1989. Hypertexture. Computer Graphics (Proc. of SIGGRAPH) 23, 3 (July), 253–262.
- PHARR, M., AND HUMPHREYS, G. 2010. Physically Based Rendering: From Theory to Implementation, 2nd ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- RAAB, M., SEIBERT, D., AND KELLER, A. 2008. Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer, 591–606.
- SKULLERUD, H. R. 1968. The stochastic computer simulation of ion motion in a gas subjected to a constant electric field. *Journal* of Physics D: Applied Physics 1, 11, 1567–1568.
- SZIRMAY-KALOS, L., TÓTH, B., AND MAGDICS, M. 2011. Free path sampling in high resolution inhomogeneous participating media. *Computer Graphics Forum 30*, 1, 85–97.
- VEACH, E., AND GUIBAS, L. J. 1994. Bidirectional estimators for light transport. In Proc. of Eurographics Rendering Workshop 1994, 147–162.
- VON NEUMANN, J. 1951. Various techniques used in connection with random digits. *Journal of Research of the National Bureau* of Standards, Appl. Math. Series 12, 36–38.
- WOODCOCK, E., MURPHY, T., HEMMINGS, P., AND T.C., L. 1965. Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry. In *Applications of Computing Methods to Reactor Problems*, Argonne National Laboratory.
- YUE, Y., IWASAKI, K., CHEN, B.-Y., DOBASHI, Y., AND NISHITA, T. 2010. Unbiased, adaptive stochastic sampling for rendering inhomogeneous participating media. ACM TOG (Proc. of SIGGRAPH Asia) 29, 6 (Dec.), 177:1–177:8.