

Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes

PETER KUTZ, Walt Disney Animation Studios
RALF HABEL, Walt Disney Animation Studios
YINING KARL LI, Walt Disney Animation Studios
JAN NOVÁK, Disney Research




Figure 1. A cloudscape rendered with a combination of our spectral and decomposition tracking techniques, which gracefully handle chromatic media and reduce collision coefficient evaluations. The insets on the right were computed in equal time, with our method yielding 3.5× lower MSE than delta tracking.

We present two novel unbiased techniques for sampling free paths in heterogeneous participating media. Our *decomposition tracking* accelerates free-path construction by splitting the medium into a control component and a residual component and sampling each of them separately. To minimize expensive evaluations of spatially varying collision coefficients, we define the control component to allow constructing free paths in closed form. The residual heterogeneous component is then homogenized by adding a fictitious medium and handled using weighted delta tracking, which removes the need for computing strict bounds of the extinction function. Our second contribution, *spectral tracking*, enables efficient light transport simulation in chromatic media. We modify free-path distributions to minimize the fluctuation of path throughputs and thereby reduce the estimation variance. To demonstrate the correctness of our algorithms, we derive them *directly* from the radiative transfer equation by extending the integral formulation of null-collision algorithms recently developed in reactor physics. This mathematical framework, which we thoroughly review, encompasses existing trackers and postulates an entire family of new estimators for solving transport problems; our algorithms are examples of such. We analyze the proposed methods in canonical settings and on production scenes, and compare to the current state of the art in simulating light transport in heterogeneous participating media.

CCS Concepts: • **Computing methodologies** → **Rendering**; *Ray tracing*;

Additional Key Words and Phrases: participating media, volume rendering, free-path sampling, transmittance, delta tracking, ratio tracking, color

ACM Reference format:

Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes. *ACM Trans. Graph.* 36, 4, Article 111 (July 2017), 16 pages.
DOI: <http://dx.doi.org/10.1145/3072959.3073665>

1 INTRODUCTION

Accurate and efficient simulation of radiative transfer in participating media is essential in many domains, such as nuclear reactor design, medical imaging, scientific visualization, and realistic image synthesis. The animation and visual effects industry, in particular, employs rich and complex volumetric structures (such as smoke, fire, or clouds) and translucent materials (such as marble, wax, or skin) to create aesthetically pleasing visuals and convey natural depth and scale cues. Simulating radiative transfer in such materials, however, presents a significant computational challenge. Monte Carlo methods based on path tracing—a popular tool for solving transport problems—require constructing a massive number of light paths. When the appearance of the medium is dominated by high-order scattering, each such light path can consist of up to thousands of scattering events. Modeling these scattering events—by means of sampling free paths—is the main contributor to the cost of synthesizing an image, and the main focus of this paper.

In continuous homogeneous media, stochastically constructing free paths is relatively cheap as we can obtain distance samples analytically. The general case of heterogeneous media, however, requires a specific approach—delta tracking—that homogenizes the volume parameters by adding a *fictitious* medium that interacts with light only formally. The properties of the fictitious medium are defined such that, upon a collision, light scatters forward with its direction and radiant intensity unaltered. While the inclusion of these *null collisions* is of no use for the principal description of radiative transfer, it does enable formulating Monte Carlo methods for sampling free paths. The main drawback of delta tracking is the large number of expensive lookups of spatially varying volume parameters.

We propose to accelerate delta tracking by decomposing the scattering medium into a homogeneous *control* component and a heterogeneous *residual* component. Individual free paths, or their segments, are then constructed using only one component. If the control component is used, the calculation is fully analytic and we avoid expensive memory lookups or procedural evaluations, keeping the cost of constructing paths low (see Figure 2). A similar decomposition was previously explored in the context of transmittance estimation [Novák et al. 2014]; we investigate the application of this concept to free-path sampling.

We also address the challenges of efficiently handling wavelength-dependent collision coefficients. We analyze several vectorized variants of weighted tracking and propose a history-aware strategy for setting collision probabilities to prevent geometric growth of the path throughput.

In order to demonstrate that our techniques are unbiased, we derive them directly from the radiative transfer equation (RTE) using the *integral formulation* of null-collision algorithms recently proposed by Galtier et al. [2013]. This mathematical framework defines an entire, largely unexplored family of Monte Carlo methods for solving radiative transfer in heterogeneous media, opening opportunities for developing new practical algorithms of which our decomposition and spectral trackers are examples. Such improvements would be difficult to achieve, if not impossible, when postulating the algorithms using physical interpretations. Instead, we derive our trackers directly from the RTE, obviating the need for more involved mathematical proofs [Coleman 1968; Miller 1967].

In what follows, we review several existing trackers for constructing free paths (Section 2), rederive Galtier et al.’s reformulation of the RTE (Section 3), and discuss configurations that lead to a number of known, previously published algorithms. We then introduce the concept of decomposing a scattering medium to sample free paths more efficiently (Section 4). For didactic purposes, we start with an intuitive, albeit constrained solution, which we later generalize by direct derivation from the RTE. The resulting algorithm reduces evaluations of spatially varying coefficients by obtaining some of the free-path samples analytically. We further improve the algorithm using a complementary approach for efficiently handling wavelength-dependent collision coefficients (Section 5). We analyze individual techniques in isolation and combined together in canonical settings and in realistic rendering scenarios (Section 6), demonstrating multiple facets of improved performance over existing methods.




Figure 2. Our decomposition tracking constructs some of the free paths (yellow) in a heterogeneous medium without evaluating spatially varying collision coefficients. The portion of yellow segments represents the savings in memory lookups. The resulting light paths are identical to delta tracking.

2 EXISTING TRACKING ALGORITHMS

The propagation of light in a participating medium is governed by the radiative transfer equation [Chandrasekhar 1960], which relates the change in radiative energy to collision coefficients of the medium: the absorption coefficient $\mu_a(\mathbf{x})$, the scattering coefficient $\mu_s(\mathbf{x})$, and their sum, the extinction coefficient $\mu_t(\mathbf{x}) = \mu_a(\mathbf{x}) + \mu_s(\mathbf{x})$. Each of these coefficients quantifies the distance density (with physical units of inverse length) of the corresponding physical process. A key tool for numerically solving the RTE is a stochastic construction of light particles’ trajectories that consist of free paths, i.e. straight flight paths between subsequent collisions with the medium. Unbiased estimators that rely on constructing such paths can be classified as either *analog* methods, that directly model the actual physical process, or *non-analog* methods, that bias the free-path distribution and reweight samples to produce unbiased results with lower variance or without the constraints of the analog process.

Closed-form Tracking. In simple volumes, e.g. those with constant, polynomial, or exponentially varying extinction, free paths can be sampled using the inverse transform sampling, first applied to Monte Carlo integration by Ulam [Eckhardt 1987]. We define a probability density function (PDF) by normalizing the transmittance function $T(t) = \exp(-\int_0^t \mu_t(\mathbf{x}_t) dt)$ along a ray parameterized by t . If the corresponding cumulative distribution function (CDF) is analytically invertible, then the free-path distance t' can be sampled analytically using a single random number ζ and a closed-form expression, e.g. $t' = -\log(1 - \zeta) / \mu_t$ for homogeneous volumes. In cases when the CDF is not invertible, but an analytic solution for the optical thickness $\tau(t) = \int_0^t \mu_t(\mathbf{x}_t) dt$ still exists, one can find the solution iteratively using Newton’s method [Brown and Martin 2003].

Regular Tracking. If the medium consists of piecewise homogeneous regions, the free path can be found by visiting boundaries along the ray and integrating optical thickness of individual regions until its negative exponentiated value reaches a stochastically chosen value of transmittance. This procedure, known as regular tracking [Sutton et al. 1999] or surface tracking [Leppänen 2010], quickly becomes expensive as we need to identify all boundaries between individual regions along the ray. Unless they can be discovered quickly, as with grids [Amanatides and Woo 1987] or octrees [Hubert-Tremblay et al. 2006], the method is relatively slow, especially when free paths cross many boundaries.

Ray Marching. A straightforward approach to reduce the cost of regular tracking is to ignore the boundaries and march along the ray with a fixed stride [Perlin and Hoffert 1989]. However, ignoring the true heterogeneity and assuming a constant extinction along each step leads to bias, which persists even if the marching is randomly jittered [Raab et al. 2008]. The bias is characterized by Jensen’s inequality [Jensen 1906], $\exp(E[-\tau]) \leq E[\exp(-\tau)]$, and can be reduced by taking smaller stride steps (ideally conforming to the Nyquist frequency), but this significantly reduces performance.

Delta Tracking. Unbiased sampling of free paths in heterogeneous volumes can be achieved using delta tracking. The method is based on von Neumann’s [1951] rejection sampling and strictly adheres to the physical process, hence matching the classification of analog estimators. It was independently developed in neutron transport [Bertini 1963; Woodcock et al. 1965; Zerby et al. 1961] and in plasma physics [Skullerud 1968], and is also known as Woodcock tracking, the null-collision algorithm, or pseudo scattering.

The main idea of delta tracking is to introduce a fictitious medium, represented by the *null-collision coefficient* $\mu_n(\mathbf{x})$, that homogenizes the total collision density in order to enable analytic sampling of free paths. The shortening of free paths due to increasing the combined collision coefficient is counteracted by rejecting some collisions and continuing the free flight forward. The rejected collisions—referred to as *null collisions*—occur with probability equal to the ratio of the null-collision coefficient to the combined (real + null) collision coefficient. The fictitious medium can be interpreted as being perfectly transparent and thereby having no effect on light transport; its sole purpose is to allow analytic distance sampling. The correctness of the method was rigorously proven using two different mathematical approaches [Coleman 1968; Miller 1967].

Since null collisions incur computational overhead, several approaches were developed to lower the null-collision coefficient, typically by partitioning the medium, e.g. using kd-trees [Yue et al. 2010, 2011] or grids [Szirmay-Kalos et al. 2011], and optimizing the amount of added fictitious medium for each region independently. Our decomposition tracking is complementary to these approaches and further reduces computational cost by omitting queries of $\mu_t(\mathbf{x})$.

Weighted Tracking. Another way to reduce null collisions is to allow the combined coefficient to occasionally underestimate the real extinction. To remain unbiased, the resulting non-analog estimator must appropriately reweight samples that step into regions with negative values of the null-collision coefficient [Carter et al. 1972; Cramer 1978]. Various weighting schemes were also proposed to artificially increase collision rates in measurement regions [Spanier and Gelbard 1969; Steen 1966], or to ensure that free paths reach a certain distance in deep-penetration problems [Cramer 1978; Morgan and Kotlyar 2015] and transmittance calculations [Novák et al. 2014] by replacing the rejection with a statistical weight. Eymet et al. [2013] investigated the benefits of combining weighted tracking with energy partitioning [Shamsundar et al. 1973].

While reasoning about negative extinctions using physical arguments is difficult, a purely mathematical view [Galtier et al. 2013] reveals an entire family of weighted algorithms that all lead to the correct solution; we review this integral framework in Section 3.

Multi-volume Tracking. Until now we have only considered scattering volumes that consist of a single, spatially varying medium. For materials where absorption can occur due to multiple transitions, Galtier et al. [2016] proposed to reduce the tracking cost by probabilistically evaluating only a single term in the sum. We combine this concept with the control/residual decomposition, which was proposed for transmittance estimation [Novák et al. 2014; Szirmay-Kalos et al. 2011]. Thanks to its multiplicative nature, transmittance can be calculated as the product of per-component estimates, reducing variance if the decomposition is chosen carefully. Our application to free-path sampling (Section 4) relies on a different theoretical justification, and, instead of reducing sample variance, we reduce the per-sample cost.

Multi-wavelength Tracking. Collision coefficients that vary spectrally require special handling. Raab et al. [2008] recommended avoiding wavelength-dependent extinction coefficients when possible, and otherwise tracing separate light paths for each wavelength [Eymet et al. 2013]. In the latter case, the tracking cost grows linearly with the number of wavelengths.

Tracing multiple wavelengths together can increase efficiency and reduce distracting color noise. Wilkie et al. [2014] described a technique in which a single “hero” wavelength is chosen for path construction, secondary wavelengths are carried along, and their (potentially high) variance is kept low using multiple importance sampling (MIS) [Veach 1997]. A similar approach was used by Chiang et al. [2016] who derive wavelength selection probabilities from per-wavelength throughput and single-scattering albedo. These techniques work well in homogeneous media. The hero-wavelength method can be extended to heterogeneous volumes, however, it employs delta tracking *and* ray marching for calculating a single free-path sample and the proposed implementation is biased. In contrast, our multi-wavelength tracking (Section 5) is unbiased and does not require explicit selection of a hero wavelength. We also discuss situations that permit combining multiple path-construction strategies, that employ an unbiased rejection-based tracker, using MIS without explicitly evaluating or approximating distance PDFs.

3 RADIATIVE TRANSFER WITH NULL COLLISIONS

In this section, we derive a version of the RTE that includes null collisions. We then extend it by incorporating the work of Galtier et al. [2013] to obtain the integral formulation of null-collision algorithms. Finally, we elaborate on specific configurations that yield delta tracking and weighted delta tracking and discuss other existing trackers in appendices. Readers familiar with the framework may skip directly to Section 4.

3.1 Integral Formulation of RTE with Null Collisions

The radiative transfer equation in its differential form

$$(\omega \cdot \nabla)L(\mathbf{x}, \omega) = -\mu_t(\mathbf{x})L(\mathbf{x}, \omega) + \mu_a(\mathbf{x})L_e(\mathbf{x}, \omega) + \mu_s(\mathbf{x}) \int_{S^2} f_p(\omega, \bar{\omega})L(\mathbf{x}, \bar{\omega}) d\bar{\omega} \quad (1)$$

describes the equilibrium radiance $L(\mathbf{x}, \omega)$ parameterized by point \mathbf{x} and direction of travel ω . The first term on the right side of Equation (1) expresses the losses due to absorption and out-scattering.

The gains due to radiant emission $L_e(\mathbf{x}, \omega)$ and inscattering are described in the second and third term, respectively. The phase function $f_p(\omega, \bar{\omega})$ quantifies the directional density of scattered light.

In order to incorporate null collisions, we can formally write the transport due to the null-collision coefficient $\mu_n(\mathbf{x})$ as

$$-\mu_n(\mathbf{x})L(\mathbf{x}, \omega) + \mu_n(\mathbf{x}) \int_{S^2} \delta(\omega - \bar{\omega})L(\mathbf{x}, \bar{\omega}) d\bar{\omega} = 0. \quad (2)$$

The losses due to the first term equal to the null-collided radiance that is “inscattered” using the Dirac delta function. Since the two terms cancel out, we can add them to the RTE without changing its validity, yielding [Cramer 1978]

$$\begin{aligned} (\omega \cdot \nabla)L(\mathbf{x}, \omega) = & -[\mu_t(\mathbf{x}) + \mu_n(\mathbf{x})]L(\mathbf{x}, \omega) + \mu_a(\mathbf{x})L_e(\mathbf{x}, \omega) \\ & + \mu_s(\mathbf{x}) \int_{S^2} f_p(\omega, \bar{\omega})L(\mathbf{x}, \bar{\omega}) d\bar{\omega} \\ & + \mu_n(\mathbf{x}) \int_{S^2} \delta(\omega - \bar{\omega})L(\mathbf{x}, \bar{\omega}) d\bar{\omega}. \end{aligned} \quad (3)$$

We will refer to the term on the last line as the *null-collided radiance*. By integrating both sides spatially along ω , introducing the shorthand notation $L_s(\mathbf{x}, \omega) = \int_{S^2} f_p(\omega, \bar{\omega})L(\mathbf{x}, \bar{\omega}) d\bar{\omega}$, and solving the Dirac integral we obtain the integral form of the RTE including null collisions:

$$\begin{aligned} L(\mathbf{x}, \omega) = & \int_0^\infty \exp\left(-\int_0^t \mu_t(\mathbf{x}_s) + \mu_n(\mathbf{x}_s) ds\right) \left[\mu_a(\mathbf{x}_t)L_e(\mathbf{x}_t, \omega) \right. \\ & \left. + \mu_s(\mathbf{x}_t)L_s(\mathbf{x}_t, \omega) + \mu_n(\mathbf{x}_t)L(\mathbf{x}_t, \omega) \right] dt, \end{aligned} \quad (4)$$

where $\mathbf{x}_t = \mathbf{x} - t\omega$ and $\mathbf{x}_s = \mathbf{x} - s\omega$. In this formulation, the change in transmittance (the exponential term) due to null collisions is always compensated for by the null-collided radiance: if $\mu_n(\mathbf{x})$ is positive, the reduction in transmittance is counteracted by adding the *positive* null-collided radiance. If $\mu_n(\mathbf{x})$ is negative, then the surplus of radiance due to overestimating transmittance is removed by adding the *negative* null-collided radiance.

Since negative values of collision coefficients are undefined in the physical sense, we must abandon a purely physical interpretation of Equation (4) here. The equation is still a valid RTE reformulation though, as null collisions have no impact on light transport.

3.2 Integral Formulation of Tracking Methods

In order to provide an algorithmic recipe for tracking algorithms in heterogeneous media, Galtier et al. [2013] introduce a number of identities that allow for translating Equation (4) directly into code for solving it. We loosely follow their derivation for absorptive-only media and extend it to include the process of scattering [Eymet et al. 2013]. We first define the probability density function (PDF)

$$p(t) = \bar{\mu}(\mathbf{x}_t) \exp\left(-\int_0^t \bar{\mu}(\mathbf{x}_s) ds\right), \quad (5)$$

which will be used to sample free-path lengths $t \in (0, \infty)$ along the source-seeking ray $(\mathbf{x}, -\omega)$ in closed form. We refer to $\bar{\mu}(\mathbf{x}) = \mu_t(\mathbf{x}) + \mu_n(\mathbf{x})$ as the *free-path-sampling coefficient*, which combines the extinction and the null-collision coefficients. Substituting the




Figure 3. A light path $X = \{\mathbf{x}_0 \dots \mathbf{x}_k\}$ in a heterogeneous medium consists of absorption, scattering, and null-collision events that determine the Monte Carlo score $w(X)$.

PDF for the exponential term in Equation (4) yields

$$\begin{aligned} L(\mathbf{x}, \omega) = & \int_0^\infty p(t) \left[\frac{\mu_a(\mathbf{x}_t)}{\bar{\mu}(\mathbf{x}_t)} L_e(\mathbf{x}_t, \omega) + \frac{\mu_s(\mathbf{x}_t)}{\bar{\mu}(\mathbf{x}_t)} L_s(\mathbf{x}_t, \omega) \right. \\ & \left. + \frac{\mu_n(\mathbf{x}_t)}{\bar{\mu}(\mathbf{x}_t)} L(\mathbf{x}_t, \omega) \right] dt. \end{aligned} \quad (6)$$

Next, we introduce probabilities $P_a(\mathbf{x})$, $P_s(\mathbf{x})$, and $P_n(\mathbf{x})$ to allow probabilistically evaluating each of the emission, inscattering, and null-collided radiance terms, respectively. To distinguish the individual free-path segments that are built to numerically solve the recursive RTE, we denote $\mathbf{x}_0 \equiv \mathbf{x}$ and $\mathbf{x}_{j+1} = \mathbf{x}_j - t_j\omega_j$:

$$\begin{aligned} L(\mathbf{x}_j, \omega_j) = & \int_0^\infty p(t_j) \left[P_a(\mathbf{x}_{j+1}) \frac{\mu_a(\mathbf{x}_{j+1})L_e(\mathbf{x}_{j+1}, \omega_j)}{\bar{\mu}(\mathbf{x}_{j+1})P_a(\mathbf{x}_{j+1})} \right. \\ & + P_s(\mathbf{x}_{j+1}) \frac{\mu_s(\mathbf{x}_{j+1})L_s(\mathbf{x}_{j+1}, \omega_j)}{\bar{\mu}(\mathbf{x}_{j+1})P_s(\mathbf{x}_{j+1})} \\ & \left. + P_n(\mathbf{x}_{j+1}) \frac{\mu_n(\mathbf{x}_{j+1})L(\mathbf{x}_{j+1}, \omega_j)}{\bar{\mu}(\mathbf{x}_{j+1})P_n(\mathbf{x}_{j+1})} \right] dt_j. \end{aligned} \quad (7)$$

Figure 3 shows a light path consisting of a number of free-path segments between different types of collision events, indexed by j .

Similarly to Galtier et al. [2013], we formalize the notion of probabilistic testing using the Heaviside function, which returns 1 if the condition is true, and 0 otherwise. A probabilistic evaluation of a random variable X can be written as the product

$$PX = \int_0^1 \mathcal{H}[y < P]X dy, \quad (8)$$

where P is the probability of evaluating X . Introducing \mathcal{H} and substituting $\mathbf{x}_+ = \mathbf{x}_{j+1}$ for readability, we obtain

$$\begin{aligned} L(\mathbf{x}_j, \omega_j) = & \int_0^\infty p(t_j) \\ & \times \left[\int_0^1 \mathcal{H}[\xi_e < P_a(\mathbf{x}_+)] \frac{\mu_a(\mathbf{x}_+)}{\bar{\mu}(\mathbf{x}_+)P_a(\mathbf{x}_+)} L_e(\mathbf{x}_+, \omega_j) d\xi_e \right. \\ & + \int_0^1 \mathcal{H}[\xi_s < P_s(\mathbf{x}_+)] \frac{\mu_s(\mathbf{x}_+)}{\bar{\mu}(\mathbf{x}_+)P_s(\mathbf{x}_+)} L_s(\mathbf{x}_+, \omega_j) d\xi_s \\ & \left. + \int_0^1 \mathcal{H}[\xi_n < P_n(\mathbf{x}_+)] \frac{\mu_n(\mathbf{x}_+)}{\bar{\mu}(\mathbf{x}_+)P_n(\mathbf{x}_+)} L(\mathbf{x}_+, \omega_j) d\xi_n \right] dt_j. \end{aligned} \quad (9)$$

By replacing the integrals with Monte Carlo estimators, this version of the RTE can be directly translated into a multitude of recursive algorithms that all yield the correct solution *regardless* of the choice of $\mu_n(\mathbf{x})$ and probabilities $P_a(\mathbf{x})$, $P_s(\mathbf{x})$, and $P_n(\mathbf{x})$. While these degrees of freedom may be uninteresting from a purely theoretical point of view, their presence is of high interest for numerical methods

Algorithm 1. *Delta tracking and weighted delta tracking.* The combined free-path-sampling coefficient $\bar{\mu}$ is assumed to be constant along the ray $(\mathbf{x}, -\omega)$. ζ and ξ represent uniformly distributed random numbers.

<pre> 1 function DELTATRACKING(x, ω) 2 while true do 3 t ← -ln(1-ζ)/μ̄ 4 x ← x - t × ω 5 if ξ < μ_a(x)/μ̄ then 6 return L_e(x, ω) 7 else if ξ < 1 - μ_n(x)/μ̄ then 8 ω ← sample ∝ f_p(ω) 9 else 10 continue </pre>	<pre> 1 function WGTDELTA TRACKING(x, ω) 2 w ← 1 3 while true do 4 t ← -ln(1-ζ)/μ̄ 5 x ← x - t × ω 6 if ξ < P_a(x) then 7 return w × μ_a(x)L_e(x, ω) 8 else if ξ < 1 - P_n(x) then 9 ω ← sample ∝ f_p(ω) 10 w ← w × μ_s(x)/(μ̄P_s(x)) 11 else 12 w ← w × μ_n(x)/(μ̄P_n(x)) </pre>
---	--

as different choices lead to vastly different convergence characteristics and algorithmic constraints. Next, we review two existing techniques deriving them directly from Equation (9).

3.3 Delta Tracking

In standard delta tracking, the probabilities for evaluating the emitted, in-scattered, and null-collided radiance are strictly proportional to the relative value of the corresponding coefficient:

$$P_a(\mathbf{x}) = \frac{\mu_a(\mathbf{x})}{\bar{\mu}(\mathbf{x})}, \quad P_s(\mathbf{x}) = \frac{\mu_s(\mathbf{x})}{\bar{\mu}(\mathbf{x})}, \quad P_n(\mathbf{x}) = \frac{\mu_n(\mathbf{x})}{\bar{\mu}(\mathbf{x})}. \quad (10)$$

In order for these probabilities to be mathematically valid (i.e. in $[0, 1]$), the free-path-sampling coefficient $\bar{\mu}(\mathbf{x})$ is set to a *majorant* of $\mu_t(\mathbf{x})$, typically represented by a piecewise constant function. These probabilities perfectly cancel the fractions first introduced in Equation (7), leaving “only” the unweighted radiance functions to be evaluated. To this end, most implementations employ *branching* to evaluate only one of the radiance functions and thus avoid a geometric growth of function calls. Since the probabilities sum to one, and exactly one decision test is true at any time, Equation (9) simplifies to

$$L(\mathbf{x}_j, \omega_j) = \int_0^\infty p(t_j) \int_0^1 \mathcal{H}[\xi < P_a(\mathbf{x}_+)] L_e(\mathbf{x}_+, \omega_j) + \mathcal{H}[P_a(\mathbf{x}_+) < \xi < 1 - P_n(\mathbf{x}_+)] L_s(\mathbf{x}_+, \omega_j) + \mathcal{H}[1 - P_n(\mathbf{x}_+) < \xi] L(\mathbf{x}_+, \omega_j) d\xi dt_j. \quad (11)$$

Figure 1 (left) shows a direct translation of Equation (11) into the code of a Monte Carlo estimator. First, a random distance t is sampled from $p(t_j)$ to estimate the outer integral. Then a uniformly distributed random variable ξ selects one of the **if** (line 5), **else if** (line 7), and **else** (line 9) branches that correspond to the three Heaviside functions in the inner integral.

3.4 Weighted Delta Tracking

One way to remove the burden of finding tight majorants on $\mu_t(\mathbf{x})$ is to introduce a weighted Monte Carlo scheme [Carter et al. 1972; Cramer 1978; Galtier et al. 2013]. Notice that Equation (6) is still valid even if the free-path-sampling coefficient is not a majorant, i.e. $\bar{\mu}(\mathbf{x}) < \mu_t(\mathbf{x})$, as long as we provide mathematically valid definitions of the branching probabilities. Freeing themselves from a physical

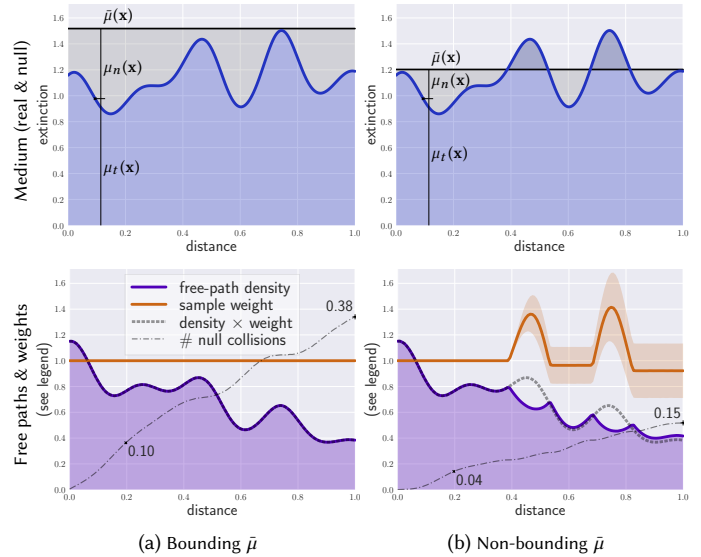


Figure 4. Weighted delta tracking in a medium with bounding and non-bounding free-path-sampling coefficient. The stretching of the free-path probability density in (b) is compensated for by weighting the samples. The red line, illustrating the weight, starts diverging from 1 as we enter regions where $\bar{\mu} < \mu_t(\mathbf{x})$. The density \times weight product is in both cases the same, but the statistical error is higher in (b) due to the very noisy weight; the red band illustrates its standard deviation. The dashed line shows the average number of null-collisions, which is reduced in (b) thanks to the lower $\bar{\mu}$.

interpretation, Cramer [1978] and Galtier et al. [2013] exploit this observation and propose to set

$$P_a(\mathbf{x}) = \frac{\mu_a(\mathbf{x})}{\mu_t(\mathbf{x}) + |\mu_n(\mathbf{x})|}, \quad (12)$$

$$P_s(\mathbf{x}) = \frac{\mu_s(\mathbf{x})}{\mu_t(\mathbf{x}) + |\mu_n(\mathbf{x})|}, \quad (13)$$

$$P_n(\mathbf{x}) = \frac{|\mu_n(\mathbf{x})|}{\mu_t(\mathbf{x}) + |\mu_n(\mathbf{x})|}. \quad (14)$$

The stretching of the free-path PDF is compensated for by properly reweighting the radiance functions; see Figure 4 for an illustration. Without restricting ourselves to a particular definition of the branching probabilities, we can express the Monte Carlo score $w(\mathbf{X})$ of path sample \mathbf{X} as

$$w(\mathbf{X}) = L_e(\mathbf{x}_k, \omega_k) w_a(\mathbf{x}_k) \prod_{j=1}^{k-1} w_\star(\mathbf{x}_j), \quad (15)$$

$$w_\star(\mathbf{x}_j) = \frac{\mu_\star(\mathbf{x}_j)}{\bar{\mu}(\mathbf{x}_j) P_\star(\mathbf{x}_j)}, \quad (16)$$

where $\mathbf{X} = \{\mathbf{x}_0 \dots \mathbf{x}_k\}$, k is the index of the first absorption collision—the termination point—and w_\star is a local collision weight with the subscript \star representing a , s , or n , depending on whether the corresponding event was an absorptive, scattering, or null collision. The product of all collision weights is referred to as the path throughput. Algorithm 1 (right) shows a pseudocode that implements a branching version of Equation (9) and computes $w(\mathbf{X})$ progressively.

Table 1. Collision coefficients used by the decomposition trackers.

Symbol & Value	Description
$\mu_a(\mathbf{x}) \in [0, \infty)$	absorption coef. of original volume
$\mu_s(\mathbf{x}) \in [0, \infty)$	scattering coef. of original volume
$\mu_t(\mathbf{x}) = \mu_a(\mathbf{x}) + \mu_s(\mathbf{x})$	extinction coef. of original volume
$\mu_a^c \in [0, \infty)$	absorption coef. of control volume
$\mu_s^c \in [0, \infty)$	scattering coef. of control volume
$\mu_t^c = \mu_a^c + \mu_s^c$	extinction coef. of control volume
$\mu_a^r(\mathbf{x}) = \mu_a(\mathbf{x}) - \mu_a^c$	absorption coef. of residual volume
$\mu_s^r(\mathbf{x}) = \mu_s(\mathbf{x}) - \mu_s^c$	scattering coef. of residual volume
$\mu_t^r(\mathbf{x}) = \mu_t(\mathbf{x}) - \mu_t^c$	extinction coef. of residual volume
$\bar{\mu} \in (\mu_t^c, \infty)$	free-path-sampling coefficient
$\mu_n(\mathbf{x}) = \bar{\mu} - \mu_t^c - \mu_t^r(\mathbf{x})$	null-collision coefficient

Discussion. Not having to define the free-path-sampling coefficient as a majorant of $\mu_t(\mathbf{x})$ eases the search for an analytically integrable function $\bar{\mu}(\mathbf{x})$ that closely approximates $\mu_t(\mathbf{x})$. However, numerous analyses [Carter et al. 1972; Cramer 1978; Galtier et al. 2013; Novák et al. 2014] demonstrate that the variance of the path throughput increases rapidly with non-bounding $\bar{\mu}(\mathbf{x})$. In fact, the throughput will alternate sign whenever a null collision occurs and $\mu_n(\mathbf{x})$ is negative. It is thus still important to seek free-path-sampling coefficients that bound $\mu_t(\mathbf{x})$ in most parts of the volume. Having the ability to occasionally violate this rule is nonetheless very practical.

Other Estimators. In Appendix A we discuss the configuration that yields the free-path-sampling technique briefly discussed by Novák et al. [2014]. Numerous transmittance estimators employing various forms of weighting [Cramer 1978; Novák et al. 2014] can also be formalized using the aforementioned integral framework, as long as the underlying equations are adapted to calculate only transmittance; we show how to do this in Appendix B.

4 DECOMPOSITION TRACKING

In this section, we present a volume-decomposition technique for accelerating free-path sampling. Our approach is motivated by the benefits of applying such decomposition to transmittance estimation [Novák et al. 2014], where it lowers variance and reduces memory accesses due to handling a part of the volume—the control component—analytically.

Similarly to previous work, we decompose the original volume into a homogeneous *control component* $\{\mu_a^c, \mu_s^c\}$ and the remaining heterogeneous *residual component* $\{\mu_a^r(\mathbf{x}), \mu_s^r(\mathbf{x})\}$; see Table 1 for symbol overview and Figure 5 for an illustration.

In the next two sections, we describe how to generate free-path samples using the two components. We begin with an intuitive, albeit limited approach (Section 4.1) that cannot incorporate sample weighting, to later derive a more flexible formulation (Section 4.2) using the integral framework from Section 3.




Figure 5. We decompose a medium (a) into a homogeneous control component (b) and a heterogeneous residual component (c). A free-path sample in the original medium can then be obtained by taking the minimum of free-path samples in individual components.

4.1 Analog Decomposition Tracking

We observe that a free-path sample in the original medium can be obtained by taking the *minimum* of free-path samples in individual components. This concept is illustrated in Figure 5 and can be understood intuitively by considering the first collision of a ray with two separate sets of particles. If we intersect the union of the two sets, only the collision at the first, or minimum, free-flight distance from either set actually matters. This is similar to finding the nearest hit in surface rendering. We will now prove rigorously that free-path samples obtained by taking the minimum follow the desired distribution. An extended version of this proof can be found in the supplemental material.

THEOREM 1. *Let non-negative extinction coefficients $\mu_A(\mathbf{x}), \mu_B(\mathbf{x})$ be combined to $\mu_C(\mathbf{x}) = \mu_A(\mathbf{x}) + \mu_B(\mathbf{x})$. Let A, B , and C be independent random variables distributed according to the following CDF $F_X(t) = 1 - \exp(-\int_0^t \mu_X(\mathbf{x}_s) ds)$ with $X \in \{A, B, C\}$. The CDFs of C and $\min(A, B)$ are identical.*

PROOF. Consider a random variable $D = \min(A, B)$. According to basic probability theory, the CDF of D reads

$$F_D(t) = F_A(t) + F_B(t) - F_A(t)F_B(t). \quad (17)$$

Substituting the CDFs of A and B into the previous relation and using $T_X(t) = \exp(-\int_0^t \mu_X(\mathbf{x}_s) ds)$, we obtain

$$\begin{aligned} F_D(t) &= [1 - T_A(t)] + [1 - T_B(t)] - [1 - T_A(t)][1 - T_B(t)] \\ &= 1 - T_A(t)T_B(t). \end{aligned} \quad (18)$$

Since $T_A(t)$ and $T_B(t)$ are exponential functions and integration is a linear operator, Equation (18) simplifies to

$$F_D(t) = 1 - T_A(t)T_B(t) = 1 - T_C(t) = F_C(t).$$

Hence we have demonstrated that the CDF of C is indeed equal to the CDF of $\min(A, B)$. \square

The combination is therefore very elegant—we create free-path samples for each component and simply take the minimum of them to arrive at the combined estimate. However, the proof assumes A, B , and C to be distributed according to transmittance, i.e. they obey an exponential distribution parameterized by negative optical thickness. As such, we can use the aforementioned min-based combination only with analog trackers; weighted trackers that produce arbitrary weighted distributions must be combined differently.

Algorithm 2. *Analog decomposition tracking*. ζ , ψ , and ξ are uniformly distributed random numbers.

```

1 function ANALOGDECOMPOSITIONTRACKING(x, ω)
2   while true do
3     tc ← -ln(1-ζ) / μtc
4     tr ← 0
5     while true do
6       tr ← tr - ln(1-ψ) / (μ̄ - μtc)
7       if tr > tc then
8         x ← x - tc × ω
9         if ξ < μac(x) / μtc then
10          return Le(x, ω)
11        else
12          ω ← sample ∝ fp(ω)
13          break
14      else
15        if ξ < μar(x) / (μ̄ - μtc) then
16          x ← x - tr × ω
17          return Le(x, ω)
18        else if ξ < 1 - μn(x) / (μ̄ - μtc) then
19          x ← x - tr × ω
20          ω ← sample ∝ fp(ω)
21        break

```

The analog decomposition tracking requires both the residual extinction coefficient $\mu_t^r(\mathbf{x})$ and the null-collision coefficient $\mu_n(\mathbf{x})$ to be always non-negative. The control extinction μ_t^c thus needs to be a *minorant* of $\mu_t(\mathbf{x})$ to ensure $\mu_t^r(\mathbf{x}) \geq 0$, and the free-path-sampling coefficient of the residual component μ^r must be a *majorant* of $\mu_t^r(\mathbf{x})$ to ensure $\mu_n(\mathbf{x}) \geq 0$.

Efficient Implementation. Given that the decomposition requires drawing two free-path samples, it may appear less efficient than delta tracking in the original volume. However, a simple short-circuit optimization can save computation and most of the expensive extinction queries. We first analytically sample the free path in the control component. Then we use delta tracking to construct a free path in the residual component, but as soon as the delta tracker exceeds the control free path, we stop and return the control sample. This avoids the lookups needed to finish the free-path sampling in the residual component. Algorithm 2 presents the algorithm of the described approach. The technique can also be adjusted to combine distance samples from multiple separate superimposed volumes, without having all of these volumes in memory at the same time, facilitating easy integration into streaming renderers.

Limitations. In addition to requiring the majorant, the main drawback of the analog method is that it also needs a minorant for decomposing the medium. The best performance requires both bounds to be tight. As mentioned before, computing tight bounds is challenging in practice—sometimes nearly impossible if the function is defined procedurally. Removing this requirement would make the decomposition tracking a lot more practical. Furthermore, since the analog decomposition relies on delta tracking, it can handle only a single wavelength (discussed in Section 5), and cannot easily accommodate additional information to further reduce variance. We overcome these limitations in the next section.

Algorithm 3. *Weighted decomposition tracking*. F is a CDF and the assignment operator \leftarrow returns the assigned value.

```

1 function WEIGHTEDDECOMPOSITIONTRACKING(x, ω)
2   w ← 1
3   while true do
4     t ← -ln(1-ξ) / μ̄
5     x ← x - t × ω
6     F ← 0
7     if ξ < (F ← F + Pac(x)) then
8       return w × μac(x) / μ̄ × Le(x, ω)
9     else if ξ < (F ← F + Psc(x)) then
10      ω ← sample ∝ fp(ω)
11      w ← w × μsc(x) / μ̄ Psc(x)
12     else if ξ < (F ← F + Par(x)) then
13      return w × μar(x) / μ̄ Par(x) × Le(x, ω)
14     else if ξ < (F ← F + Psr(x)) then
15      ω ← sample ∝ fp(ω)
16      w ← w × μsr(x) / μ̄ Psr(x)
17     else
18      w ← w × μn(x) / μ̄ Pn(x)

```

4.2 Weighted Decomposition Tracking

In this section, we derive a *weighted* version of the decomposition tracker that can handle non-bounding control extinctions and free-path-sampling coefficients, and enables efficient handling of chromatic collision coefficients—both of which are problematic with the analog version. The RTE with null collisions from Equation (4) can be rewritten for a medium with k absorptive and/or scattering components [Galtier et al. 2016]:

$$\begin{aligned}
L(\mathbf{x}, \omega) = & \int_0^\infty \exp\left(-\int_0^t \sum_{i=1}^k \mu_t^i(\mathbf{x}_s) + \mu_n(\mathbf{x}_s) ds\right) \\
& \times \left[\sum_{i=1}^k \mu_a^i(\mathbf{x}_t) L_e(\mathbf{x}_t, \omega) + \sum_{i=1}^k \mu_s^i(\mathbf{x}_t) L_s(\mathbf{x}_t, \omega) \right. \\
& \left. + \mu_n(\mathbf{x}_t) L(\mathbf{x}_t, \omega) \right] dt, \tag{19}
\end{aligned}$$

where the superscript i is used to index individual components. Substituting in $p(t)$ and incorporating probabilities and their probabilistic evaluation from Section 3.2 yields

$$\begin{aligned}
L(\mathbf{x}_j, \omega_j) = & \int_0^\infty p(t_j) \\
& \times \left[\sum_{i=1}^k \int_0^1 \mathcal{H}[\xi_e^i < P_a^i(\mathbf{x}_+)] \frac{\mu_a^i(\mathbf{x}_+)}{\mu(\mathbf{x}_+) P_a^i(\mathbf{x}_+)} L_e(\mathbf{x}_+, \omega_j) d\xi_e^i \right. \\
& + \sum_{i=1}^k \int_0^1 \mathcal{H}[\xi_s^i < P_s^i(\mathbf{x}_+)] \frac{\mu_s^i(\mathbf{x}_+)}{\mu(\mathbf{x}_+) P_s^i(\mathbf{x}_+)} L_s(\mathbf{x}_+, \omega_j) d\xi_s^i \\
& \left. + \int_0^1 \mathcal{H}[\xi_n < P_n(\mathbf{x}_+)] \frac{\mu_n(\mathbf{x}_+)}{\mu(\mathbf{x}_+) P_n(\mathbf{x}_+)} L(\mathbf{x}_+, \omega_j) d\xi_n \right] dt_j. \tag{20}
\end{aligned}$$

In our case, Equation (20) can be simplified since we have only the two previously introduced components, the *control* and the *residual*. We thus have exactly five terms to evaluate: emitted and inscattered radiance of the control component, emitted and inscattered radiance of the residual component, and the null-collided radiance.

Control Probabilities. We shall now define the probabilities of evaluating each of the five terms. Similarly to previous tracking methods, we enable branching by forcing the probabilities to sum to 1 and relate them to the respective coefficients to keep the local collision weights close to 1. Most importantly though, we define the probabilities of sampling the control component using only the constant coefficients, avoiding possibly expensive lookups of spatially varying coefficients:

$$P_a^c = P_c P_{a|c} = \frac{\mu_t^c \mu_a^c}{\bar{\mu} \mu_t^c} = \frac{\mu_a^c}{\bar{\mu}}, \quad (21)$$

$$P_s^c = P_c P_{s|c} = \frac{\mu_t^c \mu_s^c}{\bar{\mu} \mu_t^c} = \frac{\mu_s^c}{\bar{\mu}}. \quad (22)$$

In words, the probability of sampling emission or inscattering using the control component is set to the product of the marginal probability $P_c = \mu_t^c / \bar{\mu}$ of choosing the control component, and the conditional probability of choosing the respective term. Note that the definition of P_c requires $\bar{\mu} > \mu_t^c$:

Residual and Null-collision Probabilities. The residual and null-collision components are handled analogously, except that we employ absolute values (as in Section 3.4) to handle situations when either of the components is negative:

$$P_a^r(\mathbf{x}) = \left(1 - \frac{\mu_t^c}{\bar{\mu}}\right) \frac{|\mu_a^r(\mathbf{x})|}{|\mu_a^r(\mathbf{x})| + |\mu_s^r(\mathbf{x})| + |\mu_n(\mathbf{x})|}, \quad (23)$$

$$P_s^r(\mathbf{x}) = \left(1 - \frac{\mu_t^c}{\bar{\mu}}\right) \frac{|\mu_s^r(\mathbf{x})|}{|\mu_a^r(\mathbf{x})| + |\mu_s^r(\mathbf{x})| + |\mu_n(\mathbf{x})|}, \quad (24)$$

$$P_n(\mathbf{x}) = \left(1 - \frac{\mu_t^c}{\bar{\mu}}\right) \frac{|\mu_n(\mathbf{x})|}{|\mu_a^r(\mathbf{x})| + |\mu_s^r(\mathbf{x})| + |\mu_n(\mathbf{x})|}. \quad (25)$$

The first term always represents the marginal probability of *not* sampling the control component. The second term then defines the conditional probability of selecting the residual absorption, residual scattering, and null collision. The algorithmic realization of weighted decomposition tracking is shown in Algorithm 3.

The control probabilities, as defined in Equations (21) and (22), always cancel out with the fractions in Equation (20). Taking one of the two control branches therefore does not change the path throughput. This is also the case for the other branches when all of the residual and null coefficients are positive; otherwise, the path throughput is adjusted to account for negative values. Unlike the analog tracker, the weighted version allows handling chromatic coefficients, which we describe in Section 5.

4.3 Analysis & Discussion

In the analog approach, the residual component is sampled using a much lower free-path-sampling coefficient than in the case of the original volume. Delta tracking on the residual component thus performs longer steps, which, given that we stop the tracker as soon as it exceeds the control free path, reduces the number of lookups of spatially-varying coefficients; see Figure 6 for a visualization. Despite the weighted formulation using a different procedure for constructing the free path, it produces *identical* lookup statistics as the analog version—the statistics stem from the parameters of the volume, not the procedure itself.

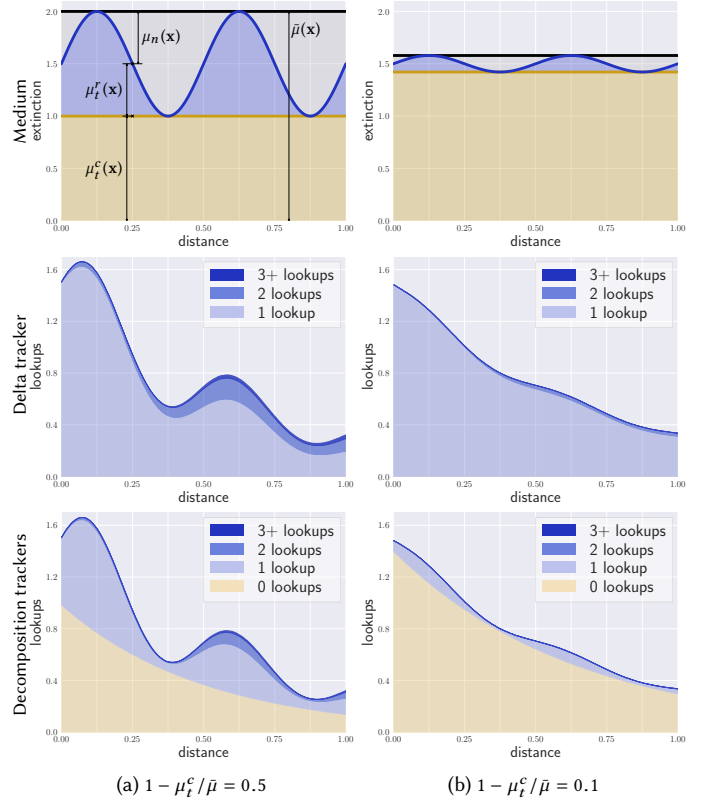


Figure 6. Number of extinction lookups required without and with decomposition in two media with high and low degree of heterogeneity. The middle and bottom rows depict free-path densities visualizing the portions of free-path samples according to the number of lookups they performed. The bottom row represents both analog and weighted decomposition tracking.

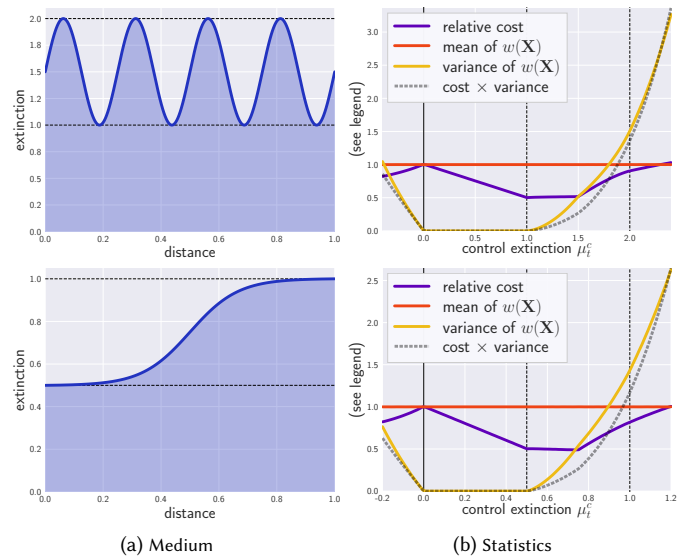


Figure 7. Visualization of the relative cost ρ (purple) of *weighted* decomposition tracking, the mean and the variance of path throughput until the first real collision (orange), and their product (dashed); all plotted w.r.t. μ_t^c .

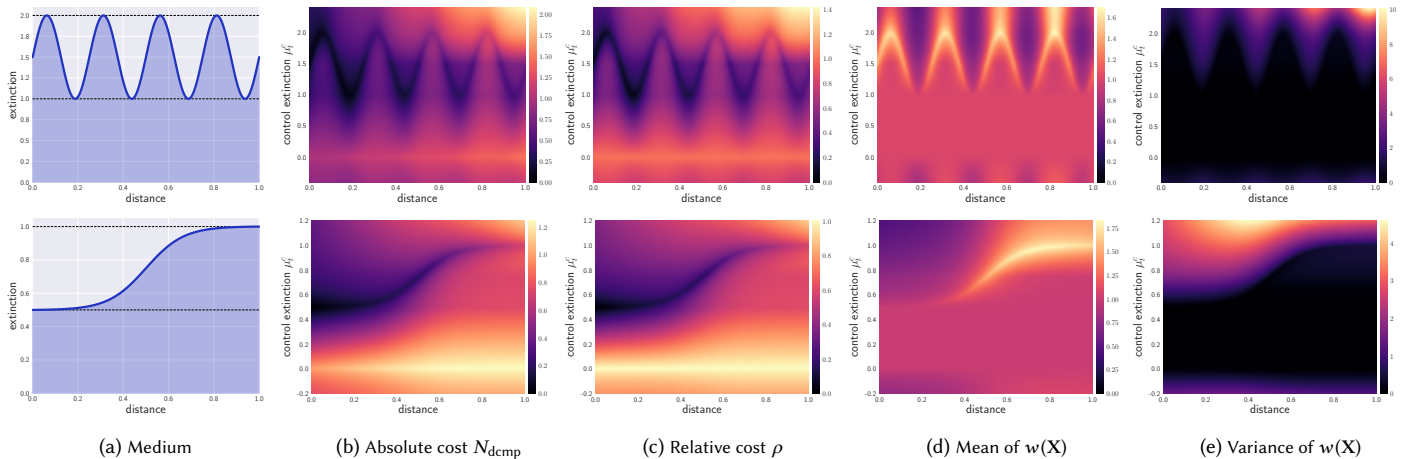


Figure 8. For media shown in column (a), we visualize the absolute cost of our *weighted-decomposition* tracking in column (b), the relative cost of our *weighted* decomposition tracking w.r.t. delta tracking ($\rho = N_{dcmp}/N_{delta}$) in column (c), the mean path throughput $w(X)$ after the first real collision in column (d), and the variance of $w(X)$ in column (e).

Relative Cost. Denoting the number of spatially-varying extinction lookups performed by our decomposition tracker N_{dcmp} , and the number of lookups performed by delta tracker in the original volume N_{delta} , we can define the relative cost of decomposition tracking to be $\rho = N_{dcmp}/N_{delta} = 1 - \mu_t^c/\bar{\mu}$. This relation, which holds only when the control extinction $\mu_t^c \in [0, \min(\mu_t(x))]$, i.e. μ_t^c is a minorant of $\mu_t(x)$, is depicted in Figure 7(b) by the straight purple line in interval $[0, \min(\mu_t(x))]$. When $\mu_t^c > \min(\mu_t(x))$, the relative cost increases to eventually surpass 1, i.e. the decomposition tracking becomes more expensive than delta tracking. This is to be expected as the control medium at that point significantly overestimates the density of the original medium.

Mean and Variance of $w(X)$. The statistics in Figure 7 were collected by invoking the tracker multiple times and recording, in addition to the cost, the path throughput after the first real collision. The orange and yellow curves depict the mean and the variance of the path throughput, respectively. Because the tracker is unbiased, the *mean* of the path throughput $w(X)$ at the first collision (over all paths) is always 1. Furthermore, if $\mu_t^c \in [0, \min(\mu_t(x))]$, then the collision weight w_\star is always 1 and the *variance* of $w(X)$ is thus zero.* Similarly to the case of the relative cost, the tracker becomes inefficient when $\mu_t^c > \min(\mu_t(x))$, in which case the variance of $w(X)$ increases rapidly.

Optimal Configuration. In Figure 8, we further study the cost and the path throughput as a function of distance (horizontal axis) and the value of control extinction μ_t^c (vertical axis). The absolute cost N_{dcmp} in column (b) is minimized when μ_t^c equals the local extinction coefficient $\mu_t(x)$. This is also the situation in which decomposition tracking brings the largest cost savings over delta tracking, as can be seen in column (c) that visualizes ρ . In the same configuration, however, the throughputs tends to have significantly more extreme

*Note that zero variance of the path throughput does not imply zero variance in actual estimation problems, where the randomness of the path's termination point can contribute significantly to the estimation variance.

values than elsewhere, and often above 1. Paths undergoing several collisions with such a configuration will tend to suffer from geometric growth of the throughput. This explains the steep increase of variance (yellow curve) in Figure 7(b) when $\mu_t^c > \min(\mu_t(x))$.

Based on these experiments, and the plots in Figure 7, we conclude that the best performance is achieved by setting the control extinction to the tight lower bound of the extinction function, i.e. $\mu_t^c = \min(\mu_t(x))$. However, having the option of overestimating $\min(\mu_t(x))$ and still converging to the correct solution is one of the main advantages of the weighted decomposition tracker over the analog version.

It is also worth noting that as the residual part reduces to 0, the performance of decomposition tracking approaches that of closed-form tracking with the relative number of lookups approaching 0, i.e. $\rho \rightarrow 0$. As the control part reduces to 0, the performance of the algorithm approaches that of delta tracking with $\rho \rightarrow 1$. The decomposition tracking thus bridges the performance gap between closed-form tracking and delta tracking, with the weighted variant allowing arbitrary values of μ_t^c and $\bar{\mu}$ to be handled at the expense of increased cost and weight variance.

Numerical Equivalence with Delta Tracking. Either form of decomposition tracking can be used as a drop-in replacement for delta tracking. In fact, the weighted decomposition tracking can produce path samples that are identical to delta tracking (just at a lower cost). Both trackers (c.f. Algorithms 1 and 3) draw the same pair of random numbers ζ and ξ . In both cases ζ is used to sample a tentative distance, and ξ is used to select an interaction type. As long as ζ and ξ stem from the same random-number sequences, the free-path-sampling coefficients are the same, and the interaction types are *ordered analogously* in both implementations (which we omit in the pseudocode for readability), both delta tracking and decomposition tracking produce identical results.

5 SPECTRAL TRACKING

As described in Section 3, and illustrated in Figure 4, weighted tracking can produce correct results even if the distribution of distance samples does not obey Beer’s law: the bias in the distribution is counteracted by properly reweighting the samples. We leverage this ability to efficiently handle spectrally resolved collision coefficients.

We define a single distribution for sampling collisions and counteract the discrepancy between this distribution and the true, per-wavelength free-path distribution by reweighting. We invoke the tracker only once for all wavelengths and, at each collision, we compute a spectrally resolved collision weight. The weight can be expressed using a spectral basis or simply as a vector of N_λ wavelengths: we use three primaries—red, green, and blue (RGB)—in all experiments but our formulas and proofs hold for arbitrary counts. Continuous spectra can be handled by carrying a discrete set of wavelengths per-path, e.g. by using the wavelength selection and bundling techniques presented by Wilkie et al. [2014]. Algorithm 4 shows a vectorized weighted delta tracking that was used to sample (weighted) free paths in a medium with independently varying RGB extinction coefficients in Figure 9.

Free-path-sampling Coefficient. In order to keep the variance of the estimation low, we should ideally set the free-path-sampling coefficient $\bar{\mu}(\mathbf{x})$ to bound the maximum value of $\mu_t(\mathbf{x}, \lambda)$ across all wavelengths λ , or as close to this value as possible, as discussed in Section 3.4.

5.1 Collision Probabilities

The variance of the spectral tracker largely depends on the variance of the spectrally resolved collision weights. The weight will inevitably deviate from 1 as we cannot set the collision probabilities to cancel out the $\mu_\star(\mathbf{x}, \lambda)/\bar{\mu}(\mathbf{x})$ fractions for all wavelengths simultaneously. The worst situations occur when $P_\star(\mathbf{x}) \ll |\mu_\star(\mathbf{x}, \lambda)/\bar{\mu}(\mathbf{x})|$, for which $|w_\star(\mathbf{x}, \lambda)| \gg 1$. Should there be multiple such collisions, the path throughput would grow geometrically for that wavelength. However, by carefully setting the collision probabilities, we can bound the local collision weight and also the throughput of the path.

5.1.1 Maximum- and Average-based Probabilities. We can bound the magnitude of the local collision weight $|w_\star(\mathbf{x}, \lambda)|$ for all λ by deriving the probabilities from *maxima* over all wavelengths:

$$P_a(\mathbf{x}) = \max_{\lambda} (|\mu_a(\mathbf{x}, \lambda)|) c^{-1}, \quad (26)$$

$$P_s(\mathbf{x}) = \max_{\lambda} (|\mu_s(\mathbf{x}, \lambda)|) c^{-1}, \quad (27)$$

$$P_n(\mathbf{x}) = \max_{\lambda} (|\mu_n(\mathbf{x}, \lambda)|) c^{-1}, \quad (28)$$

$$c = \max_{\lambda} (|\mu_a(\mathbf{x}, \lambda)|) + \max_{\lambda} (|\mu_s(\mathbf{x}, \lambda)|) + \max_{\lambda} (|\mu_n(\mathbf{x}, \lambda)|), \quad (29)$$

where c is a normalization constant. When all collision coefficients are non-negative, it can be shown that the proposed probabilities keep the local collision weight at or below the *number of collision types*, i.e. $w_\star(\mathbf{x}, \lambda) \leq 3$, for all λ regardless of the total number of wavelengths N_λ being traced; see the supplementary material for a proof. In contrast, replacing the $\max()$ function by $\text{avg}()$ would keep the local collision weight at or below the *number of wavelengths* N_λ .

Algorithm 4. *Spectral tracking.* The hat ($\hat{\cdot}$) denotes vectorized quantities that are handled using element-wise multiplication and division.

```

1 function SPECTRALTRACKING( $\mathbf{x}, \omega$ )
2    $\hat{w} \leftarrow (1, \dots, 1)_{N_\lambda}$ 
3   while true do
4      $t \leftarrow -\frac{\ln(1-\xi)}{\bar{\mu}}$ 
5      $\mathbf{x} \leftarrow \mathbf{x} - t \times \omega$ 
6     if  $\xi < P_a(\mathbf{x})$  then
7       return  $\hat{w} \circ \frac{\hat{\mu}_a(\mathbf{x})}{\hat{\mu} P_a(\mathbf{x})} \circ \widehat{L}_e(\mathbf{x}, \omega)$ 
8     else if  $\xi < 1 - P_n(\mathbf{x})$  then
9        $\omega \leftarrow \text{sample} \llcorner f_p(\omega)$ 
10       $\hat{w} \leftarrow \hat{w} \times \frac{\hat{\mu}_s(\mathbf{x})}{\hat{\mu} P_s(\mathbf{x})}$ 
11    else
12       $\hat{w} \leftarrow \hat{w} \times \frac{\hat{\mu}_n(\mathbf{x})}{\hat{\mu} P_n(\mathbf{x})}$ 

```

5.1.2 Incorporating Path History. Despite being bounded, the local collision weights can still exceed 1. Multiplying a sequence of them together may therefore result in a geometric growth of the path throughput. We can suppress this trend by incorporating the path throughput—path history—into the *average-based* collision probabilities. At each j -th vertex of the path, we include the path throughput $w(\mathbf{X}_{j-1}, \lambda) = \prod_{i=1}^{j-1} w_\star(\mathbf{x}_i, \lambda)$, of the already built subpath \mathbf{X}_{j-1} . The probabilities then become:

$$P_a(\mathbf{x}_j) = \text{avg}_{\lambda} (|\mu_a(\mathbf{x}_j, \lambda) w(\mathbf{X}_{j-1}, \lambda)|) c^{-1}, \quad (30)$$

$$P_s(\mathbf{x}_j) = \text{avg}_{\lambda} (|\mu_s(\mathbf{x}_j, \lambda) w(\mathbf{X}_{j-1}, \lambda)|) c^{-1}, \quad (31)$$

$$P_n(\mathbf{x}_j) = \text{avg}_{\lambda} (|\mu_n(\mathbf{x}_j, \lambda) w(\mathbf{X}_{j-1}, \lambda)|) c^{-1}, \quad (32)$$

$$\begin{aligned}
c &= \text{avg}_{\lambda} (|\mu_a(\mathbf{x}_j, \lambda) w(\mathbf{X}_{j-1}, \lambda)|) \\
&\quad + \text{avg}_{\lambda} (|\mu_s(\mathbf{x}_j, \lambda) w(\mathbf{X}_{j-1}, \lambda)|) \\
&\quad + \text{avg}_{\lambda} (|\mu_n(\mathbf{x}_j, \lambda) w(\mathbf{X}_{j-1}, \lambda)|). \quad (33)
\end{aligned}$$

With these probabilities the sum of the throughputs over all wavelengths is always equal to the *number of wavelengths* N_λ , as long as $\bar{\mu}(\mathbf{x}_j)$ is a majorant; a proof can be found in the supplemental material. This property implies that the average of the throughputs over all of the wavelengths is always—for all path vertices—equal to 1, and that the maximum throughput for any wavelength is N_λ . The geometric growth of the path throughput during successive collisions is thereby prevented.

Incorporating the path history into maximum-based probabilities only slows down the growth but does not limit the throughput.

Discussion. In practice, these probabilities reduce the magnitude of undesired color noise and avoid bright chromatic outliers (“fireflies”). The ability to keep the average throughput at 1 is analogous to delta tracking, where one wavelength is selected randomly and the throughput divided by the selection probability, which amounts to multiplying by N_λ . In fact, in the extreme case of completely uncorrelated chromatic coefficients, our technique gracefully regresses to naive (weighted) delta tracking and automatically settles on tracing only one particular wavelength at a time. When this occurs, $\bar{\mu}(\mathbf{x}_j)$ can be reduced to bound only the one relevant wavelength.

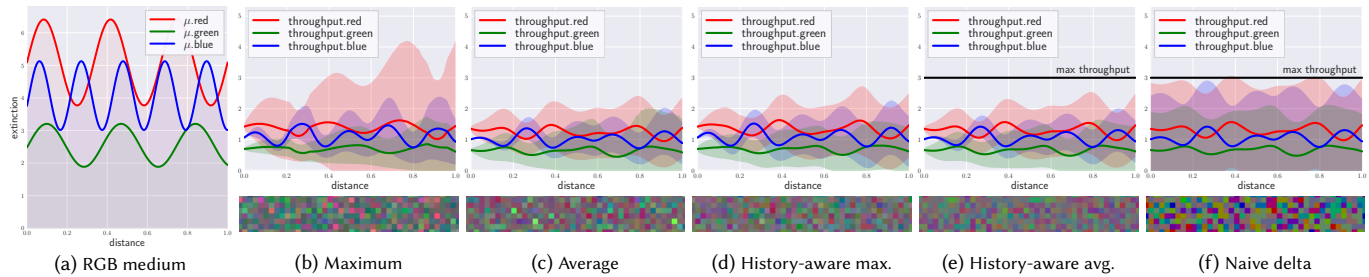


Figure 9. Comparison of four variants of our spectral tracking (b)-(e), described in Section 5.1, and naive delta tracking that handles each wavelength separately (f) in a perfectly forward-scattering, non-absorptive medium with spectrally and spatially varying $\mu_s(\mathbf{x}, \lambda)$ (a). The top row visualizes the mean path throughput with ± 1 standard deviation band for each wavelength along the line of travel. The bottom row shows noise exemplars produced by each technique.

In the other extreme, when the extinction coefficient is achromatic, our technique produces the same noise as standard delta tracking, provided that the latter used only one simulation for all wavelengths. In between these extremes, our technique provides advantages of both techniques, automatically and robustly bridging the gap between the “one-wavelength-at-a-time” and “one-tracking-for-all-wavelengths” variants of delta tracking.

5.1.3 Reduced Termination Rates. If the medium is *non-emissive*, we can further reduce the throughput if desired by *not* simulating absorption collisions and redistributing the absorption probability among the other two collision types. We simply set $P_a(\mathbf{x}_j) = 0$, remove the corresponding term from c , and recompute $P_s(\mathbf{x}_j)$ and $P_n(\mathbf{x}_j)$. Since $P_s(\mathbf{x}_j)$ and $P_n(\mathbf{x}_j)$, which appear in denominators in Equation (20), are increased, the throughput upon a scattering or null-collision reduces. While disabling absorption increases the computational cost—we need to simulate longer paths—it can reduce variance, and the magnitude of outliers if $\bar{\mu}$ does not bound $\mu_t(\mathbf{x})$.

5.2 Spectral and Decomposition Tracking

There are numerous ways in which the decomposition technique from Section 4 can be applied to spectral tracking. In our implementation, we opted for the following straightforward approach:

- set $\bar{\mu}$ to bound the maximum value of $\mu_t(\mathbf{x}, \lambda)$ across all wavelengths λ ,
- set each component of μ_a^c and μ_s^c to the minimum value of $\mu_a(\mathbf{x}, \lambda)$ and $\mu_s(\mathbf{x}, \lambda)$, respectively, across all wavelengths,
- set the control probabilities to vectorized versions of Equations (21) and (22), and
- use Equations (30), (31), and (32), modified to account only for the residual component, to set residual probabilities.

If a control component is sampled, the local collision weight is $(1, \dots, 1)_{N_\lambda}$. When a residual component is sampled, the desirable properties afforded by the probability schemes presented in Sections 5.1.1 and 5.1.2 still apply.

In summary, to efficiently handle chromatic media, we vectorize the collision coefficients, decompose the medium into an achromatic control component and a chromatic residual component, derive the collision probabilities from the maxima or averages across all wavelengths, incorporate the path history, and, if the medium is non-emissive, reduce termination rates in favor of decreasing the growth of path-throughput weights.

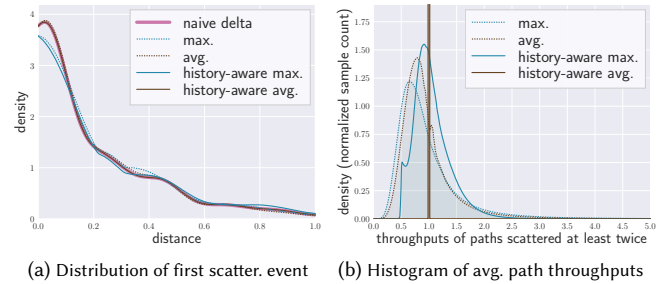


Figure 10. Distributions of the first scattering event (a) visualized for the five spectral trackers compared in Figure 9. The histograms of path throughputs in (b) illustrate how the history-aware average-based probabilities bound the path throughput regardless of the number of scattering events and the spatial and spectral variations of the collision coefficients.

5.3 Analysis & Discussion

Figure 9 visualizes multi-bounce path throughputs—their means and standard-deviation bands—as a function of distance obtained with (non-decomposed) spectral tracking with various schemes for setting the collision probabilities. We also include the naive delta tracking algorithm that handles each wavelength independently. In order to emphasize the differences in the weighting schemes, the medium is non-absorptive, i.e. albedo $\alpha = 1$, with largely uncorrelated scattering coefficients. Each spectral tracker produces paths with a different average number of scattering events, which are all around 4 in this case. For the trackers that bound the throughput, the throughput ceiling is visualized. The history-aware trackers provide a better-behaved path throughput and lower estimation noise than the other techniques in this test.

Figure 10(a) shows the distribution of the first scattering event for the setting in Figure 9. Note that the history-aware average-based probability scheme produces the same distribution of distances as naive delta tracking. Since all trackers are unbiased, multiplying the distance distributions in Figure 10(a) by the corresponding RGB throughputs always yields the same weighted distribution.

Figure 10(b) shows histograms of path throughputs after multiple (forward) scattering collisions. The wider the tail of the histogram, the more prone the tracker is to fireflies. In this scene, the history-aware average-based tracker is immune to such problems as it keeps the throughput at 1 during the path construction.

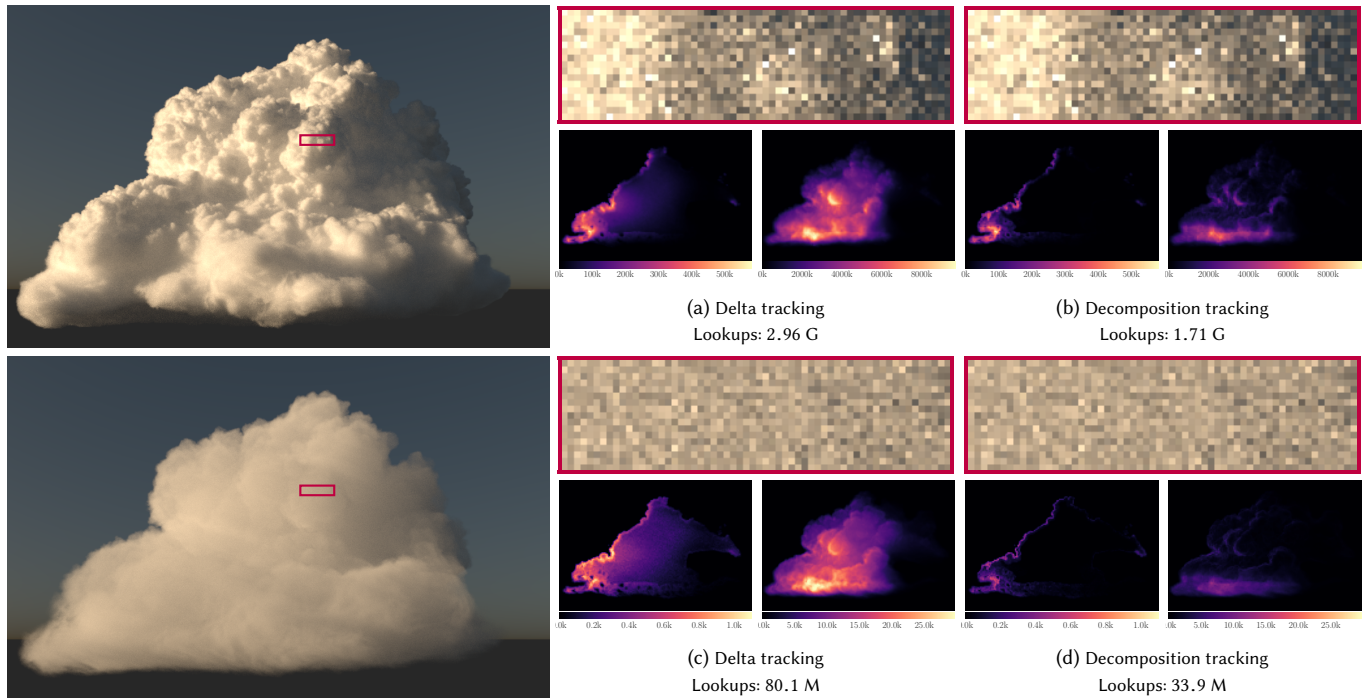


Figure 11. Delta tracking in (a), (c) and decomposition tracking in (b), (d) produce the same image up to identical noise patterns (see the top insets). However, decomposition tracking requires 42% and 58% fewer extinction lookups to render the optically thicker (top) and thinner (bottom) cloud, respectively. The lookup counts within a thin slice and across the full volume are visualized in the left and right false-color images, respectively, for renderings with 1024 spp.

6 IMPLEMENTATION & RESULTS

In addition to implementing the various trackers in a simple Monte Carlo framework for running the canonical experiments, we integrated them into a production path tracer for generating the results in this section. We use an octree as an acceleration structure to cull empty space and store approximate local minorants and majorants of the underlying volumes. The octree adapts to the homogeneity of the volume. We use VDB [Museth 2013] for storing voxel-based volumetric datasets.

We compare the performance using the number of lookups, the render time, the root mean square error (RMSE), and the “lookups-to-unit-variance” (LTUV) metric, which is computed as the product of the number of lookups and variance. For all of the aforementioned metrics, lower values represent better performance. The timings were measured on a dual 12-core 2.50 GHz Intel Xeon E5-2680 v3 but we report them in single-core equivalents.

Cloud Scene. Figure 11 shows an optically thick and an optically thin cloud with $\alpha = 1$, a Henyey–Greenstein phase function [Henyey and Greenstein 1941] with $g = 0.877$, and an octree depth of 8. The insets were rendered with 32 samples per pixel (spp), and the statistics are based on these renders. For the two configurations, our decomposition tracking avoids 42% and 58% of the extinction lookups required by delta tracking, respectively. Note that the noise produced by both trackers is identical as our decomposition tracking *does not* change the light paths, but generates them with significantly fewer lookups.

Table 2. Comparisons of delta and decomposition tracking performance for various octree depths. The scene is the thick cloud from Figure 11, rendered with 32 spp. The last two rows report the non-lookup tracker time.

Octree depth	1	4	8	16
Octree leaves visited	1.86 G	2.06 G	3.11 G	4.05 G
Lookup num. (Delta)	106 G	22.8 G	2.96 G	2.49 G
Lookup num. (Decomp)	106 G	22.6 G	1.71 G	1.03 G
Lookup time (Delta)	20065 s	5785 s	1030 s	876 s
Lookup time (Decomp)	20065 s	5472 s	536 s	336 s
Octree time (Delta)	64 s	315 s	734 s	916 s
Octree time (Decomp)	64 s	319 s	714 s	918 s
Tracker time (Delta)	6108 s	1389 s	299 s	321 s
Tracker time (Decomp)	6108 s	1379 s	280 s	300 s

The false-color images visualize the number of lookups on a plane cutting through the volume (left) and integrated over the entire volume by sweeping the plane through it (right). Note that the number of lookups is reduced most in the interior of the cloud where the extinction varies less.

In Table 2, we study the impact of the octree depth on the performance of the thick cloud. With no spatial partitioning (depth 1), delta and decomposition tracking exhibit near-identical performance; the control component is 0. Increasing the depth allows the




Figure 12. A colorful explosion with chromatic extinction rendered using 8 samples with delta tracking (b) in 198 core seconds, with spectral tracking using history-aware maximum-based probabilities (c) in 245 seconds, and with spectral tracking using history-aware average-based probabilities (d) in 266 seconds.

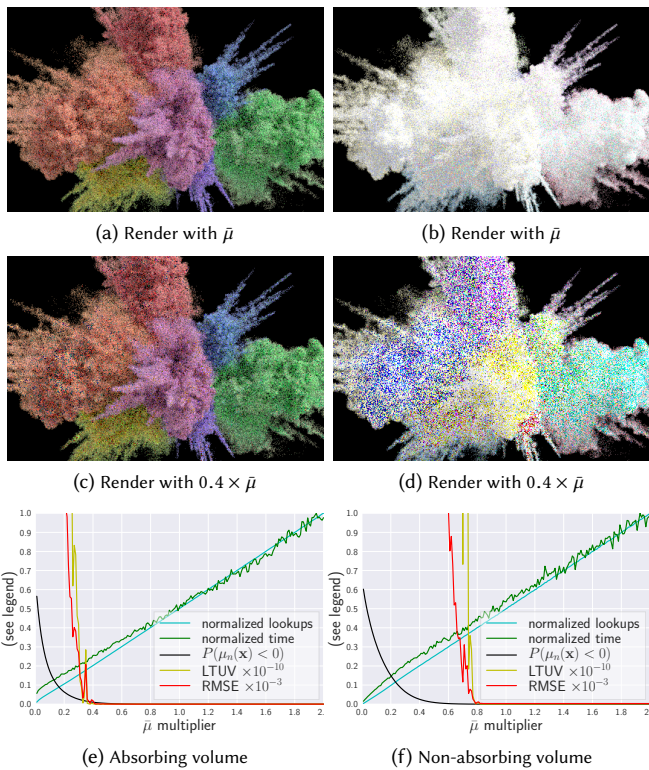


Figure 13. Effects of non-bounding free-path-sampling coefficients on the variance of rendering an absorbing (left) and a non-absorbing (right) color explosion scene. Both versions use the same colored extinction coefficients—only the albedo differs. In (a) and (b), we used a bounding $\bar{\mu}$. In (c) and (d), the $\bar{\mu}$ is scaled down by 0.4 leading to increased variance. The plots show various statistics as functions of the relative free-path-sampling coefficient.

decomposition to reduce the lookups more effectively since the volume within each octree leaf becomes more homogeneous; however, octree traversal eventually begins to dominate the render time. We found that maximum depths of 8–10 typically strike a good balance in such scenes. The decomposition reduces the overall render time by 10–20% as the octree traversal, directional sampling, and other overheads take significant time.




Figure 14. Comparison of the history-aware probabilities on the non-absorbing color explosion scene with adjusted exposure to visualize fireflies.

Color Explosion. In Figure 12, we show a medium with spectrally varying absorption, scattering, and extinction coefficients. We rendered the scene with 8 spp using one-wavelength-at-a-time delta tracking and our two history-aware spectral trackers. For the same number of light paths, spectral tracking increases the overall number of path segments, the number of lookups, and the render time all by 25%–30% in this scene. This is mainly due to taking smaller steps on average as the free-path-sampling coefficient must bound the extinction across *all* wavelengths in the highly chromatic medium. The slightly higher cost is compensated for by the 40%–45% lower RMSE, which amounts to $2.7\times$ – $3.4\times$ lower variance for the average-based and maximum-based probabilities, respectively. The overall improvement over delta tracking is $2.06\times$ – $2.46\times$ lower LTUV with the maximum-based probabilities yielding slightly better performance than the average-based ones in this particular scene.

In Figure 13, we study how various statistics of the spectral tracker depend on the value of the free-path-sampling coefficient with focus on situations when $\bar{\mu}(x) < \mu_t(x)$. We analyze two configurations: an absorbing (left) and a non-absorbing (right) version of the color explosion. The 16-spp images in the first and second rows show how the noise increases when the *bounding* piecewise constant $\bar{\mu}(x)$ is globally scaled down; see the supplementary material for more examples. As shown in the plots, lower values of the multiplier, i.e. more frequent and more significant underestimation of the bound, reduce the tracking cost. However, the variance and the LTUV rise abruptly due to many positive and negative fireflies, especially when the volume is non-absorbing.

In Figure 14, we compare the history-aware maximum-based and history-aware average-based probabilities on the non-absorbing color explosion scene, illustrating the robustness of the latter to high albedos. See the supplemental material for full images.

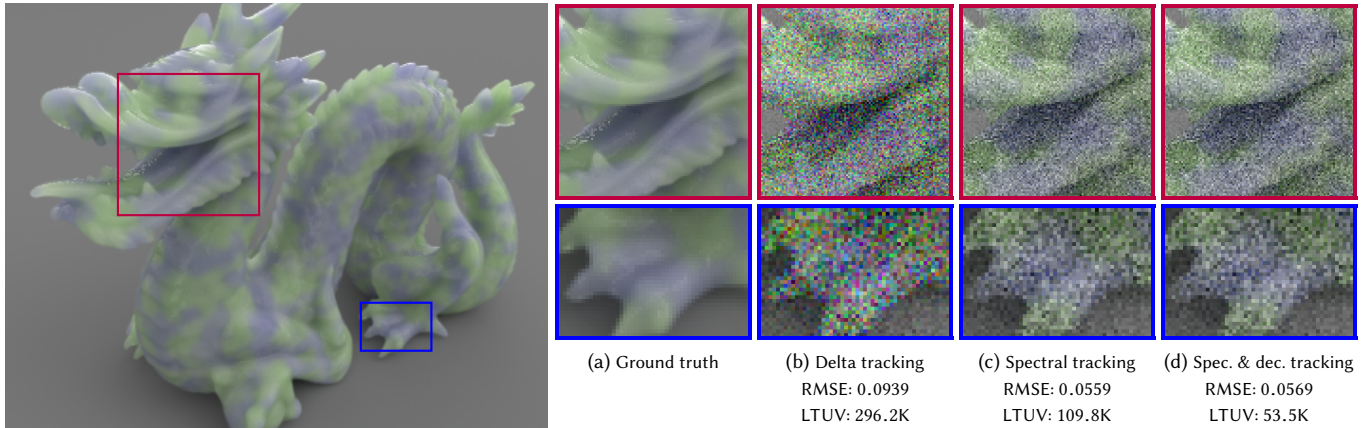


Figure 15. A translucent heterogeneous Stanford dragon with chromatic extinction rendered using 40 samples with delta tracking (b) in 133 core seconds, spectral tracking (c) in 112 core seconds, and spectral & decomposition tracking (d) in 88 core seconds.

Translucent Dragon. Our techniques can also be used for simulating subsurface scattering. Figure 15 shows a translucent dragon with a smooth dielectric boundary enclosing a colored, procedural medium. In this case, no acceleration data structure was used—the control and free-path-sampling coefficients are based on approximate minimum and maximum values over the entire medium, found by random sampling in a short pre-pass. In this figure, handling the spectrally resolved coefficients using our spectral tracking reduces the LTUV by 2.70 \times , and our weighted decomposition tracking further reduces it by 2.05 \times . Together, the spectral and decomposition techniques yield 5.54 \times lower LTUV than delta tracking and 1.65 \times lower RMSE.

The time spent simulating subsurface scattering reduces from 133 core seconds with delta tracking, to 112 with spectral tracking, and to 88 with spectral and decomposition tracking. The decrease in render time with spectral tracking is caused by different path distributions and by lower throughputs allowing Russian roulette to terminate paths earlier. The decrease in render time with decomposition easily compensates for the corresponding slight increase in RMSE. In all cases, the total render time is about 80 core seconds higher than the subsurface scattering time. Spectral and decomposition tracking together reduce the total render time by 22%.

7 DISCUSSION & FUTURE WORK

Performance. While our techniques—combined or standalone—outperform the previous state of the art in many situations, certain scene configurations may not harvest the full benefit. For instance, the efficiency of decomposition tracking largely depends on how tightly we can bound the extinction coefficient from below. This can become challenging when the medium exhibits a high degree of heterogeneity, in which case the performance largely depends on how well the spatial structure isolates and subdivides heterogeneous regions. We experimented with uniform grids and octrees, with the latter providing generally better performance. Higher efficiency could possibly be obtained with kd-trees [Yue et al. 2010] and polynomial bounds [Szirmay-Kalos et al. 2011].

An important feature of our weighted decomposition and spectral trackers is the ability to handle non-bounding free-path-sampling coefficients and control extinctions. Not having to worry about occasionally violating the desired inequalities $\mu_t^c(\mathbf{x}) \leq \min(\mu_t(\mathbf{x}))$, $\bar{\mu}(\mathbf{x}) \geq \max(\mu_t(\mathbf{x}))$ enables fast precomputation of near-optimal control and free-path-sampling coefficients, even for (non-pathological) procedural volumes.

In some situations, spatially varying lookups may not be the bottleneck, in which case decomposition tracking can only offer a rather small improvement in overall render time. Decomposition tracking excels in situations where memory lookups are expensive, data sets are large and prone to cache thrashing, or expensive high-order interpolation is used. Especially if the volume is modeled with complex procedural functions, the savings can be considerable.

It is also worth noting that our methods retain correctness and efficiency with arbitrary (valid) values of the collision coefficients and the albedo, as long as the free-path-sampling coefficient (approximately) bounds the extinction coefficient.

Unknown Distance PDF and MIS. The PDF of the free-path sample cannot be evaluated exactly (only estimated or approximated). This drawback is common to all rejection-based trackers and might appear as preventing the combination of multiple path-construction strategies using MIS, as in next-event estimation (NEE) or bidirectional techniques. However, we can still compute the optimal balance-heuristic weight if we use the *same* unbiased tracker for free-path and transmittance sampling; their PDFs are identical and cancel out in the calculation of the MIS weight. This is true even for our spectral tracker since the distance PDFs are not wavelength-dependent. The balance-heuristic weight is thus effectively based only on directional PDFs; this is what we use for volumetric NEE.

Note that the cancellation occurs only for certain MIS heuristics (e.g. balance or power) and only if we use the same tracker for free-path and transmittance sampling. This recipe thus does not generalize. Nevertheless, a near-optimal weight can almost always be computed using a deterministic approximation of the PDF (obtained e.g. by an extra ray-marching pass [Wilkie et al. 2014]).

7.1 Future Work

Adjoint-guided Tracking. As demonstrated in Section 5, the probabilities P_a , P_s , and P_n can be adjusted in many ways to achieve certain goals—in our case to reduce the weight when tracking multiple wavelengths. We have merely scratched the surface of the large body of problems that could benefit from having this freedom. For instance, if information about the distribution of light in the scene is available, the probabilities could be biased towards increasing the absorption rates [Spanier and Gelbard 1969; Steen 1966] in emissive parts, or scattering more in brightly lit regions. Incorporating techniques for importance sampling single or double scattering [Georgiev et al. 2013; Kalli and Cashwell 1977; Kulla and Fajardo 2012] into the tracker, and investigating advanced zero-variance schemes [Dwivedi 1982] would be interesting future avenues.

Dynamically Refined Extinction Bounds. To reduce the significance of finding tight extinction bounds a priori, the extinction bounds could be refined progressively based on actual extinction lookups during rendering. The initial extinction bounds could be set conservatively wide (slowing down the simulation at first), or to rough estimates (temporarily increasing variance), and then dynamically refined. This would be highly desirable for large production assets where lengthy precomputations impede productivity.

8 CONCLUSION

We presented two complementary techniques, the decomposition tracking for reducing costs, and the spectral tracking for reducing variance when simulating light transport in heterogeneous volumes. The combination of the two yields higher performance than previous state-of-the-art methods and gracefully handles spectrally resolved collision coefficients. We derived both approaches directly from the RTE using the recently proposed integral formulation of null-collision algorithms. We believe that importing this framework into computer graphics will stimulate further explorations in applications that are unique to rendering, but also increase the permeability between fields by enabling easy exchange of novel ideas.

ACKNOWLEDGMENTS

We thank Patrick Kelly for implementing various components of our volume rendering system, as well as for numerous discussions, and David Adler for help with profiling the code. Alex Nijmeh and Henrik Falt created the cloud scenes used in Figure 1 and Figure 11. Jesse Erickson created the colored explosion in Figure 12. The geometry of the dragon in Figure 15 is from the Stanford Computer Graphics Laboratory. We thank Thomas Müller, Marios Papas, Géraldine Conti, and David Adler for proofreading, and Brent Burley, Michael Kaschalk, and Rajesh Sharma for providing support, encouragement, and resources for this project. All images were rendered using Disney's Hyperion Renderer.

REFERENCES

- J. Amanatides and A. Woo. 1987. A fast voxel traversal algorithm for ray tracing. In *Eurographics '87*. 3–10.
- H. W. Bertini. 1963. *Monte Carlo simulations on intranuclear cascades*. Technical Report ORNL-3383. Oak Ridge National Laboratory, Oak Ridge, TN, USA. DOI: <https://doi.org/10.2172/4692927>
- F. B. Brown and W. R. Martin. 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.
- L. L. Carter, E. D. Cashwell, and W. M. Taylor. 1972. Monte Carlo sampling with continuously varying cross sections along flight paths. *Nuclear Science and Engineering* 48, 4 (1972), 403–411. DOI: <https://doi.org/10.13182/NSE72-1>
- S. Chandrasekhar. 1960. *Radiative transfer*. Dover Publications.
- M. J. Chiang, P. Kutz, and B. Burley. 2016. Practical and controllable subsurface scattering for production path tracing. In *ACM SIGGRAPH 2016 Talks (SIGGRAPH '16)*. ACM, New York, NY, USA, Article 49, 2 pages. DOI: <https://doi.org/10.1145/2897839.2927433>
- W. A. Coleman. 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 (April 1968), 76–81. DOI: <https://doi.org/10.13182/NSE68-1>
- S. N. Cramer. 1978. Application of the fictitious scattering radiation transport model for deep-penetration Monte Carlo calculations. *Nuclear Science and Engineering* 65, 2 (1978), 237–253.
- S. R. Dwivedi. 1982. Zero variance biasing schemes for Monte Carlo calculations of neutron and radiation transport problems. *Nuclear Science and Engineering* 80, 1 (1982), 172–178.
- R. Eckhardt. 1987. Stan Ulam, John von Neumann, and the Monte Carlo Method. *Los Alamos Science, Special Issue* (1987), 131–137.
- V. Eymet, D. Poitou, M. Galtier, M. El Hafi, G. Terrée, and R. Fournier. 2013. Null-collision meshless Monte-Carlo—Application to the validation of fast radiative transfer solvers embedded in combustion simulators. *Journal of Quantitative Spectroscopy and Radiative Transfer* 129 (April 2013), 145–157. DOI: <https://doi.org/10.1016/j.jqsrt.2013.06.004>
- M. Galtier, S. Blanco, C. Caliot, C. Coustet, J. Dauchet, M. El Hafi, V. Eymet, R. Fournier, J. Gautrais, A. Khuong, B. Piaud, and G. Terrée. 2013. Integral formulation of null-collision Monte Carlo algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer* 125 (April 2013), 57–68. DOI: <https://doi.org/10.1016/j.jqsrt.2013.04.001>
- M. Galtier, S. Blanco, J. Dauchet, M. El Hafi, V. Eymet, R. Fournier, M. Roger, C. Spiesser, and G. Terrée. 2016. Radiative transfer and spectroscopic databases: A line-sampling Monte Carlo approach. *Journal of Quantitative Spectroscopy and Radiative Transfer* 172 (March 2016), 83–97. DOI: <https://doi.org/10.1016/j.jqsrt.2015.10.016>
- I. Georgiev, J. Krivánek, T. Hachisuka, D. Nowrouzezahrai, and W. Jarosz. 2013. Joint importance sampling of low-order volumetric scattering. *ACM TOG (Proc. of SIGGRAPH Asia)* 32, 6 (Nov. 2013), 164:1–164:14. DOI: <https://doi.org/10.1145/2508363.2508411>
- L. G. Henyey and J. L. Greenstein. 1941. Diffuse radiation in the galaxy. *Astrophysical Journal* 93 (Jan. 1941), 70–83. DOI: <https://doi.org/10.1086/144246>
- V. Hubert-Tremblay, L. Archambault, D. Tubic, R. Roy, and L. Beaulieu. 2006. Octree indexing of DICOM images for voxel number reduction and improvement of Monte Carlo simulation computing efficiency. *Medical Physics* 33, 8 (2006), 2819–2831. DOI: <https://doi.org/10.1118/1.2214305>
- J. L. W. V. Jensen. 1906. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Math.* 30 (1906), 175–193. DOI: <https://doi.org/10.1007/BF02418571>
- H. J. Kalli and E. D. Cashwell. 1977. *Evaluation of three Monte Carlo estimation schemes for flux at a point*. Technical Report LA-6865-MS. Los Alamos Scientific Lab.
- C. Kulla and M. Fajardo. 2012. Importance sampling techniques for path tracing in participating media. *CGF (Proc. of Eurographics Symposium on Rendering)* 31, 4 (June 2012), 1519–1528. DOI: <https://doi.org/10.1111/j.1467-8659.2012.03148.x>
- J. Leppänen. 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 (2010), 715 – 722. DOI: <https://doi.org/10.1016/j.anucene.2010.01.011>
- L. B. Miller. 1967. *Monte Carlo analysis of reactivity coefficients in fast reactors; general theory and applications*. Technical Report ANL-7307. Argonne National Laboratory, Argonne, IL, USA. DOI: <https://doi.org/10.2172/4500721>
- L. W. G. Morgan and D. Kotlyar. 2015. Weighted-delta-tracking for Monte Carlo particle transport. *Annals of Nuclear Energy* 85 (2015), 1184–1188. DOI: <https://doi.org/10.1016/j.anucene.2015.07.038>
- K. Museth. 2013. VDB: High-resolution sparse volumes with dynamic topology. *ACM TOG* 32, 3 (July 2013), 27:1–27:22. DOI: <https://doi.org/10.1145/2487228.2487235>
- J. Novák, A. Selle, and W. Jarosz. 2014. Residual ratio tracking for estimating attenuation in participating media. *ACM TOG (Proc. of SIGGRAPH Asia)* 33, 6 (Nov. 2014), 179:1–179:11. DOI: <https://doi.org/10.1145/2661229.2661292>
- K. H. Perlin and E. M. Hoffert. 1989. Hypertexture. *Computer Graphics (Proc. of SIGGRAPH)* 23, 3 (July 1989), 253–262. DOI: <https://doi.org/10.1145/74334.74359>
- M. Raab, D. Seibert, and A. Keller. 2008. Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer, 591–606. DOI: https://doi.org/10.1007/978-3-540-74496-2_35
- N. Shamsundar, E. M. Sparrow, and R. P. Heimsch. 1973. Monte Carlo radiation solutions—effect of energy partitioning and number of rays. *International Journal of Heat and Mass Transfer* 16, 3 (1973), 690–694. DOI: [https://doi.org/10.1016/0017-9310\(73\)90236-6](https://doi.org/10.1016/0017-9310(73)90236-6)

- H. R. Skullerud. 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 (1968), 1567–1568. DOI: <https://doi.org/10.1088/0022-3727/1/11/423>
- J. Spanier and E. M. Gelbard. 1969. *Monte carlo principles and neutron transport problems*. Addison-Wesley Pub. Co.
- N. M. Steen. 1966. *A simple method to improve the efficiency of the Σ_a/Σ_t estimator in certain Monte Carlo programs*. Technical Report WAPD-TM-609. Bettis Atomic Power Lab., Pittsburgh, PA, USA.
- T. M. Sutton, F. B. Brown, F. G. Bischoff, D. B. MacMillan, C. L. Ellis, J. T. Ward, C. T. Ballinger, D. J. Kelly, and L. Schindler. 1999. *The physical models and statistical procedures used in the RACER Monte Carlo Code*. Technical Report KAPL-4840. Knolls Atomic Power Laboratory, Niskayuna, NY, USA. DOI: <https://doi.org/10.2172/767449>
- L. Szirmay-Kalos, B. Tóth, and M. Magdics. 2011. Free path sampling in high resolution inhomogeneous participating media. *Computer Graphics Forum* 30, 1 (2011), 85–97. DOI: <https://doi.org/10.1111/j.1467-8659.2010.01831.x>
- E. Veach. 1997. *Robust Monte Carlo methods for light transport simulation*. Ph.D. Dissertation. Stanford University, Stanford, CA, USA.
- J. von Neumann. 1951. Various techniques used in connection with random digits. *Journal of Research of the National Bureau of Standards, Appl. Math. Series 12* (1951), 36–38.
- A. Wilkie, S. Nawaz, M. Droske, A. Weidlich, and J. Hanika. 2014. Hero wavelength spectral sampling. *CGF (Proc. of Eurographics Symposium on Rendering)* 33, 4 (June 2014), 123–131. DOI: <https://doi.org/10.1111/cgf.12419>
- E. R. Woodcock, T. Murphy, P. J. Hemmings, and T. C. Longworth. 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.
- Y. Yue, K. Iwasaki, B. Chen, Y. Dobashi, and T. Nishita. 2010. Unbiased, adaptive stochastic sampling for rendering inhomogeneous participating media. *ACM TOG (Proc. of SIGGRAPH Asia)* 29, 6 (Dec. 2010), 177:1–177:8. DOI: <https://doi.org/10.1145/1866158.1866199>
- Y. Yue, K. Iwasaki, B. Chen, Y. Dobashi, and T. Nishita. 2011. Toward optimal space partitioning for unbiased, adaptive free path sampling of inhomogeneous participating media. *CGF (Proc. of Pacific Graphics)* 30, 7 (2011), 1911–1919. DOI: <https://doi.org/10.1111/j.1467-8659.2011.02049.x>
- C. D. Zerby, R. B. Curtis, and H. W. Bertini. 1961. *The relativistic doppler problem*. Technical Report ORNL-61-7-20. Oak Ridge National Laboratory, Oak Ridge, TN, USA. DOI: <https://doi.org/10.2172/4836227>

A FREE-PATH-SAMPLING RESIDUAL RATIO TRACKING

Novák et al. [2014] briefly discuss a potential free-path-sampling technique based on the mechanisms underlying residual ratio tracking. The authors suggest sampling free paths analytically from a representative control volume and then computing the sample weight using ratio tracking. We observe that this algorithm is a special case of weighted delta tracking with the configuration

$$\bar{\mu}(\mathbf{x}) = \mu_t^c(\mathbf{x}) + \bar{\mu}^r(\mathbf{x}), \quad (34)$$

$$P_t(\mathbf{x}) = \frac{\mu_t^c(\mathbf{x})}{\bar{\mu}(\mathbf{x})}, \quad P_n(\mathbf{x}) = \frac{\bar{\mu}^r(\mathbf{x})}{\bar{\mu}(\mathbf{x})}, \quad (35)$$

where $P_t(\mathbf{x})$ is the probability of sampling a real collision (absorption or scattering). As the authors describe, $\mu_t^c(\mathbf{x})$ and $\bar{\mu}^r(\mathbf{x})$ must allow analytic free-path sampling, and would in practice be set to constants, such as $\text{avg}(\mu_t(\mathbf{x}))$ and $\max(|\mu_t^r(\mathbf{x})|)$, respectively. Since the distribution of sampled free paths is fully determined by the relatively simple $\mu_t^c(\mathbf{x})$, it can differ drastically from the true free-path distribution, frequently resulting in extreme weights.

B INTEGRAL FORMULATION OF TRANSMITTANCE ESTIMATORS

In order to derive the integral formulation of transmittance trackers, we first modify the RTE to estimate only transmittance (i.e. uncollided radiance). We rewrite Equation (3) to account only for losses

and a point source of unit radiance at \mathbf{x}_d , $\mathbf{x}_d = \mathbf{x}_0 - d\omega$:

$$(\omega \cdot \nabla)L(\mathbf{x}, \omega) = -[\mu_t(\mathbf{x}) + \mu_n(\mathbf{x})]L(\mathbf{x}, \omega) + \delta(\mathbf{x} - \mathbf{x}_d) + \mu_n(\mathbf{x}) \int_{S^2} \delta(\omega - \bar{\omega})L(\mathbf{x}, \bar{\omega}) d\bar{\omega}. \quad (36)$$

Integrating both sides along ω ,

$$\begin{aligned} L(\mathbf{x}, \omega) &= \int_0^\infty \exp\left(-\int_0^t \bar{\mu}(\mathbf{x}_s) ds\right) \\ &\quad \times \left[\delta(\mathbf{x}_t - \mathbf{x}_d) + \mu_n(\mathbf{x}_t)L(\mathbf{x}_t, \omega)\right] dt \\ &= \int_0^\infty \exp\left(-\int_0^t \bar{\mu}(\mathbf{x}_s) ds\right) \delta(\mathbf{x}_t - \mathbf{x}_d) dt \\ &\quad + \int_0^\infty \exp\left(-\int_0^t \bar{\mu}(\mathbf{x}_s) ds\right) \mu_n(\mathbf{x}_t)L(\mathbf{x}_t, \omega) dt, \end{aligned} \quad (37)$$

solving the Dirac integral, noting that we only need to integrate up to the remaining distance d_i to the point source ($L(\mathbf{x}, \omega) = 0$ for $t > d$), and introducing $p(t)$ yields

$$L(\mathbf{x}, \omega) = \exp\left(-\int_0^{d_i} \bar{\mu}(\mathbf{x}_s) ds\right) + \int_0^{d_i} p(t) \frac{\mu_n(\mathbf{x}_t)}{\bar{\mu}(\mathbf{x}_t)} L(\mathbf{x}_t, \omega) dt, \quad (38)$$

which represents the *weighted next-flight estimator* presented by Cramer [1978]. At each collision point, the estimator performs a next-flight estimation by adding an underestimated uncollided contribution—the transmittance through the rest of the *combined* medium (first term)—and then proceeds with sampling null collisions to correct for the initial underestimation (second term). The Monte Carlo score can be written as

$$w(\mathbf{X}) = \sum_{i=0}^{k-1} \exp\left(-\int_0^{d_i} \bar{\mu}(\mathbf{x}_s) ds\right) \prod_{j=1}^i \frac{\mu_n(\mathbf{x}_j)}{\bar{\mu}(\mathbf{x}_j)}, \quad (39)$$

where k is the index of the first collision after exceeding distance d from \mathbf{x}_0 . To the best of our knowledge, this is the first derivation of the next-flight estimator directly from a modified version of the RTE, and its first description in computer graphics literature.

Cramer [1978] also analyzed a simpler version, which removes the next-flight estimation and tracks only the product of relative concentrations of null particles. By observing that the transmittance from 0 to d_i is equivalent to the integral of $p(t)$ from d_i to ∞ , we can rewrite Equation (38) to

$$L(\mathbf{x}, \omega) = \int_0^{d_i} p(t) \frac{\mu_n(\mathbf{x}_t)}{\bar{\mu}(\mathbf{x}_t)} L(\mathbf{x}_t, \omega) dt + \int_{d_i}^\infty p(t) dt. \quad (40)$$

The two integrals in Equation (40) each contain the same PDF and have disjoint domains whose union is the whole domain of the PDF, leading to a natural translation into a Monte Carlo estimator. In particular, we sample a distance t according to $p(t)$, then return the remaining factors of the corresponding integrand: $\frac{\mu_n(\mathbf{x}_t)}{\bar{\mu}(\mathbf{x}_t)} L(\mathbf{x}_t, \omega)$ if $t < d_i$ and 1 otherwise. The Monte Carlo score simplifies to

$$w(\mathbf{X}) = \prod_{j=1}^{k-1} \frac{\mu_n(\mathbf{x}_j)}{\bar{\mu}(\mathbf{x}_j)}. \quad (41)$$

In the context of rendering, this algorithm was reinvented and studied by Novák et al. [2014] under the name *ratio tracking*.