

A null-scattering path integral formulation of light transport

BAILEY MILLER*, Dartmouth College, USA
 ILIYAN GEORGIEV*, Autodesk, United Kingdom
 WOJCIECH JAROSZ, Dartmouth College, USA




Fig. 1. Unbiased rendering of spectrally and spatially varying participating media could previously be accomplished using delta tracking separately for each color component (leftmost inset), but this leads to strong color noise. Spectral tracking [Kutz et al. 2017] can reduce this noise by rendering all color components together (middle left inset), but at the cost of sampling distances based on the densest color component of the medium. Our theoretical framework allows us to leverage different sampling techniques across color components (middle right), or exploit next-event estimation (NEE) (far right), and combine these into a more robust, lower-variance estimator via multiple importance sampling (MIS). See Tables 1 and 2 for descriptions of the methods and the medium.

Unbiased rendering of general, heterogeneous participating media currently requires using null-collision approaches for estimating transmittance and generating free-flight distances. A long-standing limitation of these approaches, however, is that the corresponding path pdfs cannot be computed due to the black-box nature of the null-collision rejection sampling process. These techniques therefore cannot be combined with other sampling techniques via multiple importance sampling (MIS), which significantly limits their robustness and generality. Recently, Galtier et al. [2013] showed how to derive these algorithms directly from the radiative transfer equation (RTE). We build off this generalized RTE to derive a *path integral* formulation of null scattering, which reveals the sampling pdfs and allows us to devise new,

express existing, and combine complementary unbiased techniques via MIS. We demonstrate the practicality of our theory by combining, for the first time, several path sampling techniques in spatially and spectrally varying media, generalizing and outperforming the prior state of the art.

CCS Concepts: • **Computing methodologies** → **Ray tracing**.

Additional Key Words and Phrases: global illumination, light transport, participating media, null scattering, Monte Carlo integration

ACM Reference Format:

Bailey Miller, Iliyan Georgiev, and Wojciech Jarosz. 2019. A null-scattering path integral formulation of light transport. *ACM Trans. Graph.* 38, 4, Article 44 (July 2019), 13 pages. <https://doi.org/10.1145/3306346.3323025>

1 INTRODUCTION

The world around us is filled with participating media which volumetrically attenuates and scatters light as it travels from light sources to our eyes. While important in many fields, simulating this transport efficiently and accurately is unfortunately a notoriously difficult problem, since it requires solving not only the rendering equation [Kajiya 1986; Immel et al. 1986], but also its volumetric generalization, the radiative transfer equation [Chandrasekhar 1960].

Monte Carlo path sampling methods such as (bidirectional) path tracing [Kajiya 1986; Lafortune and Willems 1993; Veach and Guibas 1995] and its volumetric variants [Lafortune and Willems 1996; Georgiev et al. 2013] have been investigated in academia for decades

* Authors with equal contribution.

Authors' addresses: Bailey Miller, bailey.m.miller@dartmouth.edu, Dartmouth College, 9 Maynard Street, Hanover, NH, 03755, USA; Iliyan Georgiev, iliyan.georgiev@autodesk.com, Autodesk, 17 Broadwick Street, London, W1F 0DE, United Kingdom; Wojciech Jarosz, wojciech.k.jarosz@dartmouth.edu, Dartmouth College, 9 Maynard Street, Hanover, NH, 03755, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0730-0301/2019/7-ART44 \$15.00

<https://doi.org/10.1145/3306346.3323025>

due to their elegant simplicity, generality, and accuracy. Moreover, by now most major film production renderers have adopted such approaches as their dominant rendering algorithms [Georgiev et al. 2018; Burley et al. 2018; Fascione et al. 2018; Christensen et al. 2018; Christensen and Jarosz 2016; Fascione et al. 2017].

These approaches operate by stochastically constructing light transport paths between sensors and emitters. Thanks to decades of research [Novák et al. 2018], we now have an arsenal of strategies for constructing such paths in participating media [Raab et al. 2008; Kulla and Fajardo 2012; Georgiev et al. 2013; Krivánek et al. 2014]. One core strength of such methods is the ability to leverage *multiple importance sampling* (MIS) [Veach and Guibas 1995] to combine complementary sampling techniques into a single, provably good estimator with improved robustness. MIS operates by weighting each technique proportionally to its path sampling pdf.

Unfortunately, combining multiple unbiased sampling techniques in spatially and spectrally varying participating media has remained elusive. Such media require the use of so-called *null-collision* approaches to estimate transmittance and generate light transport paths. These approaches, also called “delta scattering” or “Woodcock tracking”, were initially developed and analyzed in the 1950s and 1960s [Butcher and Messel 1958, 1960; Zerby et al. 1961; Bertini 1963; Woodcock et al. 1965; Miller 1967; Coleman 1968]. They were only recently adopted [Raab et al. 2008] and extended upon in graphics [Jarosz et al. 2011b; Novák et al. 2014; Kutz et al. 2017; Szirmay-Kalos et al. 2011, 2017, 2018]. A derivation of these algorithms directly from the RTE [Galtier et al. 2013] was recently introduced to graphics [Kutz et al. 2017; Novák et al. 2018], demonstrating not only their correctness, but also providing a convenient framework for postulating new variants. Unfortunately, the authors also noted that “the most limiting drawback of these methods is their inability to quantify the pdf of individual samples”. Since MIS requires access to these pdfs, it has so far been limited to (piecewise) homogeneous media [Krivánek et al. 2014; Georgiev et al. 2013; Wilkie et al. 2014]. Heterogeneous media must either forego MIS [Kutz et al. 2017] or forego null-collision methods in favor of regular tracking [Sutton et al. 1999; Wilkie et al. 2014; Fascione et al. 2018] or tabulated sampling [Szirmay-Kalos et al. 2017; Gamito 2018] to support MIS.

We solve this long-standing problem by turning the null-scattering RTE into a generalized null-scattering path integral formulation of volumetric light transport. Our formulation allows unbiased path sampling and *analytic pdf evaluation* in spatially and spectrally varying media for the first time. This enables us to generalize, improve the performance of, or remove bias from recent techniques like spectral decomposition tracking [Kutz et al. 2017] and hero wavelength sampling [Wilkie et al. 2014], all while enabling new techniques and allowing their combination with complementary existing techniques like equiangular sampling [Kulla and Fajardo 2012] and ratio tracking [Novák et al. 2014] via MIS, as we demonstrate in Fig. 1.

2 RADIATIVE TRANSPORT BACKGROUND

We begin by reviewing the classic formulation of the steady-state radiance distribution in scenes containing surfaces and participating media. We then present the null-scattering extension [Galtier et al. 2013] of this recursive formulation, which we will convert to a path integral in Section 3.

2.1 Classical formulation

In scenes containing media with particles distributed statistically independently from each other, the radiance equilibrium is described by the *radiative transfer equation* (RTE) [Chandrasekhar 1960]

$$(\omega \cdot \nabla) L(\mathbf{x}, \omega) = \underbrace{-\mu_t(\mathbf{x})L(\mathbf{x}, \omega)}_{\text{losses}} + \underbrace{\mu_t(\mathbf{x})L^m(\mathbf{x}, \omega)}_{\text{gains}}, \quad (1)$$

which relates the differential change in the radiance L at a point $\mathbf{x}$ in direction ω to losses from absorption and out-scattering and gains from self-emission L_e^m and in-scattering L^m :

$$L^m(\mathbf{x}, \omega) = \frac{\mu_a(\mathbf{x})}{\mu_t(\mathbf{x})} L_e^m(\mathbf{x}, \omega) + \frac{\mu_s(\mathbf{x})}{\mu_t(\mathbf{x})} \int_{S^2} \rho_m(\omega, \mathbf{x}, \omega') L(\mathbf{x}, \omega') d\omega'. \quad (2)$$

Here S^2 is the unit sphere, ρ_m is the medium *phase function*, and $\mu_t(\mathbf{x}) = \mu_a(\mathbf{x}) + \mu_s(\mathbf{x})$ is the medium *extinction coefficient*, where μ_a and μ_s are the *absorption* and *scattering coefficients*, respectively.

Solving this ordinary linear differential equation for the radiance $L(\mathbf{x}, \omega)$, and using the outgoing radiance $L^s(\mathbf{z}, \omega)$ at the first visible surface point $\mathbf{z} = \mathbf{x} - z\omega$ as a boundary condition, yields the *volume rendering equation* (VRE) [Arvo 1993]:

$$L(\mathbf{x}, \omega) = \int_0^z T(\mathbf{x}, \mathbf{y}) \mu_t(\mathbf{y}) L^m(\mathbf{y}, \omega) d\mathbf{y} + T(\mathbf{x}, \mathbf{z}) L^s(\mathbf{z}, \omega), \quad (3)$$

where the medium radiance L^m is integrated at points $\mathbf{y} = \mathbf{x} - y\omega$ along the line, and like L^m (2), L^s is a sum of self-emitted and in-scattered radiance [Kajiya 1986; Immel et al. 1986]:

$$L^s(\mathbf{z}, \omega) = L_e^s(\mathbf{z}, \omega) + \int_{S^2} \rho_s(\omega, \mathbf{z}, -\omega') L(\mathbf{z}, \omega') |n(\mathbf{z}) \cdot \omega'| d\omega'. \quad (4)$$

Here ρ_s is the *bidirectional scattering distribution function* (BSDF) and $n(\mathbf{z})$ is the surface normal at $\mathbf{z}$. Finally, the medium *extinction transmittance* is defined as

$$T(\mathbf{x}, \mathbf{y}) = e^{-\tau(\mathbf{x}, \mathbf{y})} = e^{-\int_0^y \mu_t(\mathbf{x} - s\omega) ds}, \quad (5)$$

where $\tau(\mathbf{x}, \mathbf{y})$ is the medium *extinction optical thickness* between $\mathbf{x}$ and $\mathbf{y}$. Note that there is one VRE (3) for each color component, i.e. wavelength or red, green, and blue (R, G, B) channel.

To compute the radiance at a point, the VRE can be estimated recursively via Monte Carlo integration. The traditional approach, known as analog sampling, is to importance sample the non-recursive terms under the integrals in Eqs. (2) to (4) [Novák et al. 2018]. For Eq. (3) this involves sampling a distance y with density $T(\mathbf{x}, \mathbf{y}) \mu_t(\mathbf{y})$. This, however, is possible only when the corresponding cumulative distribution function is invertible, e.g. in media with (piece-wise) simple density (see Fig. 2, left). In general heterogeneous media, where analytic transmittance sampling is unavailable, other distance sampling techniques can be used, e.g. equiangular sampling [Kulla and Fajardo 2012]. This, however, requires the explicit evaluation of the transmittance $T(\mathbf{x}, \mathbf{y})$, which is also unavailable in closed form. Numerical estimation and sampling of the optical thickness $\tau(\mathbf{x}, \mathbf{y})$, e.g. via ray marching [Perlin and Hoffert 1989], even if unbiased, leads to a biased estimate of $T(\mathbf{x}, \mathbf{y})$ [Raab et al. 2008].

The only known class of techniques for unbiased light transport estimation in heterogeneous media are the so-called *null-scattering methods*. We will show that these operate in a higher-dimensional sampling space and are most naturally expressed as Monte Carlo estimators of a corresponding extension of the integral VRE (3).

2.2 Null-scattering formulation

The idea of the null-scattering formulation is to alter the medium density in a way that preserves the radiance equilibrium while enabling the use of analytical sampling techniques to estimate the equilibrium in an unbiased manner. Conceptually, this is achieved by introducing fictitious matter with (signed) density μ_n that, upon collision, scatters light forward with unchanged direction and intensity; hence the term “null scattering”. Formally, we introduce the *combined extinction* medium coefficient

$$\bar{\mu}_t(\mathbf{x}) = \mu_t(\mathbf{x}) + \mu_n(\mathbf{x}) \quad (6)$$

to write a null-scattering extension to the RTE (1):

$$(\omega \cdot \nabla) L(\mathbf{x}, \omega) = -\bar{\mu}_t(\mathbf{x})L(\mathbf{x}, \omega) + \bar{\mu}_t(\mathbf{x})\bar{L}^{\bar{m}}(\mathbf{x}, \omega). \quad (7)$$

Here, the *combined medium radiance* is

$$\bar{L}^{\bar{m}}(\mathbf{x}, \omega) = \quad (8)$$

$$\underbrace{\frac{\mu_a(\mathbf{x})}{\bar{\mu}_t(\mathbf{x})} L_e^m(\mathbf{x}, \omega)}_{\text{emission}} + \underbrace{\frac{\mu_s(\mathbf{x})}{\bar{\mu}_t(\mathbf{x})} \int_{S^2} \rho_m(\omega, \mathbf{x}, \omega') L(\mathbf{x}, \omega') d\omega'}_{\text{real scattering}} + \underbrace{\frac{\mu_n(\mathbf{x})}{\bar{\mu}_t(\mathbf{x})} L(\mathbf{x}, \omega)}_{\text{null scattering}},$$

where the individual contributions are weighted by corresponding combined-medium albedos and now include null-scattered radiance.

Solving the null-scattering RTE (7) for $L(\mathbf{x}, \omega)$, and using the same boundary condition as for Eq. (3), we obtain a null-scattering extension to the VRE:

$$L(\mathbf{x}, \omega) = \int_0^z \bar{T}(\mathbf{x}, \mathbf{y}) \bar{\mu}_t(\mathbf{y}) \bar{L}^{\bar{m}}(\mathbf{y}, \omega) d\mathbf{y} + \bar{T}(\mathbf{x}, \mathbf{z}) L^s(\mathbf{z}, \omega). \quad (9)$$


The null-scattering VRE (9) has the same basic structure as the classic VRE (3) but has one major advantage—the spatially varying free parameter μ_n . This parameter can be used to transform the *combined transmittance*

$$\bar{T}(\mathbf{x}, \mathbf{y}) = e^{-\bar{T}(\mathbf{x}, \mathbf{y})} = e^{-\int_0^y \bar{\mu}_t(\mathbf{x}-s\omega) ds} \quad (10)$$

into a form suitable for analytic sampling and evaluation, e.g. by setting $\mu_n(\mathbf{x}) = C - \mu_t(\mathbf{x})$, where C is a constant, effectively “homogenizing” the combined extinction $\bar{\mu}_t$. This makes the VRE (9) suitable for direct recursive Monte Carlo estimation. Existing null-scattering methods, such as delta tracking, can be formulated as estimators for that equation, performing a series of distance sampling decisions through null collisions, followed by direction sampling at a real-scattering collision. We refer to Novák et al. [2018] for an overview of such methods, and illustrate delta tracking in Fig. 2, right.

2.3 Discussion

Like the classical formulation, the null-scattering VRE is restricted to pure unidirectional estimation, e.g. via delta or spectral tracking [Kutz et al. 2017]. The path integral formulation of Pauly et al. [2000] enables next-event and bidirectional estimation and, importantly, their MIS combination. However, that formulation does not encompass null-scattering techniques as it is derived from the classical VRE. Specialized formulations exist for techniques such as ratio tracking [Kutz et al. 2017, Appendix B], but their utility is limited as they do not encompass other techniques either, e.g. equiangular sampling [Kulla and Fajardo 2012]. In the following section, we present a null-scattering path integral formulation that fills this gap.






Fig. 3. Illustration of a light transport path in our integral formulation, which explicitly considers both real and null scattering at path vertices. Prior formulations consider real scattering only, where this length-8 path corresponds to the length-4 path $\mathbf{x}_{i_0} \mathbf{x}_{r_1} \mathbf{x}_{r_2} \mathbf{x}_{r_3} \mathbf{x}_{r_4}$.

connecting every two consecutive vertices, which does not permit analytic evaluation or sampling in general heterogeneous media.

To address this inconvenience, we derive a path integral expression with the same form as Eq. (12) but starting from the null-scattering VRE (9). In contrast to the formulation of Pauly et al. [2000], our formulation considers both real and null scattering at path vertices, and replaces the extinction transmittance T by the analytically evaluable combined transmittance $\bar{T}$ (10).

Path space and measure. To properly handle the geometry of null scattering, we isolate such events into a *null-scattering volume* $\mathcal{V}_\delta$ which is simply a copy of $\mathcal{V}$. This extends the traditional path space and its corresponding differential measure to

$$\mathcal{P} = \bigcup_{k=1}^{\infty} (\mathcal{A} \cup \mathcal{V} \cup \mathcal{V}_\delta)^{k+1}, \quad d\bar{\mathbf{x}} = \prod_{i=0}^k d\mathbf{x}_i, \quad d\mathbf{x}_i = \begin{cases} dA(\mathbf{x}_i), & \text{if } \mathbf{x}_i \in \mathcal{A}, \\ dV(\mathbf{x}_i), & \text{if } \mathbf{x}_i \in \mathcal{V}, \\ dV_\delta(\mathbf{x}_i), & \text{if } \mathbf{x}_i \in \mathcal{V}_\delta. \end{cases} \quad (13)$$

Null-scattering vertices $\mathbf{x}_i$ are measured along the line connecting the preceding and succeeding real scattering vertices:

$$dV_\delta(\mathbf{x}_i) = d\delta_{\mathbf{x}_i^- \leftrightarrow \mathbf{x}_i^+}(\mathbf{x}_i), \quad (14)$$

where $\delta_{\mathbf{x}_i^- \leftrightarrow \mathbf{x}_i^+}(\mathbf{x}_i)$ is a Dirac measure restricting the integration along the line segment connecting the preceding and succeeding real-scattering vertices $\mathbf{x}_i^-$ and $\mathbf{x}_i^+$, respectively. The path length k is the number of segments between consecutive scattering events of any kind. Figure 3 illustrates a path of length 8 in the space $\mathcal{P}$.

Measurement contribution. For each path length k , the *measurement contribution function* is defined as

$$f(\bar{\mathbf{x}}) = W_e(\mathbf{x}_0, \omega_{\mathbf{x}_1 \mathbf{x}_0}) \cdot \left(\prod_{i=0}^{r-1} G(\mathbf{x}_{r_i}, \mathbf{x}_{r_{i+1}}) \right) \cdot \left(\prod_{i=0}^{k-1} \bar{T}(\mathbf{x}_i, \mathbf{x}_{i+1}) \right) \cdot \left(\prod_{i=1}^{k-1} \rho(\omega_{\mathbf{x}_i \mathbf{x}_{i-1}}, \mathbf{x}_i, \omega_{\mathbf{x}_{i+1} \mathbf{x}_i}) \right) \cdot L_e(\mathbf{x}_k, \omega_{\mathbf{x}_k \mathbf{x}_{k-1}}), \quad (15)$$

where $\omega_{\mathbf{x}\mathbf{y}}$ is the (unit) direction from $\mathbf{x}$ to $\mathbf{y}$. The $r+1$ real-scattering vertices on the path are indexed by r_i , with $\mathbf{x}_0 \equiv \mathbf{x}_{r_0}$ and $\mathbf{x}_k \equiv \mathbf{x}_{r_r}$ being the camera and light vertices, respectively, which we regard as real-scattering (see Fig. 3). The combined transmittance $\bar{T}(\mathbf{x}, \mathbf{y})$




Fig. 4. The classical path integral formulation (top) considers only real scattering and thus has to evaluate the extinction transmittance T between such events. Our formulation (bottom) instead evaluates the combined transmittance $\bar{T}$ as it considers null-scattering events explicitly.

is given by Eq. (10), and

$$G(\mathbf{x}, \mathbf{y}) = \frac{D(\mathbf{x}, \omega_{\mathbf{x}\mathbf{y}}) V(\mathbf{x}, \mathbf{y}) D(\mathbf{y}, \omega_{\mathbf{y}\mathbf{x}})}{\|\mathbf{x} - \mathbf{y}\|^2} \quad (16)$$

$$D(\mathbf{x}, \omega) = \begin{cases} |n(\mathbf{x}) \cdot \omega|, & \text{if } \mathbf{x} \in \mathcal{A}, \\ 1, & \text{if } \mathbf{x} \in \mathcal{V} \end{cases} \quad (17)$$

$$L_e(\mathbf{x}, \omega) = \begin{cases} L_e(\mathbf{x}, \omega_{\mathbf{x}\mathbf{y}}), & \text{if } \mathbf{x} \in \mathcal{A}, \\ \mu_a(\mathbf{x}) L_e(\mathbf{x}, \omega_{\mathbf{x}\mathbf{y}}), & \text{if } \mathbf{x} \in \mathcal{V} \end{cases} \quad (18)$$

$$\rho(\omega, \mathbf{x}, \omega') = \begin{cases} \rho_s(\omega, \mathbf{x}, -\omega'), & \text{if } \mathbf{y} \in \mathcal{A}, \\ \mu_s(\mathbf{x}) \rho_m(\omega, \mathbf{x}, \omega'), & \text{if } \mathbf{x} \in \mathcal{V}, \\ \mu_n(\mathbf{x}) H(\omega \cdot \omega'), & \text{if } \mathbf{x} \in \mathcal{V}_\delta, \end{cases} \quad (19)$$

where $V(\mathbf{x}, \mathbf{y})$ is the binary visibility function between $\mathbf{x}$ and $\mathbf{y}$. In contrast to prior definitions, our generalized scattering term ρ explicitly considers null scattering, where H is the heaviside function which enforces the ordering of the null vertices¹.

Note that the geometry term $G(\mathbf{x}, \mathbf{y})$ is evaluated only between real-scattering events $\mathbf{x}, \mathbf{y} \in \mathcal{A} \cup \mathcal{V}$. Null-scattering vertices are constrained to lie on the polyline connecting real-scattering events which is effectively a path-space manifold. This is similar to the manifolds studied by Jakob and Marschner [2012] and the changes in path density through chains of specular (i.e. delta) surface reflection and refraction between two scattering events. In our case, the geometry term through a null-scattering chain has a simple form.

3.2 Discussion

By explicitly accounting for null scattering, the problematic extinction transmittance term T is replaced by the combined transmittance $\bar{T}$ in our formulation. The main advantage is that $\bar{T}$ can be made analytically evaluable by appropriately setting the null-scattering density parameter μ_n . Conversely, setting $\mu_n \equiv 0$ makes $\bar{T} = T$, reducing our formulation to the classical one. This is useful for media with analytically evaluable extinction optical thickness τ , e.g. with homogeneous density, as it allows only real scattering to be sampled. Figure 4 highlights the differences between the two formulations.

Lastly, note that in general there is one path integral (12) for every color component. A different null density μ_n can be used for each.

¹Alternatively, this ordering can be encoded into the null-vertex measure by using $dV_\delta(\mathbf{x}_i) = d\delta_{\mathbf{x}_{i-1} \leftrightarrow \mathbf{x}_{i+1}}(\mathbf{x}_i)$ in Eq. (13), eliminating the need for the heaviside function in Eq. (19).

4 NULL-SCATTERING PATH INTEGRAL DERIVATION

In this section we derive our path integral formulation (12) from the null-scattering VRE (9). We first expand the recursions in the VRE, followed by a change of variables in the resulting high-dimensional integrals, which we ultimately merge into one path-space integral. These are the same general steps done in derivations based on the real-scattering VRE (3) [Pauly et al. 2000; Jakob 2013]. However, in our case the added null-scattering recursion in Eq. (9) increases the complexity of the expansion, and the resulting null-scattering integrals require an appropriate change of variables. Readers not interested in these technical details may skip over to Section 5, where we discuss the practical applications of our formulation.

We begin by writing the null-scattering VRE in a compact form:

$$L(\mathbf{x}, \omega) = \int_0^z \bar{T}(\mathbf{x}, \mathbf{y}) L_o(\mathbf{y}, \omega) d\mathbf{y} + \bar{T}(\mathbf{x}, \mathbf{z}) L_o(\mathbf{z}, \omega) \quad (20)$$

$$L_o(\mathbf{x}, \omega) = L_e(\mathbf{x}, \omega) + \int_{S^2} \rho(\omega, \mathbf{x}, \omega') L(\mathbf{x}, \omega') D(\mathbf{x}, \omega') d\omega' + \mu_n(\mathbf{x}) L(\mathbf{x}, \omega),$$

where $\bar{\mu}_t$ from Eqs. (8) and (9) cancels out and where we use the notation from Eqs. (17) to (19) to express the contributions from medium points $\mathbf{y}$ and surface points $\mathbf{z} = \mathbf{x} - z\omega$ in $L(\mathbf{x}, \omega)$ using a common *outgoing radiance* term L_o . We do not consider null scattering (i.e. transparency) at surfaces, thus $\mu_n(\mathbf{x}) = 0$ for $\mathbf{x} \in \mathcal{A}$.

4.1 Operator formulation

As in prior formulations [Pauly et al. 2000; Jakob 2013], we will express the pixel measurement I as a path integral by recursively expanding the radiance L (20) in Eq. (11). To express this expansion succinctly, we make use of linear operators [Veach 1997; Arvo 1995]. Substituting L into L_o replaces the last two terms of L_o by four new terms, from which we extract four operators:

$$(\mathbf{R}_m h)(\mathbf{x}, \omega) = \int_{S^2} \int_0^z \rho(\omega, \mathbf{x}, \omega') \bar{T}(\mathbf{x}, \mathbf{y}) D(\mathbf{x}, \omega') h(\mathbf{y}, \omega') d\mathbf{y} d\omega' \quad (21)$$

$$(\mathbf{R}_s h)(\mathbf{x}, \omega) = \int_{S^2} \rho(\omega, \mathbf{x}, \omega') \bar{T}(\mathbf{x}, \mathbf{z}) D(\mathbf{x}, \omega') h(\mathbf{z}, \omega') d\omega' \quad (22)$$

$$(\mathbf{N}_m h)(\mathbf{x}, \omega) = \mu_n(\mathbf{x}) \int_0^z \bar{T}(\mathbf{x}, \mathbf{y}) h(\mathbf{y}, \omega) d\mathbf{y} \quad (23)$$

$$(\mathbf{N}_s h)(\mathbf{x}, \omega) = \mu_n(\mathbf{x}) \bar{T}(\mathbf{x}, \mathbf{z}) h(\mathbf{z}, \omega). \quad (24)$$

We then define the real- and null-scattering operators, respectively

$$(\mathbf{R}h)(\mathbf{x}, \omega) = ((\mathbf{R}_m + \mathbf{R}_s)h)(\mathbf{x}, \omega), \quad (\mathbf{N}h)(\mathbf{x}, \omega) = ((\mathbf{N}_m + \mathbf{N}_s)h)(\mathbf{x}, \omega). \quad (25)$$

Note that, like L_o , $\mathbf{R}$ operates on both medium and surface points $\mathbf{x}$, while $\mathbf{N}$ is non-zero only at medium points.

Radiance equilibrium. We can now write L_o in operator form in terms of $\mathbf{R}$ and $\mathbf{N}$, and then expand the tail recursion:

$$L_o = L_e + (\mathbf{R} + \mathbf{N})L_o = \sum_{k=0}^{\infty} (\mathbf{R} + \mathbf{N})^k L_e = \sum_{k=0}^{\infty} \sum_{S_k \in \{\mathbf{R}, \mathbf{N}\}^k} S_k L_e. \quad (26)$$

The result is an expression for the outgoing radiance as a sum of emitted radiance scattered arbitrarily many times. On the right-hand side we have expanded $(\mathbf{R} + \mathbf{N})^k$ into a sum of 2^k composite operators S_k , each representing one possible series of k (real or null) scattering events. An example of such operator is $S_5 = \mathbf{RNRNN}$, where the first two events (in direction of light flow) are null scattering, followed by real scattering, and so on.

Pixel measurement. Our next step is to write the pixel measurement (11) in non-recursive operator form, as a sum of nested integrals. To that end, including W_e temporarily in the definition of the scattering function ρ (19) allows us to treat the spherical integral in Eq. (11) as a real-scattering event, such that expanding L in Eq. (11) using Eq. (20) and then expressing L_o using Eq. (26) yields

$$I = \sum_{k=0}^{\infty} \sum_{S_k \in \{\mathbf{R}, \mathbf{N}\}^k} \int_{\mathcal{A}} (\mathbf{R}S_k L_e)(\mathbf{x}, \cdot) d\mathbf{x} = \sum_{\mathbf{P} \in \Omega_P} \int_{\mathcal{A}} (\mathbf{P}L_e)(\mathbf{x}, \cdot) d\mathbf{x}, \quad (27)$$

where the operator under the integral is evaluated with a dummy direction. The set Ω_P includes all *path operators* of the form $\mathbf{P} = \mathbf{R}S_k$, where the “camera” real-scattering event is followed (in direction opposite of the light flow) by k scene scattering events of any type.

4.2 Scattering chain decomposition

To express the pixel measurement (27) as integration over a product (path) space of surface area and volume, we need to perform an appropriate change of variables in every path operator $\mathbf{P} = \mathbf{R}S_k \in \Omega_P$. For every k there are 2^k possible operators $\mathbf{P}$, each corresponding to a different k -sequence of real- and null-scattering events. To handle this combinatorial explosion, we make the key observation that every such operator can be written as a sequence of scattering chains, each starting with a real-scattering event:

$$\mathbf{P} = \mathbf{R} \overbrace{\mathbf{N} \dots \mathbf{N}}^{n_1 \text{ times}} \overbrace{\mathbf{R} \mathbf{N} \dots \mathbf{N}}^{n_2 \text{ times}} \dots \overbrace{\mathbf{R} \mathbf{N} \dots \mathbf{N}}^{n_r \text{ times}} = \mathbf{R} \mathbf{N}^{n_1} \mathbf{R} \mathbf{N}^{n_2} \dots \mathbf{R} \mathbf{N}^{n_r}, \quad (28)$$

where the number of real-scattering events is r and the total number of null-scattering events is $\sum_{i=1}^r n_i = k - r$, with $n_i \geq 0$. It thus suffices to find the change of variables for a general chain $\mathbf{R} \mathbf{N}^n$, which can then be applied to each chain in every path operator $\mathbf{P}$.

In the absence of null scattering, i.e. when $n = 0$, the chain operator $\mathbf{R} \mathbf{N}^n$ simplifies to $\mathbf{R}$; we will address this special case in Section 4.3 below. When $n > 0$, assuming no null scattering at surfaces, we can expand $\mathbf{R} \mathbf{N}^n$ using Eq. (25):

$$(\mathbf{R} \mathbf{N}^n h)(\mathbf{x}, \omega) = (\mathbf{R}_m \mathbf{N}^n h)(\mathbf{x}, \omega) + (\mathbf{R}_s \mathbf{N}^n h)(\mathbf{x}, \omega) \quad (29)$$

$$= (\mathbf{R}_m (\mathbf{N}_m + \mathbf{N}_s)^n h)(\mathbf{x}, \omega) = (\mathbf{R}_m (\mathbf{N}_m + \mathbf{N}_s)^{n-1} (\mathbf{N}_m + \mathbf{N}_s) h)(\mathbf{x}, \omega)$$

$$= (\mathbf{R}_m \mathbf{N}_m^{n-1} h)(\mathbf{x}, \omega) + (\mathbf{R}_m \mathbf{N}_m^{n-1} \mathbf{N}_s h)(\mathbf{x}, \omega). \quad (30)$$

With no null scattering at surfaces, only the medium contribution $\mathbf{R}_m \mathbf{N}^n$ in Eq. (29) can be non-zero. For the same reason, operator $\mathbf{N}_s$ can have non-zero contribution only when it appears at the end of a null-scattering chain. Thus, when $(\mathbf{N}_m + \mathbf{N}_s)^n$ is fully expanded into a sum of 2^n chains, only two non-zero chains remain in Eq. (30), where h is evaluated in a medium and at a surface, respectively. (The start of the chain can itself be in a medium or on a surface; this is handled by $\mathbf{R}$ via the common notation from Eqs. (17) and (19).) Next, we write out the two chain operators in Eq. (30) and bring them in similar forms to prepare them for a change of variables.

Operator $\mathbf{R}_m \mathbf{N}_m^n$. This operator evaluates h at a medium point $\mathbf{x}_{n+1}$ after a series of n medium null-scattering events $\mathbf{x}_1, \dots, \mathbf{x}_n$. We first write out $\mathbf{N}_m^n$ by expanding the n operators $\mathbf{N}_m$ from right to left:

$$(\mathbf{N}_m^n h)(\mathbf{x}_1, \omega) = \int_{\mathbf{x}_1}^z \dots \int_{\mathbf{x}_n}^z \left(\prod_{i=1}^n \mu_n(\mathbf{x}_i) \bar{T}(\mathbf{x}_i, \mathbf{x}_{i+1}) \right) h(\mathbf{x}_{n+1}, \omega) d\mathbf{x}_{n+1} \dots d\mathbf{x}_2. \quad (31)$$

$c_n(\mathbf{x}_1, \dots, \mathbf{x}_{n+1}, \omega)$

We then write out the full operator $\mathbf{R}_m \mathbf{N}_m^n$:

$$\begin{aligned} & (\mathbf{R}_m \mathbf{N}_m^n h)(\mathbf{x}_0, \omega) \\ &= \int_{S^2} \int_0^z \frac{c_r(\mathbf{x}_0 \mathbf{x}_1, \omega')}{\rho(\omega, \mathbf{x}_0, \omega') \bar{T}(\mathbf{x}_0, \mathbf{x}_1) D(\mathbf{x}_0, \omega')} (\mathbf{N}_m^n h)(\mathbf{x}_1, \omega') d\mathbf{x}_1 d\omega' \quad (32) \\ &= \int_{S^2} \int_0^z \int_{x_1}^z \cdots \int_{x_n}^z \frac{c_r(\mathbf{x}_0 \mathbf{x}_1, \omega') c_n(\mathbf{x}_1 \dots \mathbf{x}_{n+1}, \omega')}{c(\mathbf{x}_0 \dots \mathbf{x}_{n+1}, \omega')} d\mathbf{x}_{n+1} \cdots d\mathbf{x}_2 d\mathbf{x}_1 d\omega'. \end{aligned}$$

Note that the nested line integration effectively runs over the simplex $0 \leq x_1 \leq \dots \leq x_{n+1} \leq z$. We can switch the integration order of the simplex coordinates x_i without changing the result:

$$= \int_{S^2} \int_0^z \int_0^{x_{n+1}} \int_{x_1}^{x_{n+1}} \cdots \int_{x_{n-1}}^{x_{n+1}} c(\mathbf{x}_0 \dots \mathbf{x}_{n+1}, \omega') d\mathbf{x}_n \cdots d\mathbf{x}_2 d\mathbf{x}_1 d\mathbf{x}_{n+1} d\omega'. \quad (33)$$

Note that now the outermost line integral determines the location of the real-scattering vertex $\mathbf{x}_{n+1}$, which does not any more depend on the locations of all null-scattering vertices.

Operator $\mathbf{R}_m \mathbf{N}_m^{n-1} \mathbf{N}_s$. This operator evaluates h at a surface point $\mathbf{x}_{n+1}$ after a series of n medium null-scattering events $\mathbf{x}_1, \dots, \mathbf{x}_n$. Its expansion is almost identical to that of $\mathbf{R}_m \mathbf{N}_m^n$ in Eq. (33) above:

$$\begin{aligned} & (\mathbf{R}_m \mathbf{N}_m^{n-1} \mathbf{N}_s h)(\mathbf{x}_0, \omega) \\ &= \int_{S^2} \int_0^{x_{n+1}} \int_{x_1}^{x_{n+1}} \cdots \int_{x_{n-1}}^{x_{n+1}} c(\mathbf{x}_0 \dots \mathbf{x}_{n+1}, \omega') d\mathbf{x}_n \cdots d\mathbf{x}_2 d\mathbf{x}_1 d\omega', \quad (34) \end{aligned}$$

but without the line integral over the position of point $\mathbf{x}_{n+1}$, which in this case is fixed at the distance z to the nearest surface.

4.3 Change of variables

We are finally ready to perform the change of variables. We can do this simultaneously for Eqs. (33) and (34), thanks to their almost identical forms. For Eq. (33) we merge the two outermost integrals into volume integration ($\mathcal{V}$), and for Eq. (34) we switch from sphere to area integration ($\mathcal{A}$). We can then merge the two resulting high-dimensional integrals into one, with $\mathcal{V} \cup \mathcal{A}$ as the domain of the outermost integral:

$$\begin{aligned} & (\mathbf{R}_m^n h)(\mathbf{x}_0, \omega) = (\mathbf{R}_m \mathbf{N}_m^n h)(\mathbf{x}, \omega) + (\mathbf{R}_m \mathbf{N}_m^{n-1} \mathbf{N}_s h)(\mathbf{x}, \omega) \quad (35) \\ &= \int_{\mathcal{V} \cup \mathcal{A}} \int_0^{x_{n+1}} \cdots \int_{x_{n-1}}^{x_{n+1}} c(\mathbf{x}_0 \dots \mathbf{x}_{n+1}, \omega_{\mathbf{x}\mathbf{x}_{n+1}}) G(\mathbf{x}_0, \mathbf{x}_{n+1}) d\mathbf{x}_n \cdots d\mathbf{x}_1 d\mathbf{x}_{n+1}, \quad (36) \end{aligned}$$

where as a result the geometry Jacobian term $G(\mathbf{x}_0, \mathbf{x}_{n+1})$ appears. Note that G is evaluated between the two real-scattering vertices, bypassing the n null-scattering vertices $\mathbf{x}_1, \dots, \mathbf{x}_n$.

Recall that we derived Eq. (36) for the null-scattering case of $n > 0$. In the case of $n = 0$, we have $\mathbf{R} \mathbf{N}^n = \mathbf{R} = \mathbf{R}_m + \mathbf{R}_s$, where the change of variables is as simple as switching the integration in $\mathbf{R}_m$ (21) and $\mathbf{R}_s$ (22) to volume ($\mathcal{V}$) and surface area ($\mathcal{A}$), respectively, and merging these two into one integral over the union $\mathcal{V} \cup \mathcal{A}$. The result is an expression identical to Eq. (36) for $n = 0$, where the line integrals vanish as in this case no null scattering occurs between the real-scattering vertices $\mathbf{x}_0$ and $\mathbf{x}_{n+1} \equiv \mathbf{x}_1$.

The final change is to switch from line integration for the null-scattering locations to integration over the null-scattering volume $\mathcal{V}_\delta$. To that end, we take advantage that our generalized scattering

term ρ from Eq. (19) considers null scattering. This allows us to write Eq. (36) as an integral over a product real-null scattering space:

$$\begin{aligned} & (\mathbf{R} \mathbf{N}^n h)(\mathbf{x}_0, \omega) = \int_{(\mathcal{V} \cup \mathcal{A}) \times \mathcal{V}_\delta^n} f_r(\mathbf{x}_0, \mathbf{x}_{n+1}) h(\mathbf{x}_{n+1}, \omega_{\mathbf{x}_{n+1} \mathbf{x}}) d\mathbf{x}_0 \cdots d\mathbf{x}_{n+1} \quad (37) \\ & f_r(\mathbf{x}_{i_0}, \mathbf{x}_{i_1}) = G(\mathbf{x}_{i_0}, \mathbf{x}_{i_1}) \prod_{i=i_0}^{i_1-1} \rho(\omega_{\mathbf{x}_{i-1} \mathbf{x}_i}, \mathbf{x}_i, \omega_{\mathbf{x}_i \mathbf{x}_{i+1}}) \bar{T}(\mathbf{x}_i, \mathbf{x}_{i+1}), \quad (38) \end{aligned}$$

where we use the convenience notation $\omega_{\mathbf{x}_0 \mathbf{x}_{-1}} \equiv \omega$. The function f_r gives the throughput between two real-scattering events. Note that no Jacobian term appears after this change as the differential null-scattering measure $d\mathcal{V}_\delta$ is still a line measure, only that it is not any more aligned with a cardinal axis.

4.4 Path integral formulation

Using the product-integral expression for the null-scattering chain operator in Eq. (37), we can express the pixel measurement (27) as a product-integral as well, by expanding $\mathbf{R} \mathbf{N}^n$ repeatedly from right to left in the chain decomposition (28) of every path operator $\mathbf{P}$:

$$I = \sum_{\mathbf{P} \in \Omega_{\mathbf{P}}} \int_{\mathcal{A}} \left(\frac{\mathbf{R} \mathbf{N}^{n_1} \cdots \mathbf{R} \mathbf{N}^{n_r} L_e}{\mathbf{P}} \right) (\mathbf{x}, \cdot) d\mathbf{x} = \sum_{\mathbf{P} \in \Omega_{\mathbf{P}}} \int_{\mathcal{P}_{\mathbf{P}}} f(\bar{\mathbf{x}}) d\bar{\mathbf{x}}, \quad (39)$$

where we have used the following relation between f (15) and f_r :

$$f(\bar{\mathbf{x}}) = L_e(\mathbf{x}_k, \mathbf{x}_k) \prod_{i=0}^{r-1} f_r(\mathbf{x}_{i_i}, \mathbf{x}_{i_{i+1}}). \quad (40)$$

For every $\mathbf{P} \in \Omega_{\mathbf{P}}$ there is a corresponding path space $\mathcal{P}_{\mathbf{P}}$ that contains all paths with certain length and sequence of scattering types. The union $\mathcal{P}$ (13) of all these spaces then contains every possible path of any length and type, and the corresponding path integral is written in Eq. (12). This concludes our derivation.

5 APPLICATIONS

Our light transport formulation from Section 3 provides a framework for devising new and expressing existing unbiased volumetric rendering methods as direct Monte Carlo estimators of a path integral. In our framework the pdfs of the corresponding sampling techniques are known, which enables the application of multiple importance sampling (MIS) [Veatch and Guibas 1995] in a straightforward manner, something that was previously considered difficult.

To showcase the capabilities of our framework, we describe a practical path tracing algorithm that combines unidirectional and next-event sampling techniques. We also incorporate hero wavelength MIS to handle spectrally varying media [Wilkie et al. 2014].

5.1 Monte Carlo estimators for the path integral

A direct Monte Carlo estimator for our path integral (12) has the same well-known general form introduced by Veatch [1997]:

$$\langle I \rangle(\bar{\mathbf{x}}) = \frac{f(\bar{\mathbf{x}})}{p(\bar{\mathbf{x}})} = \frac{f(\bar{\mathbf{x}})}{p(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k)}. \quad (41)$$

Here, $\bar{\mathbf{x}}$ is a randomly sampled path of length k , where its pdf $p(\bar{\mathbf{x}})$, i.e. the joint density of its vertices, is expressed w.r.t. the path measure $d\bar{\mathbf{x}}$ from Eq. (13).

With an appropriately chosen null density μ_n , the path contribution $f(\bar{\mathbf{x}})$ (and specifically the combined transmittance term $\bar{T}$) is

analytically evaluable and sampleable for any path in both homogeneous and heterogeneous media. This allows combining different path sampling techniques via MIS into one low-variance estimator:

$$\langle I \rangle_{\text{MIS}} = \sum_{i=1}^n \frac{1}{n_i} \sum_{j=1}^{n_i} w_i(\bar{\mathbf{x}}_{i,j}) \frac{f(\bar{\mathbf{x}}_{i,j})}{p_i(\bar{\mathbf{x}}_{i,j})}, \quad (42)$$

where n is the number of techniques and $\bar{\mathbf{x}}_{i,j}$ are n_i independently sampled paths from distribution p_i . The *power heuristic* $w_i(x) = [n_i p_i(x)]^\beta / \sum_{k=1}^n [n_k p_k(x)]^\beta$ is a provably good way to combine techniques in terms of minimizing the variance of the estimator $\langle I \rangle_{\text{MIS}}$. We use the *balance heuristic*, which corresponds to $\beta = 1$.

5.2 Path sampling techniques

When expressed in a path integral framework, differences between unbiased pixel estimators (41) can only be due to differences in their respective path sampling pdfs. These correspond to different factorizations of the joint distribution into a series of conditional vertex pdfs [Georgiev et al. 2013], e.g. unidirectional or bidirectional, as well as different choices for these vertex pdfs.

Due to their forward-scattering nature, null vertices must be sampled along segments connecting real-scattering vertices. This is typically done either along a ray sampled at one real-scattering vertex, or along a line segment resulting from a bidirectional (e.g. next-event) connection between two real-scattering vertices. The former is the continuous, volumetric analog of tracing through a specular scattering chain along a given direction in surface rendering. The latter is a rough continuous analog to connecting real-scattering vertices through a refractive boundary [Walter et al. 2009; Hanika et al. 2015]. Next, we discuss several existing sampling techniques, illustrated in Fig. 5, and express their path pdfs in our framework.

Delta tracking. This method is a direct (unidirectional) estimator for the null-scattering VRE (9) [Novák et al. 2018]. It samples the distance to the next event along the path with pdf proportional to the combined transmittance $\bar{T}$. It then selects the type of event—absorption, real or null scattering—with respective probabilities equal to the fractional albedos in Eq. (8). It samples the phase function ρ_m in the case of real scattering, or continues forward in the case of null scattering, and terminates upon an absorption event or when a light source is hit. In our framework, the path pdf of this technique, illustrated in Fig. 5, left, reads

$$p_{\text{dt}}(\bar{\mathbf{x}}) = p(\mathbf{x}_0) \prod_{i=1}^k p(\omega_i | \mathbf{x}_{i-1}, \omega_{i-1}) p(x_i | \mathbf{x}_{i-1}, \omega_i) P(e_{\mathbf{x}_i}) \left[\prod_{i=0}^r G(\mathbf{x}_{r_i}, \mathbf{x}_{r_{i+1}}) \right],$$

where $e_{\mathbf{x}} \in \{a, s, n\}$ denotes the event type at vertex $\mathbf{x}$ and where

$$p(\omega' | \mathbf{x}, \omega) \propto \rho(\omega, \mathbf{x}, \omega'), \quad p(y | \mathbf{x}, \omega) = \bar{T}(\mathbf{x}, y) \bar{\mu}_t(y), \quad P(e_{\mathbf{x}}) = \frac{\mu_{e_{\mathbf{x}}}(\mathbf{x})}{\bar{\mu}_t(\mathbf{x})}.$$

Note the occurrence of the geometry terms between every two consecutive real-scattering vertices when the path pdf is expressed w.r.t. the measure $d\bar{\mathbf{x}}$ (13). All terms in the corresponding path contribution thus cancel out in the pixel estimator (41), except for the emitted radiance $L_e(\mathbf{x}_k, \omega_{\mathbf{x}_k} \mathbf{x}_{k-1})$.




Fig. 5. Mapping of existing unidirectional, next-event, and joint path sampling methods to techniques in our integral formulation. Delta tracking (left) starts from the camera and samples the path vertices in order until termination, choosing the scattering type—absorption, real, or null scattering—at every vertex. Given a light vertex ($\mathbf{x}_5$), ratio tracking (middle) samples null-scattering vertices ($\mathbf{x}_3$ and $\mathbf{x}_4$) along the connection segment with a real-scattering vertex ($\mathbf{x}_2$) on the camera subpath. Given a light vertex ($\mathbf{x}_5$) and a camera ray ($\mathbf{x}_0, \omega$), equiangular sampling (right) determines the location of a real-scattering vertex ($\mathbf{x}_2$); the null-scattering vertices along the two resulting connection segments can be sampled via ratio tracking.

Ratio tracking. With existing path integral formulations that consider only real scattering, this technique is used as an unbiased estimator for the extinction transmittance T between two path vertices [Novák et al. 2014]. It performs a random walk along the connection segment, using the same distance sampling as delta tracking but deterministically selecting null scattering at every vertex. In our framework, this estimator can be interpreted as sampling a null-scattering subpath between two given real-scattering vertices $\mathbf{x}_i$ and $\mathbf{x}_j$. The joint pdf of these $j - i - 1$ null vertices is

$$p_{\text{rt}}(\mathbf{x}_i, \mathbf{x}_j) = \left[\prod_{l=i+1}^{j-1} \frac{\bar{T}(\mathbf{x}_{l-1}, \mathbf{x}_l) \bar{\mu}_t(\mathbf{x}_l)}{p(\mathbf{x}_l | \mathbf{x}_{l-1}, \omega_{\mathbf{x}_l \mathbf{x}_j})} \right] \bar{T}(\mathbf{x}_{j-1}, \mathbf{x}_j), \quad (43)$$

where the last transmittance term is the random walk termination probability—the distance pdf integrated from $\|\mathbf{x}_j - \mathbf{x}_{j-1}\|$ to infinity. Note that the pdf is invariant to the direction in which the random walk is performed: from $\mathbf{x}_i$ toward $\mathbf{x}_j$ or from $\mathbf{x}_i$ toward $\mathbf{x}_j$.

Ratio tracking is typically used to connect the real-scattering end vertices of delta-tracking sampled light and camera subpaths. Denoting the number of vertices on these subpaths respectively as s and t , the path pdf of this (real) *vertex connection* technique is

$$p_{s,t}(\bar{\mathbf{x}}) = p_{\text{dt}}(\mathbf{x}_{r_0} \dots \mathbf{x}_{r_{t-1}}) p_{\text{rt}}(\mathbf{x}_{r_{t-1}}, \mathbf{x}_{r_{r-s+1}}) p_{\text{dt}}(\mathbf{x}_{r_r} \dots \mathbf{x}_{r_{r-s+1}}). \quad (44)$$

The special case of $s = 1$ is known as *next-event estimation*; we illustrate this technique in Fig. 5, middle.

Equiangular sampling. In the next-event estimator, the geometry term along the connection segment does not cancel out as it does not appear in the path pdf (44). This can cause high variance when the light vertex is inside the medium. To remedy this, the equiangular sampling technique generates the position of the penultimate real-scattering vertex $\mathbf{x}_{r_{r-1}}$ along the ray ($\mathbf{x}_{r_{r-2}}, \omega_{r_{r-1}}$) with pdf proportional to that geometry term [Kulla and Fajardo 2012]:

$$p_{\text{ea}}(\mathbf{x}_{r_{r-1}} | \mathbf{x}_{r_{r-2}}, \omega_{r_{r-1}}, \mathbf{x}_k) \propto G(\mathbf{x}_{r_{r-1}}, \mathbf{x}_k) = \frac{1}{\|\mathbf{x}_{r_{r-1}} - \mathbf{x}_k\|^2}. \quad (45)$$

This results in two connection segments between the real-scattering vertices $\mathbf{x}_{r-2}$, $\mathbf{x}_{r-1}$, and $\mathbf{x}_k \equiv \mathbf{r}_r$. With prior path integral formulations, ratio tracking can be used to estimate the extinction transmittance T along these segments. In our formulation this translates to sampling null-scattering subpaths between those vertices. Combined with a delta-tracking sampled prefix subpath from the camera, the full path pdf of this technique, illustrated in Fig. 5, right, reads

$$p(\bar{\mathbf{x}}) = p_{\text{dt}}(\mathbf{x}_0 \dots \mathbf{x}_{r-2}) p(\omega_{r-1} | \mathbf{x}_{r-2}, \omega_{r-2}) p(\mathbf{x}_k) p_{\text{ea}}(\mathbf{x}_{r-1} | \mathbf{x}_{r-2}, \omega_{r-1}, \mathbf{x}_k) \cdot p_{\text{rt}}(\mathbf{x}_{r-2}, \mathbf{x}_{r-1}) p_{\text{rt}}(\mathbf{x}_{r-1}, \mathbf{x}_k) G(\mathbf{x}_{r-2}, \mathbf{x}_{r-1}). \quad (46)$$

5.3 Handling chromatic media

In media with spectrally varying coefficients μ^c , a different path integral (12) has to be estimated for every color component c :

$$\langle I_c \rangle(\bar{\mathbf{x}}) = \frac{f_c(\bar{\mathbf{x}})}{p_{\tilde{c}}(\bar{\mathbf{x}})}, \quad (47)$$

where the path pdf $p_{\tilde{c}}(\bar{\mathbf{x}})$ now depends on the chosen *sampling component* $\tilde{c}$ which drives the distance and event type sampling in the path construction via the corresponding coefficients $\mu_a^{\tilde{c}}, \mu_s^{\tilde{c}}, \mu_n^{\tilde{c}}$.

One approach is to compute a separate estimator for every component c , setting $\tilde{c} = c$ in each estimator. Unfortunately, estimating every component separately can be costly and also leads to color noise. To remedy this issue, Wilkie et al. [2014] proposed to combine all these estimators via MIS. This requires evaluating the pdf p_c for every component c , which, as the authors point out, has been unavailable for unbiased path sampling techniques in heterogeneous media. By providing analytic pdfs for such techniques, our formulation makes this MIS combination straightforward.

Another approach is to estimate all integrals I_c simultaneously using the sampling technique $p_{\tilde{c}}$ corresponding to a chosen component $\tilde{c}$. Unfortunately, this can lead to extreme variance for components c with medium density (locally) larger than that of $\tilde{c}$. Kutz et al. [2017] proposed sampling distances according to a combined extinction $\bar{\mu}_t$ that bounds the extinctions of all color components along with carefully chosen event type probabilities to limit the variance of the estimate $\langle I_c \rangle$ of each component. This conservative technique eliminates the color noise, but the dense sampling can significantly increase the computational cost when one channel has medium density significantly higher than the rest. Our framework allows us to track all color components together using distances sampled from extinctions that bound only an individual channel, while mitigating the potential for higher variance using MIS.

5.4 Unidirectional path tracing

Our path integral formulation enables the application of MIS to combine the estimators of different sampling techniques. In our implementation, sampling a path begins with selecting a random color component. We use this component to guide the distance and albedo-based scattering event type sampling for the entirety of the path. While we may employ additional strategies such as next event estimation or equiangular sampling, the sampling color component remains invariant. This approach gives us a spectral version of each path sampling technique whose pdf we must track in addition to the path contribution. Every time a full path is completed, we use these

Algorithm 1. Our spectral MIS using the balance heuristic, combining N sampling techniques, one for each color component $c_1, \dots, c_N$. The function $\text{SampleEvent}(P_1, P_2, P_3)$ chooses between three events with probabilities $P_1 + P_2 + P_3 = 1$, and the symbol $\circ$ denotes component-wise multiplication.

```

1: function UNIDIRECTIONALSPECTRALMIS( $\mathbf{x}, \omega$ )
2:    $\mathbf{f} \leftarrow \langle 1, \dots, 1 \rangle_N$  ▷ Initialize per-component path contributions
3:    $\mathbf{p} \leftarrow \langle 1, \dots, 1 \rangle_N$  ▷ Initialize per-component path pdfs
4:    $\tilde{c} \leftarrow \text{SelectRandom}(c_1, \dots, c_N)$  ▷ Select color component
5:   while true do ▷ Begin path random walk
6:      $t \leftarrow -\frac{\ln(1-\xi)}{\bar{\mu}_t^{\tilde{c}}}$  ▷ Sample medium interaction;  $\xi \in [0, 1]$ 
7:      $\mathbf{x} \leftarrow \mathbf{x} + t\omega$ 
8:      $e \leftarrow \text{SampleEvent}\left(\frac{\mu_a^{\tilde{c}}(\mathbf{x})}{\bar{\mu}_t^{\tilde{c}}}, \frac{\mu_s^{\tilde{c}}(\mathbf{x})}{\bar{\mu}_t^{\tilde{c}}}, \frac{\mu_n^{\tilde{c}}(\mathbf{x})}{\bar{\mu}_t^{\tilde{c}}}\right)$  ▷  $e \in \{a, s, n\}$ 
9:      $\mathbf{f} \leftarrow \mathbf{f} \circ \langle \bar{T}^{c_1}(t)\mu_e^{c_1}(\mathbf{x}), \dots, \bar{T}^{c_N}(t)\mu_e^{c_N}(\mathbf{x}) \rangle$  ▷ Update contrib, pdf
10:     $\mathbf{p} \leftarrow \mathbf{p} \circ \langle \bar{\mu}_t^{c_1} \bar{T}^{c_1}(t) \frac{\mu_e^{c_1}(\mathbf{x})}{\bar{\mu}_t^{c_1}}, \dots, \bar{\mu}_t^{c_N} \bar{T}^{c_N}(t) \frac{\mu_e^{c_N}(\mathbf{x})}{\bar{\mu}_t^{c_N}} \rangle$ 
11:    if  $e = a$  then
12:      return  $\frac{\mathbf{f} \circ L_e(\mathbf{x}, -\omega)}{\text{average}(\mathbf{p})}$  ▷ Balance heuristic  $\rightarrow$  average  $N$  techniques
13:    else if  $e = s$  then
14:       $\omega \leftarrow \text{sample} \propto \rho_m(\omega)$  ▷ Sample phase function

```

Algorithm 2. The shadow connection routine for our Spectral+NEE MIS method. The symbol $\circ$ denotes component-wise multiplication.

```

1: function SAMPLENEXTEVENT( $\mathbf{x}, \omega, \tilde{c}, \mathbf{f}, \mathbf{p}$ ) ▷ See Alg. 1 for  $\tilde{c}, \mathbf{f}, \mathbf{p}$ 
2:    $\omega', p_{\omega'} \leftarrow \text{SampleLightDir}(\mathbf{x})$  ▷ Sample direction toward light
3:    $\mathbf{f} \leftarrow \mathbf{f} \cdot \rho_m(\omega, \mathbf{x}, \omega')$  ▷ Update path contributions
4:    $\mathbf{p}_{\text{nee}} \leftarrow \mathbf{p} \cdot p_{\omega'}$  ▷ Initialize next-event path pdfs
5:    $\mathbf{p}_{\text{uni}} \leftarrow \mathbf{p} \cdot \rho_m(\omega, \mathbf{x}, \omega')$  ▷ Initialize unidirectional path pdfs
6:   while true do ▷ Begin ratio tracking along shadow ray
7:      $t \leftarrow -\frac{\ln(1-\xi)}{\bar{\mu}_t^{\tilde{c}}}$  ▷ Sample medium interaction;  $\xi \in [0, 1]$ 
8:      $\mathbf{y} \leftarrow \text{IntersectRay}(\mathbf{x}, \omega', t)$ 
9:     if  $\mathbf{y}$  then return  $\frac{\mathbf{f} \circ L_e(\mathbf{y}, -\omega')}{\text{average}(\mathbf{p}_{\text{uni}}, \mathbf{p}_{\text{nee}})}$  ▷ Surface hit, avg.  $2N$  techniques
10:     $\mathbf{x} \leftarrow \mathbf{x} + t\omega'$ 
11:     $\mathbf{f} \leftarrow \mathbf{f} \circ \langle \bar{T}^{c_1}(t)\mu_n^{c_1}(\mathbf{x}), \dots, \bar{T}^{c_N}(t)\mu_n^{c_N}(\mathbf{x}) \rangle$  ▷ Only evaluate null
12:     $\mathbf{p}_{\text{nee}} \leftarrow \mathbf{p}_{\text{nee}} \circ \langle \bar{\mu}_t^{c_1} \bar{T}^{c_1}(t), \dots, \bar{\mu}_t^{c_N} \bar{T}^{c_N}(t) \rangle$ 
13:     $\mathbf{p}_{\text{uni}} \leftarrow \mathbf{p}_{\text{uni}} \circ \langle \bar{\mu}_t^{c_1} \bar{T}^{c_1}(t) \frac{\mu_n^{c_1}(\mathbf{x})}{\bar{\mu}_t^{c_1}}, \dots, \bar{\mu}_t^{c_N} \bar{T}^{c_N}(t) \frac{\mu_n^{c_N}(\mathbf{x})}{\bar{\mu}_t^{c_N}} \rangle$ 

```

pdfs to calculate an MIS weight. Algorithm 1 provides pseudo-code for a unidirectional variant of this method.

In versions of our algorithm where we combine two path construction techniques, such as unidirectional sampling and next-event estimation, the number of path pdfs we need to track is doubled. Algorithm 2 lists our next-event routine.

5.5 Bidirectional path tracing

In bidirectional path tracing (BPT) based on delta-tracked camera and light subpaths and ratio-tracked vertex connections, the MIS weights have previously been restricted to using either only directional real-scattering pdfs or costly ray-marched transmittance sampling pdf estimates. By considering null scattering explicitly, our framework provides full path pdfs that are exact and cheap

to evaluate. We can use these pdfs to improve the accuracy of the technique weights and thus facilitate the variance reduction of MIS.

For a path of total length k in our framework, the number of possible BPT sampling techniques is determined by its real-scattering length r , as in prior formulations. Techniques differ by the number of real-scattering vertices on the light and camera subpaths they connect, respectively s and t (with $s + t - 1 = r$), with pdf given by Eq. (44). Applying our spectral MIS on top increases the number of techniques by a factor equal to the number of color components. Our current BPT implementation does not employ spectral MIS and handles achromatic media only.

6 RESULTS

To demonstrate the utility of our path integral framework, we compare variants of our methods to prior work on a variety of scenes that benefit from the use of MIS. We evaluate rendering efficiency by measuring images' root mean square error (RMSE) and lookups to unit variance (LTUV). The LTUV is a renderer-agnostic metric computed by multiplying the mean squared error (i.e. squared RMSE) of the image by its total number of medium lookups [Kutz et al. 2017]. Tables 1 and 2 provide descriptions of the methods and the media used in our experiments.

Next-event estimation. Figure 6 shows two bunny-shaped clouds, each with a highly forward scattering Henyey-Greenstein phase function. The thin medium on the top is illuminated from behind by a large spherical light source, and the dense medium on the bottom is lit from the side by a small spherical source. Unidirectional sampling and next-event estimation (NEE) on their own each fail on one scene. Combining them via MIS within our formulation, however, avoids these failures and handles both scenes robustly. We also compare to the directional MIS approach of Kutz et al. [2017]; it uses delta tracking along shadow rays, thus the combination weights with unidirectional sampling (also based on delta tracking) only consider the pdf of the shadow ray direction. In the dense medium, our method performs better as it can leverage the more efficient ratio tracking technique in NEE. This technique is also more expensive than delta tracking, causing a slight performance loss in our method compared to directional MIS in the thinner medium. However, this loss is relatively small compared to our improved efficiency in the dense medium.

Equiangular sampling. Our formulation provides the flexibility of choosing not only between delta tracking and ratio tracking but also other path sampling techniques. This flexibility, a result of knowing how to analytically compute path densities, is one of the key benefits of the null path formulation. Figure 7 demonstrates this by including equiangular sampling [Kulla and Fajardo 2012]. Next-event estimation (NEE) performs well in the right side of the cloud, illuminated by a distant light source, but exhibits extreme noise near the small spherical lights inside the cloud, due to not importance sampling the geometry term along the shadow ray. Our MIS combination of NEE and equiangular sampling handles both of these lighting configurations efficiently. Previously, equiangular sampling could only be combined with regular tracking-based or biased ray-marching-based transmittance sampling techniques.

Table 1. Descriptions of the various methods used in our experiments, where ours are marked in bold. BPT stands for bidirectional path tracing.

Method	Description
Unidirectional	Analog tracing: delta tracking & phase function sampling
Next-event (NEE)	As above, with ratio-tracked light source connections
Directional MIS	Unidirectional + delta-tracked light source connections
Unidir.+NEE MIS	Unidirectional + ratio-tracked light source connections
Independent tracking	Directional MIS applied separately to R,G,B components
Spectral tracking	Method of Kutz et al. [2017]; uses Directional MIS
Spectral MIS	MIS between per-component Directional MIS
Spectral+NEE MIS	MIS between per-component Unidirectional+NEE MIS
BPT: Directional MIS	Technique weights based on real-scattering pdfs only
BPT: Full MIS	Technique weights based on our full path pdfs

Table 2. The maximum extinction ($\bar{\mu}_t$), scattering albedo (μ_s/μ_t), and scattering anisotropy (g) of the various media shown in the figures throughout the paper. Spatially varying albedos are given as min:max ranges. The color bars in Figs. 1, 9 and 10 illustrate the maximum extinction ($\bar{\mu}_t$, dark) and maximum scattering ($\bar{\mu}_s$, light) coefficients for every color component.

Medium	$\bar{\mu}_t$	μ_s/μ_t	g
Fig. 1	30, 100, 30	0.001:0.95, 0.001:0.95, 0.001:0.95	-0.4
Fig. 6, top	4	1	0.99
Fig. 6, bottom	16	1	0.99
Fig. 7	20	1	0
Fig. 9, left	30, 100, 30	0.001:0.95, 0.001, 0.001:0.95	-0.4
Fig. 9, right	30, 35, 30	0.001:0.95, 0.001:0.95, 0.001:0.95	-0.4
Fig. 10	600, 50, 60	0.001:0.1, 0.0075:0.7, 0.001:0.35	0
Fig. 11	20	0.5	0

Spectral media. Our MIS framework also provides benefits when rendering spectrally varying media. Figure 8 visualizes the path throughput variation in a perfectly forward-scattering medium. The traditional independent tracking method renders the color components independently, choosing a random one to estimate for each pixel sample. Our method uses the same sampling techniques but combines them via MIS. As a result, we both reduce color noise and avoid the increased cost of spectral tracking's use of a global bounding majorant [Kutz et al. 2017]. Our method is still prone to some color noise due to under-sampling the majorant of a channel, though this noise is bounded by the number of techniques in MIS.

Figures 1 and 9 demonstrate how our framework allows us to overcome the limitations of both independent tracking and spectral tracking. In the Fig. 1 plume, the green channel is significantly more dense than the red and blue. Independent tracking shows noticeable color noise, even though it is able to estimate the pixel luminance better than spectral tracking whose sampling cost is increased by the green-channel density. Our spectral MIS method has little color noise and estimates the luminance better, which is improved further by the use of next-event estimation (NEE).

In the left plume in Fig. 9 we reduce the albedo of the green channel. Doing this reduces the efficiency of spectral tracking even further as this channel determines the global bounding majorant but its scattering contribution is negligible. Both our methods and independent tracking avoid always using this high majorant. In the right plume, the densities of the channels are similar in magnitude and their albedo ranges are identical, making the cost reduction from tight, spectrally varying majorants less significant. Our spectral MIS




Fig. 6. Unidirectional sampling and next-event estimation on their own cannot robustly handle both the back-lit thin medium (top) and side-lit dense medium (bottom). Our MIS combination of the two techniques handles both scenes efficiently, and its ratio-tracked shadow connections provide improvement in the dense bunny over directional MIS which is limited to using delta tracking along these connections.

method still performs on par with spectral tracking, and the addition of ratio-tracked NEE brings more significant improvement.

Our approach also extends to solid media like translucent stone. In Fig. 10 we model malachite by interpolating between two sets of medium parameters based on a lookup into a malachite color spline that is offset by several octaves of procedural noise [Perlin 1985]. Since the medium is enclosed by a dielectric boundary which prevents shadow connections, we use pure unidirectional path construction for all three compared methods. Thus, the differences between them are in the spectral handling only. Even though lookups in this medium are significantly cheaper than accessing the hierarchical density grid structure of the plumes, our method still provides improvements when medium parameters force spectral tracking to use an overly conservative majorant. Note that, in principle, the per-channel sampling majorants in our method can be set arbitrarily, hence spectral tracking is the special case where all color components use the same (bounding) majorant.

Bidirectional path tracing. In the scene in Fig. 11, the medium covers the light source, causing unidirectional path tracing to fail due to geometric singularities near the light source that it does not importance sample [Georgiev et al. 2013]. Bidirectional path tracing (BPT) handles these singularities effectively via MIS. We compare our weights based on full path pdfs against a traditional implementation restricted to using directional real-scattering pdfs only. By accounting for medium sampling, our weights serve as better proxies for the efficiency of the sampling techniques, which in turn improves the quality of their combination, as seen in the zoom-ins. To gain




Fig. 7. Our framework provides a principled way to combine equiangular sampling and next-event estimation in heterogeneous media via MIS. Scene modeled after Kulla and Fajardo [2012, Fig. 7].

more insight, we compute the contribution of each (s, t) technique weighted by our full-path pdf weights and by directional pdf weights. We then show the positive/negative difference between the two resulting images (in blue/purple), relative to the total contribution of the corresponding real-scattering path length $r = s + t - 1$. The set of (s, t) false-color images highlights the differences between the two weighting schemes. Being oblivious to attenuation through the medium, directional MIS prefers techniques that perform connections along long path segments, e.g. camera connections ($t = 1$). In contrast, our weights tend to prefer techniques that perform connections through the dense medium, e.g. next-event ($s = 1$), recognizing that ratio tracking can be much more efficient along such segments than a delta-tracked random walk.

7 CONCLUSIONS, DISCUSSION & FUTURE WORK

We have presented a null-scattering extension of the standard path integral. Critically, instead of treating null scattering as a black-box rejection sampling process, our extension accounts for null scattering directly in the path integral itself. This makes it possible, for the first time, to compute analytic path pdfs for null-scattering methods, enabling their combination with other complementary techniques, and with each other using MIS. Many production renderers have so far avoided null-scattering approaches, favoring more expensive regular tracking [Fascione et al. 2018] or biased ray marching [Georgiev et al. 2018; Fascione et al. 2018], since these approaches can leverage the increased robustness of MIS. With our framework, null-scattering approaches become viable in this setting. Furthermore, our theory lays out a solid foundation for interpreting current, and developing new unbiased rendering techniques of spatially and spectrally varying participating media.




Fig. 8. We demonstrate the MIS abilities of our framework in a perfectly forward-scattering, non-absorptive 1D medium with spatially and spectrally varying density. We plot the average path throughput of each color component at real-scattering events, along with a ± 1 standard deviation band. The top-left insets are noise exemplars produced by each method. Our method uses the same sampling techniques as independent tracking, but combines them via MIS. This results in a reduction of color noise comparable to that of spectral tracking but at a lower cost (i.e. number of medium lookups), thanks to avoiding the use of a common conservative majorant for all color components. Experiment designed after Kutz et al. [2017, Fig. 9].




Fig. 9. Spectral tracking uses a common majorant for all color components, which leads to high per-pixel-sample cost in the left plume whose green channel has high density but low albedo. Our MIS combination of per-channel sampling techniques outperforms spectral tracking by reducing the number of expensive medium lookups while avoiding the color noise of independent tracking. In the right plume, whose spectral variation is small, we can still outperform spectral tracking through the use of ratio-tracked next-event estimation (NEE). Table 1 describes the individual methods and Table 2 provides the medium parameters.




Fig. 10. In this procedurally generated, malachite-like material, medium lookups are relatively cheap, yet by avoiding the use of a common bounding majorant for all color components our spectral MIS method achieves noticeable variance reduction over spectral tracking. Independent tracking outperforms spectral tracking for the same reason, albeit at the expense of more color noise. Table 1 describes the individual methods and Table 2 provides the medium parameters.




Fig. 11. Our formulation provides full path pdfs that incorporate medium sampling, making the MIS weights in bidirectional path tracing (BPT) account better for sampling technique efficiency compared to prior implementations restricted to using directional real-scattering pdfs only. Unidirectional path tracing fails in this scene due to path contribution singularities it does not importance sample. On the right, we show the contributions of individual real-scattering path lengths r , along with the difference between the full- and directional-MIS-weighted contribution images for every $(s, t = r - s + 1)$ technique. The camera-connection techniques $(t = 1)$ are dominated by purple as they are preferred more by directional MIS which is oblivious to the presence of media. Our full MIS scheme prioritizes vertex connections through the dense medium (e.g. $s = 1$). Black means the technique is weighted equally by both schemes.

Other light transport algorithms. While we have demonstrated our theory in unidirectional and bidirectional path tracers, our framework can be readily used to leverage or combine null scattering within other path sampling approaches in density estimation [Jensen et al. 2001; Jensen and Christensen 1998; Jarosz et al. 2008, 2011a,b; Krivánek et al. 2014], joint path sampling [Georgiev et al. 2013], Metropolis light transport [Pauly et al. 2000], and gradient-domain rendering [Kettunen et al. 2015].

Risks of MIS. As with any application of MIS, our approach can be less efficient than one of the combined techniques alone. Since this risk grows with the number of techniques, one needs to be careful to not increase this number unnecessarily. Similarly, while our framework allows more flexibility in the sampling techniques used (e.g. using non-bounding majorants in spectrally varying media), this flexibility should be exploited judiciously as it can lead to higher variance (e.g. increased color noise in zero-density regions). Combining multiple techniques via MIS also comes with additional implementation and validation burden.

Path-space manifolds. It may be interesting to extend our theory to include null- and specular-scattering surfaces. This would provide a unified framework that considers null volume scattering, transparent surfaces, and the specular path-space manifold [Jakob 2013]. Our null-scattering path space in fact directly encodes a path-space manifold since all null vertices are restricted to lie on the line between real vertices. This means that in contrast to specular surfaces, which require numerical exploration to remain on the manifold [Jakob 2013], direct Monte Carlo sampling of our path space always remains on the manifold. Ratio tracking [Novák et al. 2014] in our framework is effectively a manifold next-event estimation [Hanika et al. 2015] technique that operates across multiple specular (null) vertices without the need for exploration to stay on the manifold.

Extending other techniques. By providing analytic pdfs, we believe that our framework may also ease the process of extending other rendering techniques to account for heterogeneous media. The recently developed photon planes and volumes [Bitterli and Jarosz 2017] extension to photon mapping currently works only in the

homogeneous setting, in part due to the inability to evaluate path pdfs for arbitrary locations on a photon plane. Path guiding techniques [Vorba et al. 2014; Herholz et al. 2016; Müller et al. 2017] could also conceivably be extended to consider propagation distances in addition to directions, combining the resulting guiding strategies with traditional path sampling techniques via MIS. By explicitly treating null collisions as additional dimensions in path space, our framework might allow for constructing better correlated offset paths in volumetric gradient-domain rendering [Kettunen et al. 2015]. Finally, it may be possible to derive a similar null-scattering path space extension to allow heterogeneous non-exponential or correlated participating media [Jarabo et al. 2018; Bitterli et al. 2018].

ACKNOWLEDGMENTS

All images have been rendered using PBRT [Pharr et al. 2016]. The media in Figs. 1, 6 and 9 are from the OpenVDB [Museth 2013] repository and the cloud model in Fig. 7 is from Disney Research. Figures 1 and 9 are illuminated by an environment map from HDRI Haven. Fig. 10 is from the PBRT scene repository, and contains a dragon model from XYZ RGB Inc. and an environment map from Bernhard Vogl. The scene in Fig. 11 is from Benedikt Bitterli and contains a smoke plume from the PBRT scene repository.

This work was supported by a gift from Autodesk and National Science Foundation Grants IIS-1812796 and CNS-1205521. The authors also thank John Miller at University of California, Berkeley for providing access to computing resources.

REFERENCES

- James Arvo. 1993. Transfer Functions in Global Illumination. In *ACM SIGGRAPH '93 Course Notes - Global Illumination*.
- James Richard Arvo. 1995. *Analytic methods for simulated light transport*. Ph.D. Dissertation. Yale University.
- Hugo W. Bertini. 1963. *Monte Carlo simulations on intranuclear cascades*. Technical Report ORNL-3383. Oak Ridge National Laboratory, Oak Ridge, TN, USA.
- Benedikt Bitterli and Wojciech Jarosz. 2017. Beyond Points and Beams: Higher-Dimensional Photon Samples for Volumetric Light Transport. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 36, 4 (July 2017).
- Benedikt Bitterli, Srinath Ravichandran, Thomas Müller, Magnus Wrenninge, Jan Novák, Steve Marschner, and Wojciech Jarosz. 2018. A radiative transfer framework for non-exponential media. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 37, 6 (Nov. 2018).

- Brent Burley, David Adler, Matt Jen-Yuan Chiang, Hank Driskill, Ralf Habel, Patrick Kelly, Peter Kutz, Yining Karl Li, and Daniel Teece. 2018. The Design and Evolution of Disney's Hyperion Renderer. *ACM Transactions on Graphics* 37, 3, Article 33 (July 2018).
- John C. Butcher and Harry Messel. 1958. Electron Number Distribution in Electron-Photon Showers. *Phys. Rev.* 112 (Dec. 1958). Issue 6.
- John C. Butcher and Harry Messel. 1960. Electron number distribution in electron-photon showers in air and aluminium absorbers. *Nuclear Physics* 20 (1960).
- Subrahmanyan Chandrasekhar. 1960. *Radiative Transfer*. Dover Publications.
- Per Christensen, Julian Fong, Jonathan Shade, Wayne Wooten, Brenden Schubert, Andrew Kensler, Stephen Friedman, Charlie Kilpatrick, Cliff Ramshaw, Marc Bannister, Brenton Rayner, Jonathan Brouillat, and Max Liani. 2018. RenderMan: An Advanced Path-Tracing Architecture for Movie Rendering. *ACM Transactions on Graphics* 37, 3, Article 30 (Aug. 2018).
- Per H. Christensen and Wojciech Jarosz. 2016. The Path to Path-Traced Movies. *Foundations and Trends in Computer Graphics and Vision* 10, 2 (Oct. 2016).
- 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).
- Luca Fascione, Johannes Hanika, Marcos Fajardo, Per Christensen, Brent Burley, Brian Green, Rob Pieké, Christopher Kulla, Christophe Hery, Ryusuke Villemin, Daniel Heckenberg, and André Mazzzone. 2017. SIGGRAPH 2017 Course Notes: Path tracing in production (Parts 1 and 2). In *SIGGRAPH Courses*.
- Luca Fascione, Johannes Hanika, Mark Leone, Marc Droske, Jorge Schwarzhaupt, Tomáš Davidović, Andrea Weidlich, and Johannes Meng. 2018. Manuka: A Batch-Shading Architecture for Spectral Path Tracing in Movie Production. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 37, 3, Article 31 (Aug. 2018).
- Mathieu Galtier, Stéphane Blanco, Cyril Caliot, Christophe Coustet, Jérémie Dauchet, Mouna El Hafi, Vincent Eymet, Richard Fournier, Jacques Gautrais, Anaïs Khuong, Benjamin Piaud, and Guillaume Terrée. 2013. Integral formulation of null-collision Monte Carlo algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer* 125 (2013).
- Manuel Gamito. 2018. Path Tracing in Production: Path Tracing the Framestorian Way. In *ACM SIGGRAPH 2018 Courses*. ACM, New York, NY, USA, Article 15.
- Iliyan Georgiev, Thiago Ize, Mike Farnsworth, Ramón Montoya-Vozmediano, Alan King, Brecht Van Lommel, Angel Jimenez, Oscar Anson, Shinji Ogaki, Eric Johnston, Adrien Herubel, Declan Russell, Frédéric Servant, and Marcos Fajardo. 2018. Arnold: A Brute-Force Production Path Tracer. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 37, 3, Article 32 (Aug. 2018).
- Iliyan Georgiev, Jaroslav Krivánek, Toshiya Hachisuka, Derek Nowrouzezahrai, and Wojciech Jarosz. 2013. Joint Importance Sampling of Low-Order Volumetric Scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 32, 6 (Nov. 2013).
- Johannes Hanika, Marc Droske, and Luca Fascione. 2015. Manifold Next Event Estimation. *Computer Graphics Forum* 34, 4 (July 2015).
- Sebastian Herholz, Oskar Elek, Jiří Vorba, Hendrik Lensch, and Jaroslav Krivánek. 2016. Product Importance Sampling for Light Transport Path Guiding. *Computer Graphics Forum* 35, 4 (2016).
- David Immel, Michael Cohen, and Donald Greenberg. 1986. A radiosity method for non-diffuse environments. *Computer Graphics (Proc. SIGGRAPH)* 20, 4 (1986).
- Wenzel Jakob. 2013. *Light Transport on Path-Space Manifolds*. Ph.D. Dissertation. Cornell University.
- Wenzel Jakob and Steve Marschner. 2012. Manifold Exploration: A Markov Chain Monte Carlo Technique for Rendering Scenes with Difficult Spectral Transport. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 31, 4, Article 58 (July 2012).
- Adrian Jarabo, Carlos Aliaga, and Diego Gutierrez. 2018. A Radiative Transfer Framework for Spatially-Correlated Materials. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 37, 4 (2018).
- Wojciech Jarosz, Derek Nowrouzezahrai, Iman Sadeghi, and Henrik Wann Jensen. 2011a. A Comprehensive Theory of Volumetric Radiance Estimation Using Photon Points and Beams. *ACM Transactions on Graphics* 30, 1 (Feb. 2011).
- Wojciech Jarosz, Derek Nowrouzezahrai, Robert Thomas, Peter-Pike Sloan, and Matthias Zwicker. 2011b. Progressive Photon Beams. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 30, 6 (Dec. 2011).
- Wojciech Jarosz, Matthias Zwicker, and Henrik Wann Jensen. 2008. The Beam Radiance Estimate for Volumetric Photon Mapping. *Computer Graphics Forum (Proc. Eurographics)* 27, 2 (April 2008).
- Henrik Wann Jensen and Per H. Christensen. 1998. Efficient Simulation of Light Transport in Scenes With Participating Media Using Photon Maps. In *Annual Conference Series (Proc. SIGGRAPH)*.
- Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. 2001. A Practical Model for Subsurface Light Transport. *Annual Conference Series (Proc. SIGGRAPH)* 35 (2001).
- James T. Kajiya. 1986. The Rendering Equation. *Computer Graphics (Proc. SIGGRAPH)* 20, 4 (Aug. 1986).
- Markus Kettunen, Marco Manzi, Miika Aittala, Jaakko Lehtinen, Frédo Durand, and Matthias Zwicker. 2015. Gradient-domain Path Tracing. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 34, 4, Article 123 (July 2015).
- Christopher Kulla and Marcos Fajardo. 2012. Importance Sampling Techniques for Path Tracing in Participating Media. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 31, 4 (2012).
- Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 36, 4 (July 2017).
- Jaroslav Krivánek, Iliyan Georgiev, Toshiya Hachisuka, Petr Vévoda, Martin Šik, Derek Nowrouzezahrai, and Wojciech Jarosz. 2014. Unifying Points, Beams, and Paths in Volumetric Light Transport Simulation. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 33, 4 (July 2014).
- Eric Lafortune and Yves Willems. 1993. Bi-directional path tracing. In *Proc. Compu-graphics*.
- Eric Lafortune and Yves Willems. 1996. Rendering participating media with bidirectional path tracing. *Photorealistic Rendering Techniques (Proc. Eurographics Workshop on Rendering)* (1996).
- Laurence 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.
- Thomas Müller, Markus Gross, and Jan Novák. 2017. Practical Path Guiding for Efficient Light-Transport Simulation. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 36, 4 (June 2017).
- Ken Museth. 2013. VDB: High-resolution Sparse Volumes with Dynamic Topology. *ACM Transactions on Graphics* 32, 3, Article 27 (July 2013).
- Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo Methods for Volumetric Light Transport Simulation. *Computer Graphics Forum* 37, 2 (May 2018).
- Jan Novák, Andrew Selle, and Wojciech Jarosz. 2014. Residual Ratio Tracking for Estimating Attenuation in Participating Media. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 33, 6 (Nov. 2014).
- Mark Pauly, Thomas Kolig, and Alexander Keller. 2000. Metropolis Light Transport for Participating media. In *Rendering Techniques (Proc. Eurographics Workshop on Rendering)*.
- Ken Perlin. 1985. An Image Synthesizer. *Computer Graphics (Proc. SIGGRAPH)* 19, 3 (July 1985).
- Ken Perlin and Eric Hoffert. 1989. Hypertexture. *Computer Graphics (Proc. SIGGRAPH)* 23, 3 (July 1989).
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically Based Rendering: From Theory To Implementation* (3rd ed.).
- Matthias Raab, Daniel Seibert, and Alexander Keller. 2008. Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer.
- 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.
- László Szirmay-Kalos, Iliyan Georgiev, Milán Magdics, Balázs Molnár, and Dávid Légrády. 2017. Unbiased Estimators to Render Procedurally Generated Inhomogeneous Participating Media. *Computer Graphics Forum (Proc. Eurographics)* 36, 2 (2017).
- László Szirmay-Kalos, Milán Magdics, and Mateu Sbert. 2018. Multiple Scattering in Inhomogeneous Participating Media Using Rao-Blackwellization and Control Variates. *Computer Graphics Forum (Proc. Eurographics)* 37, 2 (May 2018).
- László Szirmay-Kalos, Balázs Tóth, and Milán Magdics. 2011. Free path sampling in high resolution inhomogeneous participating media. *Computer Graphics Forum* 30, 1 (2011).
- Eric Veach. 1997. *Robust Monte Carlo methods for light transport simulation*. Ph.D. Dissertation. Stanford, CA, USA.
- Eric Veach and Leonidas Guibas. 1995. Optimally combining sampling techniques for Monte Carlo rendering. *Annual Conference Series (Proc. SIGGRAPH)* 29 (1995).
- Jiří Vorba, Ondřej Karlík, Martin Šik, Tobias Ritschel, and Jaroslav Krivánek. 2014. On-line Learning of Parametric Mixture Models for Light Transport Simulation. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 33, 4 (Aug. 2014).
- Bruce Walter, Shuang Zhao, Nicolas Holzschuch, and Kavita Bala. 2009. Single Scattering in Refractive Media with Triangle Mesh Boundaries. *ACM Transactions on Graphics* 28, 3, Article 92 (July 2009).
- Alexander Wilkie, Sehera Nawaz, Mark Droske, Andrea Weidlich, and Johannes Hanika. 2014. Hero wavelength spectral sampling. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 33, 4 (June 2014).
- 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.
- C. D. Zerby, R. B. Curtis, and Hugo W. Bertini. 1961. *The relativistic doppler problem*. Technical Report ORNL-61-7-20. Oak Ridge National Laboratory, TN, USA.