

Monte Carlo Methods for Volumetric Light Transport Simulation

Jan Novák^{†1} Iliyan Georgiev² Johannes Hanika³ Wojciech Jarosz⁴

¹Disney Research ²Solid Angle ³Karlsruhe Institute of Technology ⁴Dartmouth College

Abstract

*The wide adoption of path-tracing algorithms in high-end realistic rendering has stimulated many diverse research initiatives. In this paper we present a coherent survey of methods that utilize Monte Carlo integration for estimating light transport in scenes containing participating media. Our work complements the volume-rendering state-of-the-art report by Cerezo et al. [CPP*05]; we review publications accumulated since its publication over a decade ago, and include earlier methods that are key for building light transport paths in a stochastic manner. We begin by describing analog and non-analog procedures for free-path sampling and discuss various expected-value, collision, and track-length estimators for computing transmittance. We then review the various rendering algorithms that employ these as building blocks for path sampling. Special attention is devoted to null-collision methods that utilize fictitious matter to handle spatially varying densities; we import two “next-flight” estimators originally developed in nuclear sciences. Whenever possible, we draw connections between image-synthesis techniques and methods from particle physics and neutron transport to provide the reader with a broader context.*

CCS Concepts

•**Computing methodologies** → **Computer graphics; Rendering; Ray tracing;**

1. Introduction

Since their inception in the Los Alamos National Laboratory, computer-accelerated Monte Carlo (MC) methods have gained popularity in many fields. Their relevance to computer graphics—image synthesis in particular—is indisputable and underscored by the many production renderers, such as Arnold, Hyperion, Manuka, or RenderMan, all based on stochastic sampling of light transport paths connecting light sources and camera sensors [CJ16, FHF*17]. The high dimensionality of transport problems, complexity of modern scenes, and demand for physical accuracy make MC the current method of choice in production rendering.

The quest for realism in visual effects and physical plausibility across the film and animation industry demands virtual scenes to be modeled with a high amount of detail and physical accuracy. Nearly every scene nowadays contains some volumetric elements that participate in light transport by absorbing, scattering, or emitting visible light. Participating media play an important role when modeling realistic materials that produce effects such as the subtle softening of illumination on human skin, characteristic translucency of fruits, volumetric caustics on aeration under water, or atmospheric scattering in fog, clouds, or planetary nebulae.

Much work has gone into simulating light propagation in media, and much has changed since the survey of Cerezo et al. [CPP*05]

from over a decade ago. More computational power enables holistic approaches that tie volumetric effects and surface scattering together and yield simpler authoring workflows. Methods that were previously assumed to be incompatible have been unified to allow renderers to benefit from each method’s respective strengths. Generally, investigations have shifted away from specialized solutions, e.g. for single- or multiple-scattering approximations or analytical methods such as diffusion.

A lot more focus is nowadays devoted to MC integration over the complete space of all light-transport paths. In this report we review these methods and provide a coherent summary of the current state of the art in MC rendering of scenes containing participating media. We focus on physically based approaches and draw connections to work in neutron transport to paint a more complete picture of how these algorithms have been developed. For the most part, we treat the media as isotropic (but with potentially anisotropic phase functions) and place emphasis on heterogeneous volumes (i.e. with spatially varying density). We leave out approximations based on diffusion, discrete ordinates, and methods for scientific volume visualization, and keep the report focused on MC methods for realistic rendering. We hope this report can provide guidance for practitioners implementing path-tracing based algorithms and also serve as a reference for researchers interested in applications of radiative transfer to image synthesis.

The remainder of this report is structured as follows: We first summarize the optical properties (Section 2) of volumetric matter, and then introduce the theoretical foundations (Section 3) that

[†] Author names appear in alphabetical order except for the first.

most MC methods rely upon, such as the radiative transfer equation, the volume rendering equation, and the path integral formulation of light transport. The following two sections describe the two fundamental building blocks of MC light-transport algorithms for volumes: methods for free-path sampling and the general concept of sampling a distance along a ray in a medium (Section 4), and methods for estimating the medium transmittance along a finite ray segment (Section 5). We also import two new transmittance estimators to the computer graphics community from the field of neutron transport. Next, we show how such distance- and transmittance-sampling techniques are used to construct entire light transport paths (Section 6) and review advanced algorithms based on such strategies (Section 7). We also briefly touch upon emissive media such as fire, explosions, or bioluminescence. Finally, we discuss different data structures (Section 8) for accelerating distance sampling and transmittance estimation, and conclude (Section 9) by listing a number of remaining challenges and open problems.

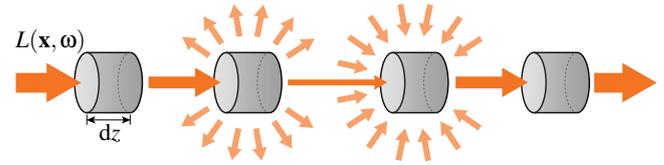
2. Optical properties

Before diving into the theory and practice of rendering scenes with participating media, we first discuss the relevant optical properties of matter.

Collision coefficients. We consider a participating medium to consist of a collection of microscopic particles which either scatter or absorb photons. Instead of modeling these particles explicitly, we describe them statistically by specifying their total *density* per unit volume ρ [m^{-3}] and their *cross-sectional areas*, σ_s and σ_a [m^2] for scattering and absorbing particles, respectively. By assuming that particle positions are *statistically independent*, we can multiply the individual cross-sections by the density to yield the *absorption coefficient* μ_a and *scattering coefficient* μ_s , each quantifying the local probability density [m^{-1}] of a photon undergoing the respective interaction per unit distance traveled. The *extinction coefficient* $\mu_t = \mu_a + \mu_s$ indicates the probability density of either type of event happening per unit distance. The ratio $\alpha = \mu_s/\mu_t$ quantifies the fraction of photons that undergo a scattering (vs. absorption) event at a collision, and it is commonly referred to as the single-scattering *albedo*.

Phase function. The directional density of light scattered at a point in a medium is given by the *phase function* f_p and generally depends on the incident and outgoing directions, with an optional spatial variation. If the medium scatters light uniformly, the phase function is *isotropic*, $f_p = 1/(4\pi)$, and otherwise it is classified as *anisotropic*. Some phase functions model the scattering density as axially symmetric, i.e. they parameterize f_p using only the angle (or its cosine) between the two directions. Examples of popular models are the Henyey-Greenstein [HG41], Rayleigh [Ray71], and Lorenz-Mie [Lor90, Mie08] phase functions; see Pegoraro [Peg16] for details and other scattering profiles.

Spatial dependence. If all of these medium properties are spatially invariant, the medium is said to be *homogeneous*. Otherwise, the medium is *inhomogeneous* or *heterogeneous*.



(a) Absorption (b) Out-scattering (c) In-scattering (d) Emission

Figure 1: Losses and gains of radiance due to absorption (a), scattering (b, c), and emission (d) within a differential volume element.

Directional dependence. When the collision coefficients μ_a and μ_s do not depend on the direction of light propagation and the phase function can be parameterized only by the angle between incident and scattered light, then the medium is said to be *isotropic*. Note that the phase function can still be anisotropic. If the collision coefficients or the phase function depend on the direction of incident or scattered light, i.e. the response of the material varies with the direction of propagation, the medium is referred to as *anisotropic*.

Assumptions. In the following, we describe mathematical foundations and algorithms for simulating light transport under a number of assumptions. As mentioned previously, we assume statistical independence of particle positions, which will give rise to the classical transport equations with exponential falloff of light. We further assume *isotropic media*. Most methods, however, can be analogously applied to anisotropic media by adjusting the definition of the collision coefficients and phase function to factor in the directional dependence [JAM*10]. We also assume that the refractive index only changes at medium/surface boundaries, which means that light travels along straight lines between interactions, ignoring effects like mirages. We further limit this survey to methods that simulate elastic scattering, where only a photon's direction, and not its energy (wavelength), may change at scattering events. This ignores effects like fluorescence and phosphorescence, and means that the transport at one wavelength can be simulated independently of other wavelengths. It is possible to incorporate all of these effects using anisotropic [JAM*10], refractive [ABW14], and multi-energy or *vector* [JA18] forms of the resulting transport equations. We encourage the interested reader to consult these works to fully comprehend the mathematics of light transport under these more general conditions.

3. Mathematical foundations for radiative transport

We now discuss mathematical formalizations of light-matter interactions and review the equations that form the basis of the MC algorithms presented later. We write points with boldface (e.g. \mathbf{x}), distances italicized (e.g. y), and sometimes use the same letter for a point and its distance from another point to emphasize the relation (e.g. $\mathbf{y} = \mathbf{x} + y\omega$, where ω is a direction vector).

3.1. Radiative transfer equation

The absorption and scattering interactions lead to three processes that occur as light propagates along a ray: absorption, out-scattering, and in-scattering, which are illustrated in Figure 1. Also

accounting for medium emission, the change in radiance traveling along direction ω through a differential volume element at point \mathbf{x} is described by the *radiative transfer equation* (RTE) [Cha60]:

$$(\omega \cdot \nabla)L(\mathbf{x}, \omega) = -\mu_a(\mathbf{x})L(\mathbf{x}, \omega) - \mu_s(\mathbf{x})L(\mathbf{x}, \omega) + \mu_s(\mathbf{x})L_s(\mathbf{x}, \omega) + \mu_a(\mathbf{x})L_e(\mathbf{x}, \omega), \quad (1)$$

where the term

$$L_s(\mathbf{x}, \omega) = \int_{S^2} f_p(\omega, \bar{\omega})L_i(\mathbf{x}, \bar{\omega})d\bar{\omega} \quad (2)$$

is commonly referred to as in-scattered radiance, which collects the incident radiance L_i from all directions on the unit sphere S^2 . The first two terms in Equation (1) correspond to losses due to absorption and out-scattering, and the third and the fourth term represent gains due to in-scattering and volumetric emission, respectively.

Integrating both sides of the differential RTE (1) along the direction ω yields the integral form:

$$L(\mathbf{x}, \omega) = \int_0^\infty T(\mathbf{x}, \mathbf{y}) [\mu_a(\mathbf{y})L_e(\mathbf{y}, \omega) + \mu_s(\mathbf{y})L_s(\mathbf{y}, \omega)] dy, \quad (3)$$

which integrates emitted and in-scattered light along the ray $(\mathbf{x}, -\omega)$; here $\mathbf{y} = \mathbf{x} - y\omega$. The *transmittance* $T(\mathbf{x}, \mathbf{y})$ is obtained by folding the loss of light along a line due to absorption and out-scattering into a single differential process and integrating it along ω , yielding

$$T(\mathbf{x}, \mathbf{y}) = e^{-\int_0^y \mu(\mathbf{x}-s\omega) ds}. \quad (4)$$

This formula is also known as the Beer-Lambert law [Lam60], and the integral in the exponent is called the *optical thickness* τ . While transmittance is a function of two points, we will often simplify the notation and denote it as a function of just a scalar distance t : $T(t) = e^{-\tau(t)} = e^{-\int_0^t \mu(\mathbf{x}-s\omega) ds}$.

In Equation (1), we adopt the convention of modulating the self-emitted radiance L_e by the absorption coefficient. This is motivated by the mere ease of succinctly formalizing certain concepts introduced later. It is worth noting that some literature uses a dedicated coefficient for emission or removes the coefficient altogether. Dropping the coefficient, however, requires L_e to have different units than L . We thus recommend either using the absorption (or emission) coefficient to force all radiance functions be in $W/m^2/sr$, or defining a source term that cannot be confused with emitted radiance. Lastly, the use of the absorption coefficient reinforces the notion of incandescence and luminescence: light can only be emitted where energy is first absorbed.

3.2. Volume rendering equation

Typical scenes contain not only participating media but also objects with hard boundaries. Equation (3) then needs to be extended to accommodate for light interactions with object surfaces. The radiative equilibrium at a surface point \mathbf{z} is described by the surface rendering equation [Kaj86]:

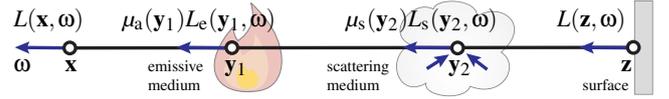
$$L(\mathbf{z}, \omega) = L_e(\mathbf{z}, \omega) + \int_{S^2} f_r(\mathbf{z}, \omega, \bar{\omega})L_i(\mathbf{z}, \bar{\omega})|n(\mathbf{z}) \cdot \bar{\omega}|d\bar{\omega}, \quad (5)$$

where $L_e(\mathbf{z}, \omega)$ represents the radiance emitted by the surface, $f_r(\mathbf{z}, \omega, \bar{\omega})$ is the bidirectional scattering distribution function (BSDF) which relates the differential outgoing radiance $dL(\mathbf{z}, \omega)$

to the incident radiance $L_i(\mathbf{z}, \bar{\omega})$, and $|n(\mathbf{z}) \cdot \bar{\omega}|$ is a foreshortening term due to the angle between the incident radiance direction $\bar{\omega}$ and the surface normal $n(\mathbf{z})$.

Equation (5) provides the boundary condition for truncating the integration bounds of Equation (3) to the nearest surface at distance z . For the surface point $\mathbf{z} = \mathbf{x} - z\omega$, the resulting *volume rendering equation* (VRE) reads

$$L(\mathbf{x}, \omega) = \int_0^z T(\mathbf{x}, \mathbf{y}) [\mu_a(\mathbf{y})L_e(\mathbf{y}, \omega) + \mu_s(\mathbf{y})L_s(\mathbf{y}, \omega)] dy + T(\mathbf{x}, \mathbf{z})L(\mathbf{z}, \omega). \quad (6)$$



The integral represents light emitted and in-scattered along the line between point \mathbf{x} and the surface point \mathbf{z} . The second term amounts to the medium-attenuated exitant radiance from the surface given by Equation (5). Equation (6) is also called the *equation of transfer* in some literature.

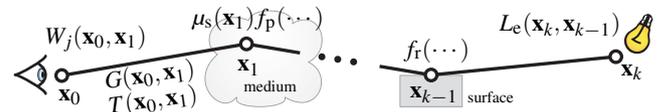
3.3. Path integral formulation

Most existing MC methods for light-transport simulation operate by sampling random light-transport trajectories, or *light paths*, in the scene and evaluating the amount of energy they bring from the light sources all the way to the camera. Veach formalized this concept in his path integral formulation of light transport [Vea97], which arises from the expansion of the recursive Equation (6) using the Neumann series [Vea97, PKK00] and formalizing the flux measurement on a camera sensor.

In the path integral formulation, the value of pixel j in the rendered image I can be described as a conceptually simple integral:

$$I_j = \int_{\mathcal{P}} f_j(\bar{\mathbf{x}})d\bar{\mathbf{x}}. \quad (7)$$

Here, $\bar{\mathbf{x}} = (\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k) \in \mathcal{P}$ is a transport path of length k (segments) and \mathcal{P} is the space of all such possible paths in the scene. The integral is to be understood as summing up contributions of all path lengths $k \in [1, \infty)$. The measurement contribution function $f_j(\bar{\mathbf{x}})$ (defined below) gives the pixel response to the differential flux carried by the given path. The differential path measure $d\bar{\mathbf{x}} = \prod_{i=0}^k d\mathbf{x}_i$ is the product of measures at individual vertices, each being either the differential-volume or the differential-area measure depending on whether \mathbf{x}_i is in a medium or on a surface.



The measurement contribution for length- k paths is defined as

$$f_j(\bar{\mathbf{x}}) = W_j(\mathbf{x}_0, \mathbf{x}_1)L_e(\mathbf{x}_k, \mathbf{x}_{k-1})G(\mathbf{x}_0, \mathbf{x}_1)T(\mathbf{x}_0, \mathbf{x}_1) \cdot \prod_{i=1}^{k-1} f_s(\mathbf{x}_{i-1}, \mathbf{x}_i, \mathbf{x}_{i+1})G(\mathbf{x}_i, \mathbf{x}_{i+1})T(\mathbf{x}_i, \mathbf{x}_{i+1}), \quad (8)$$

where $W_j(\mathbf{x}_0, \mathbf{x}_1)$ is the response of pixel j , and

$$G(\mathbf{x}, \mathbf{y}) = \frac{D(\mathbf{x}, \mathbf{y})D(\mathbf{y}, \mathbf{x})}{\|\mathbf{x} - \mathbf{y}\|^2}, \quad \text{where} \quad (9)$$

$$D(\mathbf{x}, \mathbf{y}) = \begin{cases} |n(\mathbf{x}) \cdot \omega_{\mathbf{x} \rightarrow \mathbf{y}}| & \text{if } \mathbf{x} \text{ is on a surface,} \\ 1 & \text{if } \mathbf{x} \text{ is in a medium.} \end{cases} \quad (10)$$

$$L_e(\mathbf{x}, \mathbf{y}) = \begin{cases} L_e(\mathbf{x}, \omega_{\mathbf{x} \rightarrow \mathbf{y}}) & \text{if } \mathbf{x} \text{ is on a surface,} \\ \mu_a(\mathbf{x})L_e(\mathbf{x}, \omega_{\mathbf{x} \rightarrow \mathbf{y}}) & \text{if } \mathbf{x} \text{ is in a medium.} \end{cases} \quad (11)$$

$$f_s(\mathbf{x}, \mathbf{y}, \mathbf{z}) = \begin{cases} f_r(\omega_{\mathbf{y} \rightarrow \mathbf{x}}, \mathbf{y}, \omega_{\mathbf{y} \rightarrow \mathbf{z}}) & \text{if } \mathbf{y} \text{ is on a surface,} \\ \mu_s(\mathbf{y})f_p(\omega_{\mathbf{y} \rightarrow \mathbf{x}}, \mathbf{y}, \omega_{\mathbf{y} \rightarrow \mathbf{z}}) & \text{if } \mathbf{y} \text{ is in a medium.} \end{cases} \quad (12)$$

Most of these terms may have an implicit dependency on the wavelength of light, and the integral over the measurement contribution function may be evaluated for each wavelength separately to yield color images. Note that these equations do not account for cross-talk between wavelengths, i.e. fluorescence [JA18].

3.4. Monte Carlo integration

Virtually all modern high-quality physically based rendering engines estimate the aforementioned transport integrals via MC integration. We use angle brackets $\langle \cdot \rangle$ to denote an MC estimator of some quantity and write *primary*, i.e. one-sample, estimators for brevity: $\langle F \rangle = f(x)/p(x)$, where $f(x)$ is the integrand of F and $p(x)$ is the probability density function (PDF) for sampling points x . By averaging N independent realizations of a primary estimator, one can obtain a *secondary* (i.e. multi-sample) estimator.

Some algorithms discussed later (e.g. unidirectional volumetric path tracing) adopt a *local* view of the transport and estimate the amount of radiance arriving at point \mathbf{x} on a sensor from direction ω by applying MC estimation to the volume rendering equation (6):

$$\langle L(\mathbf{x}, \omega) \rangle = \frac{T(\mathbf{x}, \mathbf{y})}{p(\mathbf{y})} \left[\mu_a(\mathbf{y})L_e(\mathbf{y}, \omega) + \mu_s(\mathbf{y})L_s(\mathbf{y}, \omega) \right] + \frac{T(\mathbf{x}, \mathbf{z})}{P(z)} L(\mathbf{z}, \omega), \quad (13)$$

where $p(\mathbf{y})$ is the PDF of sampling point \mathbf{y} , which is y units away from \mathbf{x} , and where the nearest-surface radiance can be optionally evaluated with probability $P(z)$. This estimator requires two main routines: one for sampling distances along the ray, and one for evaluating the transmittance T between two given points; we discuss these in Section 4 and Section 5 respectively.

The recursive nature of Estimator (13) promotes a somewhat localized view on the light-transport simulation—we consider one path segment at a time. In contrast, the path integral framework provides a more holistic, *global* view on the transport problem, highlighting the existence of multiple ways of constructing the same path. An estimator of Equation (7) has the following general form:

$$\langle I_j \rangle = \frac{f_j(\bar{\mathbf{x}})}{p(\bar{\mathbf{x}})}, \quad (14)$$

where $\bar{\mathbf{x}}$ is a randomly sampled path with probability density $p(\bar{\mathbf{x}})$ given by the procedure, or technique, used to construct its vertices. For a path $\bar{\mathbf{x}}$ of some length $k \geq 1$, the contribution $f_j(\bar{\mathbf{x}})$ is well-defined and given by Equation (8).

In the path-integral framework, unbiased simulation methods thus differ only in the path-sampling techniques they employ and their corresponding path PDFs $p(\bar{\mathbf{x}})$. This “flat” view of light transport enables the use of global sampling techniques that coordinate the sampling of entire sequences of vertices, in contrast to the local techniques that sample one vertex at a time. Examples of such techniques are joint importance sampling and Metropolis sampling discussed in Section 6, and some of the bidirectional algorithms reviewed in Section 7. Such algorithms often make use of multiple importance sampling (MIS) [Vea97], which is a way of optimally weighting estimators based on their PDF to reduce variance.

Regardless of the individual sampling strategy, all techniques require two basic building blocks. The first one is to sample a distance and is covered in Section 4. This is essential in the common case where a path is constructed incrementally by successively extending it from the sensor to the lights. More abstract MC techniques, which do not perform analog random walks, need to evaluate the measurement-contribution function and thus the transmittance between two points. This is sometimes possible in closed form, and sometimes using a stochastic estimate (cf. Section 5).

4. Distance sampling

In this and the following section, we discuss techniques for sampling distances and estimating transmittance along a ray. To classify distance-sampling methods, we borrow the terminology of *analog* and *non-analog* estimators from the field of neutron transport, and categorize the algorithms according to whether they strictly adhere to the physical process of light propagation or not, respectively.

Analog methods sample the distance to the next light-medium collision along the line of flight analogously to how photons interact with materials in the real world. The sampling procedure is in such cases commonly referred to as *free-path sampling* or *free-flight-distance sampling*. The distance distribution strictly adheres to the Beer-Lambert law, i.e. it has a PDF proportional to the transmittance along the given ray. Sampling can be explicit, via an inversion of the corresponding cumulative distribution function (CDF), or implicit, through probabilistic reasoning as in null-collision algorithms (discussed below). Since the distribution of particles inside the medium is assumed to be random, the free-flight distance is a random variable.

Non-analog methods have been developed to improve sampling efficiency over analog methods. They use sampling distributions that deviate from the true distribution of free paths, which is “corrected” by appropriately weighting the samples.

Both analog and non-analog methods share the common theme of sampling distances according to a certain PDF. In the following, we first review analytic and semi-analytic analog methods for media that permit free-path sampling in closed form or through a simple iterative process. Next, we discuss rejection-based analog estimators that rely on so-called *null* collisions, and review their non-analog variants that lift certain constraints and further extend the class of supported media. Lastly, we describe non-analog approaches that draw distances from carefully designed PDFs to reduce the variance in specific scenarios.

Table 1: Various expressions used for closed-form sampling of free paths in homogeneous and exponentially-varying media [BM01].

Medium	$\mu_t(t)$	$T(t)$	$p(t)$	$F(t)$	$F^{-1}(\xi)$
Homogeneous (infinite)	c	e^{-ct}	ce^{-ct}	$1 - e^{-ct}$	$-\ln(1 - \xi)/c$
Homogeneous (bounded)	$c : t < b$	e^{-ct}	$ce^{c(b-t)} / (e^{cb} - 1)$	$(e^{cb} - e^{c(b-t)}) / (e^{cb} - 1)$	$-\ln(1 - \xi(1 - e^{-cb})) / c$
Exp. decreasing (infinite)	$ce^{-at}, a > 0$	$\exp(-\frac{c}{a}(1 - e^{-at}))$	$\mu_t(t) \exp(-\frac{c}{a}(1 - e^{-at}))$	$1 - \exp(-\frac{c}{a}(1 - e^{-at}))$	$-\ln(1 - a/c \ln(\xi)) / a$

4.1. Analytic and semi-analytic free-path sampling

For estimators of the form in Equation (13), a sensible importance sampling strategy is to use a distance PDF that is proportional to transmittance. The corresponding sampling procedure can be derived from the Beer-Lambert law (Equation (4)). The law gives the fraction of particles—the transmittance $T(t)$ —that will, in expectation, travel beyond a certain distance t without colliding with the medium. Denoting X a random free-path variable, we thus demand $P(X > t) = T(t)$. Note that the CDF of sampled distances $F(t)$ is defined as the complementary probability, $P(X \leq t)$, hence

$$F(t) = 1 - T(t). \quad (15)$$

The recipe for sampling free paths can be obtained by inverting the CDF. Algorithmically, the procedure amounts to generating a random number ξ and finding the position along the ray where $1 - T(t)$ equals ξ . In the remainder of this subsection we discuss finding this location in closed form and via an iterative search. In both cases, the probability density (PDF) of each sample can be computed by differentiating the CDF:

$$p(t) = \frac{dF(t)}{dt} = \frac{d}{dt}(1 - e^{-\tau(t)}) = \frac{d\tau(t)}{dt} e^{-\tau(t)} = \mu_t(t) e^{-\tau(t)}. \quad (16)$$

4.1.1. Closed-form methods

In homogeneous media, where the extinction μ_t is spatially invariant, optical thickness is linearly proportional to the distance t along the ray: $\tau(t) = t\mu_t$. The free-path sampling CDF thus simplifies to

$$F(t) = 1 - e^{-\mu_t t}, \quad (17)$$

which can be easily solved for t , yielding an analytic formula for sampling free paths using uniform random numbers $\xi \in [0, 1)$,

$$t(\xi) = -\frac{\ln(1 - \xi)}{\mu_t}, \quad (18)$$

with PDF

$$p(t) = \mu_t e^{-\mu_t t}. \quad (19)$$

An alternative approach for deriving the same sampling routine is to first define a PDF that is proportional to transmittance. The μ_t factor in the PDF expression comes from normalizing the transmittance function, i.e. dividing it by its zeroth moment—the mean free path $1/\mu_t$. The CDF is then expressed by integrating the PDF, and its inversion yields the same closed-form solution for $t(\xi)$.

The former approach of defining the distance-sampling CDF via the transmittance function was first applied in the context of particle transport by Stanislaw Ulam in a letter to John von Neumann, laying down one of the cornerstones of solving transport problems

using MC integration. A similar derivation can be carried out for inhomogeneous media as well, the only requirement—albeit a strong one—is the existence of a differentiable and invertible CDF; a few examples are listed in Table 1. Brown and Martin [BM03] propose to use the iterative Newton-Raphson method to find the solution $t(\xi)$ numerically when closed-form inversion is not possible.

The above procedure assumes the medium is infinite. In the presence of solid objects, one option is to clamp the PDF to zero for distances beyond the nearest surface and re-derive the CDF (see Table 1). Alternatively, we can handle the interactions with the surface. The probability of hitting the surface located at distance s is equal to the probability of sampling a random distance t beyond s . This is simply the integral of the PDF from s to infinity, which by definition equals to the transmittance at s ; for homogeneous media

$$P(t > s) = \int_s^\infty p(t) dt = 1 - F(s) = e^{-\mu_t s}. \quad (20)$$

4.1.2. Regular tracking

Certain scenes can be well described as collections of “simple” volumes for which closed-form free-path sampling routines exist. Examples of such scenes are pebble-bed reactors with gas-cooled graphite pebbles, layered materials such as human skin, or anything that can be approximated well by a piecewise-constant spatial subdivision. Since each volume element allows for analytic sampling, all we need is a way to carry samples across its boundaries in case the sampled distance exceeds them. This method is known as *regular tracking* [SBB*99] or *surface tracking* [Lep10] and works as follows.

Given a ray and a sampled value of transmittance $\xi \in (0, 1]$, we sweep along the ray and iterate over individual partitions until the accumulated optical thickness first exceeds the value that corresponds to the sampled transmittance. The exact location within the last partition is found analytically such that the total optical thickness τ up to this point yields the exact transmittance value, i.e. $e^{-\tau} = \xi$. Figure 2 provides an illustration of this algorithm.

The main strength of regular tracking is its correctness, i.e. the distribution of free paths is exactly proportional to transmittance: $p(t) = \mu_t(t)T(t)$. Its main drawback is the need to “discover” all boundaries along the free-path sample. This can be expensive when free paths cross many boundaries, unless the crossings can be found quickly, such as with uniform grids [AW87], octrees [HTAT*06], or other regular spatial-subdivision structures.

4.1.3. Ray marching

In order to reduce the cost of regular tracking, one can simply ignore the boundaries and march along the ray with fixed-size steps [PH89]. This significantly simplifies the implementation. At

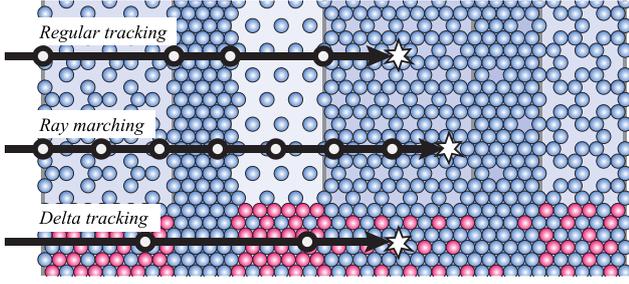


Figure 2: Illustration of free-path sampling in a medium composed of several homogeneous regions; blue and pink particles represent real and fictitious matter, respectively. Regular tracking finds all boundary crossings and solves for the position analytically. Ray marching steps with a constant stride until it accumulates the sampled optical thickness. Delta tracking first homogenizes the volume using fictitious matter and then compensates for that by explicitly handling the resulting null collisions with that matter.

each step, the algorithm queries the local medium extinction and then moves forward by a fixed distance (see Figure 2). The optical thickness is accumulated assuming the extinction is constant or linear between the marching steps. Once the marching exceeds the desired value of optical thickness, the algorithm retracts and solves for the sampled location between the last two steps, again assuming constant or linear extinction.

Approximating the actual extinction with a piecewise constant/linear function causes the distance distribution to deviate from the true, free-path distribution. The bias persists even if the stepping is randomly jittered [RSK08] and can become particularly objectionable when the medium contains thin features, which the stepping is likely to miss. In cases where the frequency spectrum of the extinction function is band-limited, one should conform to the Shannon sampling theorem and increase the stepping to $2 \times$ the Nyquist frequency. This can, however, quickly become expensive if the extinction changes rapidly along the ray. The cost can be reduced by adapting the step size locally or by employing nested and high-order ray-marching schemes [Muñ14].

4.2. Null-collision methods

The first null-collision methods were developed in the field of neutron transport and plasma physics to address the challenge of sampling free paths in heterogeneous media. These methods are based on von Neumann’s idea of generating arbitrary sample distributions via rejection sampling [vN51] and can handle arbitrary, e.g. procedurally generated, heterogeneity. They are applicable to a wider class of volumes than regular tracking and do not bias the free-path distribution like ray marching. The most limiting drawback of these methods is their inability to quantify the PDF of individual samples.

The main idea of null-collision methods is to introduce *fictitious* matter that has no impact on the light transport but enables the use of the (semi-) analytic distance sampling routines we discussed in Section 4.1. The fictitious matter exhibits so-called *null collisions* upon which light continues traveling forward with unchanged energy. We can vary the amount of fictitious matter spatially—for in-

stance to homogenize the total (real + fictitious) density—to enable sampling collisions in closed form. This idea of adding fictitious matter is common to all null-collision algorithms. Individual methods vary in the amount of fictitious matter added and the procedure for sampling collisions with the altered medium.

4.2.1. Delta tracking

To best of our knowledge, the first description of delta tracking was provided by Butcher and Messel [BM58, BM60] and Zerby et al. [ZCB61] in an application of Monte Carlo methods to electron-photon cascades. The technique was later detailed in the doctoral thesis of Bertini [Ber63] and independently mentioned by Woodcock et al. [WMHL65] introducing the term fictitious collision. Skullerud [Sku68] developed an analogous technique for applications in plasma physics. Due to the many origins and applications, the method is known under several names, the most popular being delta tracking, Woodcock tracking, pseudo scattering, and the null-collision algorithm. Despite being very intuitive, it took several years until the correctness of the algorithms was proved rigorously [Mil67, Col68]. Recently, Galtier et al. [GBC*13] presented a derivation of the algorithm directly from the RTE, demonstrating not only the correctness, but also providing a convenient framework for postulating new variants of the algorithm.

Delta tracking achieves the correct free-path distribution by artificially increasing the collision rate and rejecting some of the collisions in the spirit of a zero-sum game. The impact of the fictitious matter can be formalized with the following differential equation:

$$-\mu_n(\mathbf{x})L(\mathbf{x}, \omega) + \mu_n(\mathbf{x}) \int_{S^2} \delta(\omega - \bar{\omega})L(\mathbf{x}, \bar{\omega}) d\bar{\omega} = 0, \quad (21)$$

which states that the losses and gains due to null-collisions perfectly cancel out leaving the radiation field intact. Adding the left-hand side of Equation (21) to the right-hand side of Equation (1), integrating the resulting differential equation spatially, and solving the Dirac delta integral of null collisions yields:

$$L(\mathbf{x}, \omega) = \int_0^\infty T_{\bar{\mu}}(\mathbf{x}, \mathbf{y}) \left(\underbrace{\mu_a(\mathbf{y})L_e(\mathbf{y}, \omega)}_{\text{emission}} + \underbrace{\mu_s(\mathbf{y})L_s(\mathbf{y}, \omega)}_{\text{scattering}} + \underbrace{\mu_n(\mathbf{y})L(\mathbf{y}, \omega)}_{\text{null-collisions}} \right) dy, \quad (22)$$

where $T_{\bar{\mu}}(\mathbf{x}, \mathbf{y}) = e^{-\int_0^y \mu_a(\mathbf{z}) + \mu_s(\mathbf{z}) + \mu_n(\mathbf{z}) dz}$; we refer to the sum $\mu_a(\mathbf{x}) + \mu_s(\mathbf{x}) + \mu_n(\mathbf{x})$ as the total collision coefficient $\bar{\mu}(\mathbf{x})$. We can then turn Equation (22) into an MC estimator,

$$\langle L(\mathbf{x}, \omega) \rangle = \frac{T_{\bar{\mu}}(\mathbf{x}, \mathbf{y})}{p_{\bar{\mu}}(\mathbf{y})} \left[\mu_a(\mathbf{y})L_e(\mathbf{y}, \omega) + \mu_s(\mathbf{y})L_s(\mathbf{y}, \omega) + \mu_n(\mathbf{y})L(\mathbf{y}, \omega) \right], \quad (23)$$

where $p_{\bar{\mu}}$ denotes the PDF for sampling collisions in the combined medium. Since the fictitious matter does not impact light transport, we can freely modulate its density to ensure the collision sampling can be done in closed form, e.g. by making the total collision coefficient constant and sampling from $p_{\bar{\mu}}$ by using Equation (18) with $\bar{\mu}$ instead of μ_t . In the remainder of this section, we will focus on volumes assuming the scene contains no surfaces and use Estimator (23) as a simplified version of Estimator (13); see the work by Eymet et al. [EPG*13] for a treatment of scenes with surfaces.

Having sampled a tentative collision at \mathbf{y} , we now have three

radiance components to evaluate in Equation (23). Delta tracking randomly chooses one of them with probabilities proportional to the corresponding coefficient:

$$P_a(\mathbf{y}) = \frac{\mu_a(\mathbf{y})}{\bar{\mu}(\mathbf{y})}, \quad P_s(\mathbf{y}) = \frac{\mu_s(\mathbf{y})}{\bar{\mu}(\mathbf{y})}, \quad P_n(\mathbf{y}) = \frac{\mu_n(\mathbf{y})}{\bar{\mu}(\mathbf{y})}, \quad (24)$$

classifying the collision as either absorptive, scattering, or null, respectively. The absorptive and scattering events represent real collisions and yield valid free-path samples in the original medium. The null collision, on the other hand, requires estimating L by recursively invoking the estimator from \mathbf{y} . Delta tracking thus steps forward along the line by sampling *tentative* collisions until one is probabilistically classified as real. The sampled free-flight distance y is then the one from the ray origin to the first real collision.

Table 2 summarizes the properties of up-to-now described methods for sampling free paths. Below, we discuss the main drawbacks and challenges of delta tracking in detail.

Unknown PDF value. Delta tracking casts the problem of finding a collision in the original medium as a problem of finding a real collision in the combined medium. The distribution of distances to the *first* real collision in the combined medium corresponds to free-path distribution in the original medium, i.e. the PDF of free-path samples produced by delta tracking is $\mu_t(\mathbf{y})T(\mathbf{x}, \mathbf{y})$ [Col68, Mil67]. However, the main drawback of delta tracking is that, because of the exponentiated heterogeneous integral in $T(\mathbf{x}, \mathbf{y})$, we cannot compute the actual value of the PDF. Looking back at Equation (13), this does not seem to be a problem since the $T(\mathbf{x}, \mathbf{y})$ in the PDF is conveniently canceled out by the transmittance term in the numerator. However, the inaccessibility of the PDF value prevents us from combining delta tracking with other distance-sampling techniques via multiple importance sampling [Vea97], unless we base the MIS weight on an approximation of the PDF obtained using e.g. ray marching. The difficulty of calculating the PDF stems from the fact that the real collision can follow after an arbitrarily long chain of null collisions. The (marginal) PDF of the real collision can thus be calculated only by integrating over an infinite set of null-collision chains, which is practically infeasible.

Stratification and quasi-random numbers. Using stratified random numbers for sampling, such as quasi-Monte Carlo sequences or blue-noise patterns, requires a fixed mapping from random-number dimensions to sampling techniques that use them. Because methods like delta tracking require a variable number of random numbers, such a mapping is hard to derive. One strategy to employ some stratification is to map two stratified dimensions to sampling of the first tentative collision (distance and probabilistic classification), and construct the remaining collisions (if any) with independent uniform random numbers [RSK08].

Overhead of null collisions. Each tentative collision requires an evaluation of spatially varying coefficients that often depend on costly procedurals or slow memory access. To maximize the performance, one should therefore add no more fictitious matter than absolutely necessary to enable closed-form sampling. This amounts to finding a tight *majorant* bound of the extinction function that permits analytic collision sampling, which has been addressed using various spatial-subdivision schemes that we discuss in Section 8.

4.2.2. Weighted delta tracking

In this section, we discuss a generalization of delta tracking for postulating non-analog variants of the algorithm, which sample distances with PDFs not necessarily proportional to transmittance. To counteract the deviation from the analog process, non-analog methods weight each sample producing a *weighted* distance distribution that still corresponds to the distribution of true free paths. Beside other applications, the weighting provides means to handle media with *negative* extinction coefficients as well as cases where $\bar{\mu}(\mathbf{x})$ is not a majorant of $\mu_t(\mathbf{x})$, i.e. $\mu_n(\mathbf{x}) < 0$ for some \mathbf{x} . Lastly, weighted tracking allows for a more efficient estimation of transmittance (c.f. Section 5) and handling of multiple wavelengths (c.f. Section 6.5).

For notation brevity, we rewrite the estimator in Equation (23) as

$$\langle L(\mathbf{x}, \omega) \rangle = \frac{T_{\bar{\mu}}(\mathbf{x}, \mathbf{y})}{p_{\bar{\mu}}(\mathbf{y})} \sum_{\mathcal{C}} \mu_{\star}(\mathbf{y}) L_{\star}(\mathbf{y}, \omega), \quad (25)$$

where $\mathcal{C} = \{(\mu_a, L_e), (\mu_s, L_s), (\mu_n, L)\}$, and (μ_{\star}, L_{\star}) represents one entry in \mathcal{C} .

In order to formalize the probabilistic evaluation of individual radiance components (akin to Russian roulette), we can use the following notation:

$$\langle f(x) \rangle = \mathcal{H}[P - \xi] \frac{f(x)}{P}, \quad (26)$$

where $\xi \in (0, 1]$ is a uniform random number, P is the probability of evaluating $f(x)$, and \mathcal{H} is the Heaviside function. The estimator can then be rewritten as

$$\langle L(\mathbf{x}, \omega) \rangle = \frac{T_{\bar{\mu}}(\mathbf{x}, \mathbf{y})}{p_{\bar{\mu}}(\mathbf{y})} \sum_{\mathcal{C}} \mathcal{H}[P_{\star}(\mathbf{y}) - \xi_{\star}] \frac{\mu_{\star}(\mathbf{y})}{P_{\star}(\mathbf{y})} L_{\star}(\mathbf{y}, \omega), \quad (27)$$

where \star again marks quantities associated with a specific collision type. The values of individual ξ_{\star} are typically anti-correlated to ensure that only one component gets evaluated. For clarity, we rewrite Estimator (27) by concatenating the two fractions into a single local collision weight $w_{\star}(\mathbf{x})$:

$$\langle L(\mathbf{x}, \omega) \rangle = \sum_{\mathcal{C}} \mathcal{H}[P_{\star}(\mathbf{y}) - \xi_{\star}] w_{\star}(\mathbf{x}) L_{\star}(\mathbf{y}, \omega), \quad (28)$$

$$w_{\star}(\mathbf{x}) = \frac{T_{\bar{\mu}}(\mathbf{x}, \mathbf{y})}{p_{\bar{\mu}}(\mathbf{y})} \frac{\mu_{\star}(\mathbf{y})}{P_{\star}(\mathbf{y})}. \quad (29)$$

In standard delta tracking, the probabilities $P_{\star}(\mathbf{x})$ are proportional to the respective collision coefficient as in Equations (24), and the distance sampling PDF is $p_{\bar{\mu}}(\mathbf{y}) = \bar{\mu}(\mathbf{y}) T_{\bar{\mu}}(\mathbf{x}, \mathbf{y})$. Taking these into account, we can verify that the collision weight w_{\star} is always 1 because the transmittance $T_{\bar{\mu}}$ cancels out with the tentative-collision PDF $p_{\bar{\mu}}$ up to the factor $1/\bar{\mu}$, which together with μ_{\star} cancel out the probability P_{\star} . This confirms that probabilities that are proportional to the collision coefficients yield an analog estimator.

Several works [Cra78, GBC*13, EPG*13, NSJ14, SKGM*17, KHLN17] made the observation that probabilities P_{\star} do not in fact need to be proportional to the respective collision coefficients μ_{\star} . As long as they are valid probabilities permitting probabilistic evaluation, i.e. $P_{\star} \in (0, 1]$, they can be set arbitrarily. While it may be difficult to interpret the physical process in such case, the mathematics still hold and tailoring the probabilities to the problem at hand can often improve sampling efficiency.

Table 2: Overview of the properties of different analog methods for sampling free paths.

Closed-form tracking	Regular tracking	Ray marching	Delta tracking
simple media such as homogeneous, linear, or exponential	piecewise “simple” media	any media	any media with bounded $\mu_t(\mathbf{x})$
efficient but highly limited	iterative, inefficient if free paths cross many boundaries	iterative, inefficient for media with high frequencies	iterative, inefficient if too much fictitious matter is added
unbiased	unbiased	biased	unbiased
known PDF	known PDF	known (approximate) PDF	unknown PDF

Since deviations from the analog distribution are corrected by weighting the samples, it is illustrative to write out the weights applied to real collisions. Assuming a real collision i occurs at \mathbf{x}_i after $(i - 1)$ null collisions, the weight induced by the chain of null collisions reads:

$$w(\mathbf{x}_1 \cdots \mathbf{x}_{i-1}) = \prod_{j=1}^{i-1} \frac{T_{\bar{\mu}}(\mathbf{x}_{j-1}, \mathbf{x}_j) \mu_n(\mathbf{x}_j)}{P_{\bar{\mu}}(\mathbf{x}_j)} = \prod_{j=1}^{i-1} \frac{\mu_n(\mathbf{x}_j)}{\bar{\mu}(\mathbf{x}_j) P_n(\mathbf{x}_j)}. \quad (30)$$

This provides us with a recipe for weighting real collisions to obtain a weighted distribution that is equivalent to that of free paths in the original volume.

Weighted variants of delta tracking have been proposed to address various transport problems. Early research in neutron transport focused on increasing collision rates in measurement regions [Ste66, SG69] or artificially stretching the free-path distribution in deep-penetration problems [Cra78, MK15]. Specific schemes were also developed for handling negative values of (null-) collision coefficients [CCT72, Cra78, GBD*16, SKGM*17]. These variants allow violating the assumption that $\bar{\mu}(\mathbf{x})$ is a majorant of $\mu_t(\mathbf{x})$, which can be sometimes desirable. For instance, when the volume contains a small region with high extinction, a bounding $\bar{\mu}$ will lead to many null collisions everywhere else. Decreasing $\bar{\mu}$ reduces the number of null collisions at the cost of negative weights. The ability to handle $\bar{\mu} < \mu_t(\mathbf{x})$ is very practical also for precomputing acceleration data structures in situations when the maximum value of extinction in a finite region cannot be found exactly. Unfortunately, negative weights will rapidly increase variance, so one should allow $\bar{\mu} < \mu_t(\mathbf{x})$ only when absolutely necessary.

Non-bounding $\bar{\mu}(\mathbf{x})$ leading to negative $\mu_n(\mathbf{x})$ can be handled e.g. by setting the collision probabilities to

$$P_a(\mathbf{x}) = \frac{\mu_a(\mathbf{x})}{\mu_t(\mathbf{x}) + |\mu_n(\mathbf{x})|}, \quad (31)$$

$$P_s(\mathbf{x}) = \frac{\mu_s(\mathbf{x})}{\mu_t(\mathbf{x}) + |\mu_n(\mathbf{x})|}, \quad (32)$$

$$P_n(\mathbf{x}) = \frac{|\mu_n(\mathbf{x})|}{\mu_t(\mathbf{x}) + |\mu_n(\mathbf{x})|}. \quad (33)$$

Notice that the probabilities are identical to those in standard delta tracking from Equations (24) if $\mu_n(\mathbf{x})$ is non-negative, in which case $w_s(\mathbf{x}) = 1$. As soon as one of the tentative collisions occurs in a region where $\mu_n(\mathbf{x})$ is negative, the terms in the weight no longer cancel out perfectly and the sample weight will deviate from 1.

Another practical application of weighting was proposed by

Eymet et al. [EPG*13] who combined weighted tracking and energy partitioning [SSH73], also known as sample splitting. The double-particle model of Szirmay-Kalos et al. [SKGM*17] effectively amounts to sample splitting as well.

Discussion. The above formalization of weighted delta tracking was inspired by the integral formulation of null-collision algorithms developed in particle-transport physics by Galtier et al. [GBC*13] and recently adopted in computer graphics by Kutz et al. [KHLN17]. Since the integral formulation is derived from the RTE, it directly demonstrates the correctness of null collision algorithms without relying on constraining physical interpretations or involved mathematical proofs [Col68, Mil67]. To that end, Galtier et al. first apply a number of identities to the RTE to obtain a formulation that can be trivially turned into a Monte Carlo estimator. We took a slightly different approach here of first defining an estimator and then reasoning about the probabilistic evaluation. Our approach should feel more familiar to computer-graphics audience; nevertheless, we encourage the interested reader to consult the works of Galtier et al. [GBC*13] and Eymet et al. [EPG*13] for a slightly different perspective and additional insights for handling surfaces.

4.2.3. Decomposition tracking

The performance of (weighted) delta tracking is in practice proportional to the cost of querying spatially-varying medium collision coefficients. Kutz et al. [KHLN17] propose to reduce the number of these evaluations by a technique called decomposition tracking. The idea is to decompose the volume into two components: a (piecewise) homogeneous component μ_t^c , which can be sampled cheaply, and a residual heterogeneous component $\mu_t^r(\mathbf{x}) = \mu_t(\mathbf{x}) - \mu_t^c$, which should be queried as little as possible to reduce the cost. Given two distance samples, \mathbf{x}_c representing a free path in the homogeneous component, and \mathbf{x}_r obtained using delta tracking of the residual component, a valid free path in the *composite* volume can be obtained by taking the minimum of the two sampled distances, i.e. $\mathbf{x} = \min(\mathbf{x}_c, \mathbf{x}_r)$. Figure 3 provides an illustration.

Creating two distance samples is clearly not faster than sampling the composite volume, unless—and this is the key observation of Kutz and colleagues—the sample in the residual component is constructed lazily. The authors first sample the homogeneous component and then perform delta tracking of the heterogeneous component, but as soon as a tentative collision exceeds \mathbf{x}_c , the tracking is terminated. Some free paths are thus created without ever evaluating $\mu_t^r(\mathbf{x})$ and the number of spatially varying queries is generally

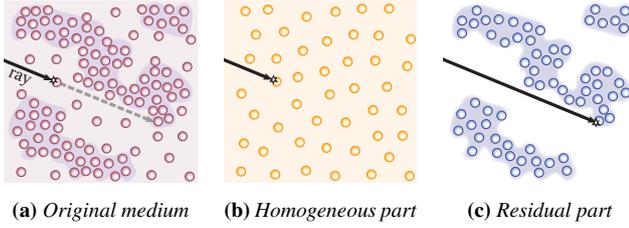


Figure 3: Decomposition tracking splits a heterogeneous medium into a homogeneous component and a residual heterogeneous component. The method then (lazily) samples free paths independently in each component and takes the shorter as the free path in the original medium. Figure reproduced from [KHLN17].

reduced. With a careful implementation, this approach can produce identical free-path samples to delta tracking without any decomposition if the same random numbers are used.

A desired property of the decomposition is to represent most of the volume using the homogeneous component—this will lead to greater savings. The goal is thus to find a tight *minorant* bound μ_t^c that permits closed-form sampling. Since delta tracking is used for the residual component, we also need to define an ideally tight majorant for the residual component. Various acceleration structures can be used to that end (see Section 8). Kutz et al. [KHLN17] also propose a weighted variant of decomposition tracking to gracefully handle situations when exact (local) minima and maxima cannot be calculated easily, such as in some procedural volumes.

4.3. Tabulation-based sampling

As already demonstrated in the sections on null-collision methods above, distances can be sampled arbitrarily as long as the samples are weighted by $T(\mathbf{x}, \mathbf{y})/p(y)$. In this section, we briefly discuss tabulation approaches designed to importance sample more terms instead of just the transmittance. Later in Section 6, when we discuss the construction of entire light transport paths, we will also examine the distance sampling PDFs used by other techniques, such as equiangular [KF12] and joint-importance sampling [GKH*13].

Since analytic sampling techniques can be employed only for a rather small class of problems, numerical recipes based on tabulation are often used in cases when the sampling CDF cannot be easily defined or inverted. In such cases, one typically builds a piecewise-polynomial PDF along the ray and finds the sampled location using a binary search over the segments, followed by solving the polynomial to find the exact location.

Kulla and Fajardo [KF12] precompute a piecewise-constant PDF proportional to the product of transmittance and scattering coefficient along the given ray using ray marching. This can be further improved by including a spatially-varying estimate of fluence [NSJ14]. Novák et al. [NNDJ12b, NNDJ12a] employed piecewise-linear PDFs to importance sample the contributions of virtual linear lights to camera rays according to the product of the two phase-functions along the lines. The proposed approach constructs the PDF on the fly, allowing only about ten evaluations of the product to keep the overhead small. This constraint was lifted

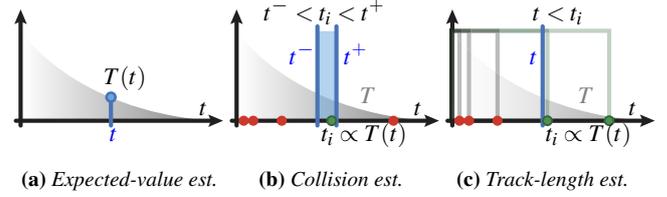


Figure 4: Expected-value estimators (a) calculate the transmittance value exactly. Collision (b) and track-length (c) estimators utilize free-path sampling techniques and check if the sample falls within an interval centered at t , or past t , respectively. Figure reproduced from [BJ17].

by Georgiev et al. [GKH*13] who re-parameterized the geometric setting to precompute a single table for all configurations of two (infinite) lines. The precomputed PDF can be sampled much more densely, providing more accurate importance sampling.

Many of the aforementioned algorithms weight distance samples to ensure unbiased results. In the next section, we detail the various approaches developed for estimating transmittance, which is needed for computing the sample weight, and for the general task of estimating transport along a given path segment.

5. Transmittance estimation

Transmittance quantifies the expected fraction of photons that travel between two points \mathbf{x} and \mathbf{y} without undergoing absorption or out-scattering. When sampling free paths, we account for transmittance implicitly through the distribution of distances. If other distance PDFs are used, or when calculating attenuation along given rays (e.g. shadow rays in a path tracer), we need to estimate transmittance explicitly. There are many different ways to estimate transmittance, and in the subsections below we broadly classify these according to how they approach the problem; Figure 4 illustrates some of these approaches.

5.1. Estimators integrating optical thickness

One broad class of techniques estimates transmittance by computing the optical-thickness integral $\tau(t) = \int_0^t \mu_t(\mathbf{z}) dz$ and exponentiate its negative to obtain the transmittance.

5.1.1. Expected-value estimator

When the extinction coefficient along the ray permits estimating optical thickness analytically, the transmittance can be calculated exactly (with zero variance and no bias). Such estimators are referred to as *expected-value* estimators [Spa66]. For homogeneous media, the expected-value estimator involves evaluating the exponential at distance t :

$$\langle T(t) \rangle_{EV} = e^{-\mu t}. \quad (34)$$

Some simple forms of heterogeneous media also accept analytic computation of optical thickness. We list a few examples in Table 1.

5.1.2. Regular tracking

Regular tracking, introduced in Section 4.1.2, can be adapted for calculating transmittance in piecewise-simple media. We sweep along the straight line from 0 to t finding all boundary crossings and for each penetrated volume i , we compute optical thickness τ_i of the overlapping ray segment. The negated total optical thickness of all k segments is then exponentiated to calculate the transmittance: $T(t) = -e^{-\sum_i^k \tau_i}$.

If the volumes/voxels are homogeneous, then we simply have $\tau_i = \mu_{t,i} \cdot \Delta_i$, where Δ_i is the distance the ray travels through voxel i and $\mu_{t,i}$ is its extinction coefficient. It is also possible to compute the transmittance in this way for higher-order voxel reconstructions such as trilinear interpolation, though the expression for τ_i becomes more complex. Since the optical thickness is computed exactly, regular tracking can be classified as an expected-value estimator.

5.1.3. Ray marching & quadrature rules

As in the case of sampling free paths, one can avoid the burden of finding all boundary crossings by using ray marching. The basic idea is to approximate the optical thickness integral by a numerical quadrature:

$$\langle T(t) \rangle_{\text{RM}} = e^{-\langle \tau(t) \rangle}. \quad (35)$$

The most straightforward approach uses a Riemann summation, which approximates τ by assuming that the extinction function is *constant* along each equal-sized step. This corresponds to using Equation (35) with:

$$\langle \tau(t) \rangle_{\text{RS}} = \sum_{i=1}^k \mu_t(t_i) \Delta_t, \quad \text{with: } t_i = (i-1 + t_0) \Delta_t, \quad (36)$$

where $\Delta_t = t/k$, and $t_0 = 0, 1/2, 1$ results in the left-, central-, and right-Riemann sum, respectively. Higher-order quadrature techniques like the trapezoid rule and Simpson's rule (which assume linear and quadratic variation of μ_t , respectively) are also possible. In addition, the step size Δ_t can be adapted to the local variation of the medium density.

Ray marching can be applied to any media with potentially spatially varying extinction; however, the assumption of constant, linear, quadratic, etc. variation of extinction leads to a systematic over-estimation of the transmittance as formally described by Jensen's inequality [Jen06]. Ray marching can therefore be viewed as a *biased* expected-value estimator.

5.1.4. Monte Carlo integration of optical thickness

An alternative to integrating the optical thickness using quadrature is to employ a MC estimator, which takes the general form

$$\langle \tau(t) \rangle_{\text{MC}} = \frac{1}{k} \sum_{i=1}^k \frac{\mu_t(t_i)}{p(t_i)}, \quad \text{with } t_i \in [0, t] \text{ and } t_i \propto p(t_i). \quad (37)$$

Note that if we use a uniform PDF and interpret $\Delta = \frac{1}{kp(t_i)}$, then this is equivalent to that of the Riemann sum above, though evaluated using stochastic samples.

Since MC integration has relatively poor convergence for such

low-dimensional integrals, Pauly et al. [PKK00] suggested a stratified MC approach which jitters an entire sequence of regularly spaced samples using a single uniform random offset $\xi \in [0, 1]$. Plugging this sequence into Equation (37) and assuming a uniform PDF yields:

$$\langle \tau(t) \rangle_{\text{SMC}} = \sum_{i=1}^k \mu_t(t_i) \Delta_t, \quad \text{with: } t_i = (i + \xi) \Delta_t. \quad (38)$$

This brings this approach even closer to the Riemann sum, but the random jitter decorrelates the error between different evaluations and replaces structured (banding) artifacts with noise. This general idea of randomly offsetting entire regular sample patterns has also more recently been used for other MC integration problems in graphics under the name "uniform jittered sampling" [RAMN12].

Even though Equations (37) and (38) are *unbiased* estimators of optical thickness, once the estimate is inserted into the transmittance exponential, the final transmittance estimate will be *biased*, overestimating the true transmittance on average. As with numerical quadrature, this can also be proven using Jensen's inequality.

5.2. Estimators using free-path sampling

It is also possible to stochastically estimate transmittance by transforming the free-path sampling techniques from the previous section into transmittance estimators.

5.2.1. Collision estimator

Collision estimators approximate the transmittance at a given distance t by relating it to the collision density at that distance. In the neutron transport literature [Spa66, SG69, LK91] these estimators are typically described as "scoring" the reciprocal of the extinction coefficient if a sample falls within some designated region. We will instead derive the collision estimator directly as a MC estimator of (potentially blurred) transmittance.

We start by noting that the convolution of any function with the delta distribution δ gives back the original function. Hence, we can write the transmittance as

$$T(t) = \int_0^\infty T(s) \delta(s-t) ds. \quad (39)$$

Given the above integral, we could imagine approximating it with a simple MC estimator that draws random distances t_i . Unfortunately, due to the delta function, such an approach would always return zero since the probability of a sample t_i falling exactly at t is zero. To avoid this problem, we can replace the delta distribution with some normalized kernel K with finite support:

$$T(t) \approx \int_0^\infty T(s) K(s-t) ds \approx \frac{T(t_i) K(t_i-t)}{p(t_i)}. \quad (40)$$

While this step introduces bias—*blurring* of transmittance by K —the one-sample MC estimator on the RHS provides a usable result.

It could seem that Equation (40) has not helped to solve our problem, since the RHS still requires evaluating the transmittance at t_i . If we sample the distances t_i using free-flight sampling with PDF

$p(t_i) = \mu_t(t_i)T(t_i)$, we get

$$\langle T(t) \rangle_C = \frac{\mathcal{T}(t_i)K(|t_i-t|)}{\mu_t(t_i)\mathcal{T}(t_i)} = \frac{K(t_i-t)}{\mu_t(t_i)}. \quad (41)$$

For the case where K is a normalized box kernel that returns $1/w$ whenever $t_i - t \in [-\frac{w}{2}, \frac{w}{2}]$ and zero otherwise, Equation (41) would return $\frac{1}{w\mu_t(t_i)}$ iff the sample t_i falls within the kernel support. In general, any normalized kernel is admissible, and as the width of the kernel approaches the delta distribution, the estimator approaches the correct expected value. Note that in neutron transport, “collision” estimators are typically used to estimate the *integral* of transmittance within a region, instead of the average transmittance. This simply corresponds to using an un-normalized kernel for K .

The collision estimator is typically used in methods such as volumetric photon mapping [Jen96, JC98], where it is often derived as a density estimator.

5.2.2. Track-length estimator

While collision estimators approximate transmittance by considering the density of photon collisions occurring at a specific location/distance, track-length estimators interpret the transmittance as the probability that a photon will reach *at least* some distance. This technique was first developed in neutron transport [Spa66, SG69, Cra78], and then independently reinvented in graphics by Jarosz et al. [JNT*11].

To obtain this estimator, let us consider the following simple modification:

$$\langle T(t) \rangle_{RR} = \begin{cases} \frac{T(t)}{P(\text{accept})} & \text{if } \text{accept} \\ 0 & \text{otherwise.} \end{cases} \quad (42)$$

The above expression evaluates the transmittances only some of the time, i.e. with probability $P(\text{accept})$, and the remaining times it just returns zero. This is an instance of Russian roulette [AK90], and it is easy to show that this estimator is *unbiased*:

$$E[\langle T(t) \rangle_{RR}] = P(\text{accept}) \cdot \frac{T(t)}{P(\text{accept})} + P(\text{reject}) \cdot 0 = T(t), \quad (43)$$

as long as $P(\text{accept}) > 0$ for all t where $T(t) > 0$.

Much as in the collision Estimator (40), the benefits of such a formulation are not immediately obvious however, since Equation (42) still requires computing $T(t)$. The trick is to design the acceptance probability $P(\text{accept})$ so that it exactly equals the transmittance $T(t)$, in which case the terms will cancel out, and we can return 1 and avoid the need to evaluate transmittance at all.

We can accomplish this by defining *accept* as the condition that we have generated a free-flight distance $t_i > t$. If the free-flight distances are distributed according to the transmittance, i.e. $p(t_i) = \mu_t(t_i)T(t_i)$, using the techniques described in the last section, the probability that a photon has traveled at least as far as t is:

$$P(t_i > t) = \int_t^\infty p(s) ds = \int_t^\infty \mu_t(s)T(s) ds = T(t). \quad (44)$$

Note that the above relationship is a generalization of Equation (20) and is true whether the medium is homogeneous or heterogeneous.

By using $P(t_i > t)$ as the acceptance probability in Equation (42), we obtain the so-called track-length estimator:

$$\langle T(t) \rangle_{TL} = \begin{cases} \frac{T(t)}{P(t_i > t)} = 1 & \text{if } t_i > t \\ 0 & \text{otherwise,} \end{cases} \quad (45)$$

which returns 1 if the sampled photon reaches t , and zero otherwise.

Bias. As shown in Equation (43), this is an unbiased estimate of transmittance if $P(t_i > t) = T(t)$, which is true for all unbiased transmittance-based free-flight sampling approaches discussed in the previous section. We could also use a biased approach like ray marching to compute free-flight distances t_i . If we do so, however, the free-flight PDF is only approximately equal, $p(t_i) \approx \mu_t(t_i)T(t_i)$, and the resulting track-length estimator becomes biased.

Discussion. In the neutron transport literature, the track-length estimator gets its name because it scores the “track-length” or distance that a photon travels within a region of interest. In our definition we return 1 instead of the length since we are estimating the transmittance at some distance, instead of the integral of the transmittance through some region. Křivánek et al. [KGH*14] were the first to establish a firm mathematical link between track-length estimators used in neutron transport, and the short photon beam estimators independently developed in graphics [JNT*11]. In fact, the track-length estimator of transmittance is used extensively in graphics beyond just photon beams: it arises whenever we perform Russian roulette to evaluate the transmittance for the reduced radiance to a surface.

5.2.3. Weighted track-length estimator

We can generalize the track-length estimator by starting with Equation (42), and defining:

$$\langle T(t) \rangle_{WTL} = \begin{cases} \frac{T(t)}{P(t_i > t)} & \text{if } t_i > t \\ 0 & \text{otherwise,} \end{cases} \quad (46)$$

where the acceptance probability in the denominator is computed by integrating some *sampling* PDF as in Equation (44), but one which does not necessarily cancel out exactly with $T(t)$.

Residual tracking. Novák et al. [NSJ14] applied this idea in the form of the so-called *residual-tracking* estimator. The technique utilizes a decomposition of the volume into two parts, much like the aforementioned decomposition tracking [KHLN17], and bears similarities with the “separation of the main part” used by Szirmay-Kalos et al. [SKTM11] for the estimation of optical thickness.

The goal of residual tracking is to account for a portion of the medium analytically, and only use the track-length estimator for the remaining portion of the medium’s extinction function. We can accomplish this by decomposing the heterogeneous extinction function $\mu_t(t)$ into a *control* $\mu_t^c(t)$ and a *residual* $\mu_t^r(t)$ extinction function, such that $\mu_t(t) = \mu_t^c(t) + \mu_t^r(t)$ and hence $T(t) = T^c(t)T^r(t)$.

The residual-tracking estimator samples distances proportional to the residual transmittance function, i.e. $p^r(t_i) = \mu_t^r(t_i)T^r(t_i)$, and evaluates Equation (46) using this sampling PDF to compute the

probability in the denominator:

$$P(t_i > t) = \int_t^\infty p^r(s) ds = T^r(t). \quad (47)$$

Hence, we have:

$$\langle T(t) \rangle_{\text{Res}} = \langle T^c(t) \rangle_{\text{EV}} \langle T^r(t) \rangle_{\text{TL}} = T^c(t) \begin{cases} \frac{T^r(t)}{P(t_i > t)} = 1 & \text{if } t_i > t \\ 0 & \text{otherwise.} \end{cases} \quad (48)$$

Note that this estimator simply returns the control transmittance—which we assume we can compute using an expected-value estimator—whenever the sample reaches t . Residual tracking can be seen as a hybrid between an expected-value and a track-length estimator. If the control component matches the actual extinction function exactly, the residual extinction is zero and Equation (48) simplifies to the expected-value estimator. Conversely, if the control component is zero and the residual matches the actual extinction exactly, then Equation (48) simplifies to the track-length estimator (45).

5.3. Estimators utilizing null collisions

Null-collision based free-flight routines, such as delta-tracking, can be used to estimate transmittance using the collision, track-length, and weighted track-length estimators described in Section 5.2. However, the concept of adding null collisions opens up the possibility for other classes of transmittance estimators, which we describe below.

5.3.1. Ratio tracking

When delta tracking is used to form a track-length transmittance estimator, $i - 1 \geq 0$ null collisions are generated in the medium prior to obtaining the first real collision at t_i ; here i indexes the tentative collisions. The algorithm probabilistically chooses whether each tentative collision is real or null by drawing a random number and terminating the procedure if a real collision is chosen (with probability $P_{a+s}(t_j) = \frac{\mu_n(t_j)}{\bar{\mu}(t_j)}$ for $j = 1, \dots, i - 1$). The transmittance estimate thus starts at 1, where it remains during all null collisions t_j , but immediately drops to 0 upon the first real collision at t_i . This probabilistic termination can be seen as a form of Russian roulette and introduces variance.

The application of delta tracking as a track-length transmittance estimator was described by Cramer [Cra78], who proposed also a weighted version for reducing variance of deep penetration estimates in nuclear physics. The weighted version was independently developed in the field of computer graphics by Novák et al. [NSJ14] under the name *ratio tracking*; we use this name to distinguish the method from other weighted approaches.

The goal of the ratio-tracking estimator is—as in the track-length estimator—to estimate the expected percentage of photons that make it beyond distance t . The main idea is to remove the probabilistic termination and instead weight the samples by the probability of continuing the walk; this concept is known as “implicit capture” [AK90] or “Rao-blackwellization” [SKGM*17].

In contrast to a track-length estimator, which returns a binary result depending on whether the real collision occurs at $t_i < t$, ratio

tracking allows *all* distance samples to reach t scoring a fractional weight. The tracking *never* terminates before t and weights the transmittance estimate at each tentative collision t_j by the “probability” of continuing forward, $\mu_n(t_j)/\bar{\mu}(t_j) = 1 - \mu(t_j)/\bar{\mu}(t_j)$, thereby maintaining the correct expected value. We used the term probability here to emphasize the relation to delta tracking, but the term is simply a weight that can become negative when $\bar{\mu}(t)$ does not bound $\mu_t(t)$.

The estimator can also be derived directly from the weighted delta tracker, see Equation (25), by enabling only null collisions before reaching t , i.e. $P_a(s) = P_s(s) = 0$, $P_n(s) = 1$ for $s < t$, and forcing the walk to terminate as soon as t has been reached. The weight Equation (30) then forms the following estimator:

$$\langle T(t) \rangle_{\text{RT}} = \prod_{j=1}^{i-1} \frac{\mu_n(t_j)}{\bar{\mu}(t_j)}. \quad (49)$$

5.3.2. Residual ratio tracking

Novák and colleagues [NSJ14] proposed to combine ratio tracking with the concept of residual estimation. The residual-ratio-tracking estimator extracts a control extinction function $\mu_t^c(t)$ and computes the control transmittance $T^c(t)$ in closed form using an expected-value estimator. The residual transmittance $T^r(t)$ is approximated by ratio tracking. The final transmittance estimator is:

$$\langle T(t) \rangle_{\text{ResRT}} = T^c(t) \prod_{j=1}^{i-1} \frac{\mu_n^r(t_j)}{\bar{\mu}^r(t_j)}. \quad (50)$$

5.3.3. Next-flight estimators

As concluded by Cramer [Cra78] and Novák et al. [NSJ14], both track-length and weighted estimators have certain drawbacks. While the first approach suffers from higher variance due to binary scores, the second can be computationally inefficient if many steps are needed to reach t , during which the weight becomes insignificant. To that end, Cramer [Cra78] proposed to equip track-length and ratio-tracking estimators with a form of next-flight estimation (a.k.a. next-event estimation) drawn from the starting point and from each tentative collision to the (end) point at distance t .

Each next-flight connection represents the *uncollided* transport through the real + fictitious medium along the remainder of the ray. We start with the uncollided contribution through the entire distance $T_{\bar{\mu}}(0, t)$. Due to the presence of null collisions, this clearly underestimates the true transmittance. We thus track forward and analogously *add* another (weighted) uncollided contribution at each tentative collision. Cramer proposed and proved correctness of two variants of such estimation.

The first variant utilizes a probabilistic termination (analogous to delta tracking) and scores a weighted free-flight estimate at each tentative collision up to the first real collision (inclusive) or up to distance t . Such *next-flight-delta-tracking* estimator reads:

$$\langle T(t) \rangle_{\text{NFDT}} = T_{\bar{\mu}}(0, t) + \sum_{j=1}^n \frac{\mu_n(t_j)}{\bar{\mu}(t_j)} T_{\bar{\mu}}(t_j, t) \quad (51)$$

where n is the index of the first collision classified as real, or the last tentative collision before reaching t , whichever occurs first.

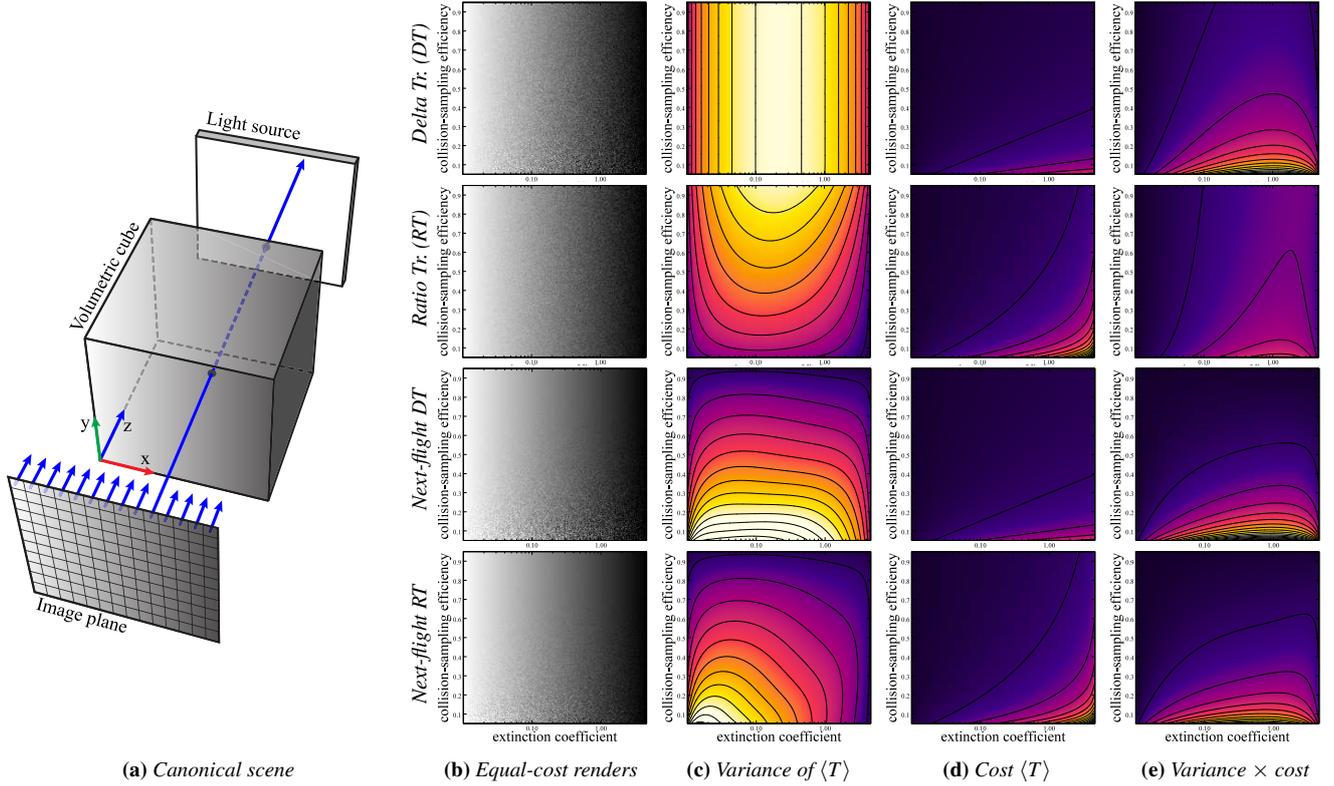


Figure 5: An analysis of four different transmittance estimators, each using roughly the same number of μ_t evaluations for rendering column (b). The test scene (a) consists of a volume with varying optical thickness (horizontal axis) such that the transmittance increases linearly. To study the impact of null collisions, we vary the amount of fictitious matter vertically such that the relative amount of real matter, which defines the collision-sampling efficiency, changes linearly from 0.05 (bottom) to 0.95 (top). False-coloring in each column uses the same scale.

The second variant relies on ratio tracking. The algorithm proceeds all the way until t with all next-flight estimates being weighted by the product of local null-collision weights encountered up to and including that point. The *next-flight-ratio-tracking* estimator reads:

$$\langle T(t) \rangle_{\text{NFRT}} = T_{\bar{\mu}}(0, t) + \sum_{j=1}^m \prod_{k=1}^j \frac{\mu_n(t_k)}{\bar{\mu}(t_k)} T_{\bar{\mu}}(t_j, t) \quad (52)$$

where m is the index of the last tentative collision before reaching t .

While it may be difficult to justify these estimators using intuition, Kutz et al. [KHLN17] demonstrate that the next-flight-ratio-tracking estimator can be derived directly from the RTE using the integral framework of Galtier et al. [GBC*13].

In Figure 5, we compare the four aforementioned estimators in a canonical setting [NSJ14] that allows identifying configurations that lead to good or poor performance with each estimator. The scene consists of an axis-aligned unit cube filled with an absorbing medium. The extinction coefficient increases exponentially along the x axis so that the transmittance (and thus the brightness of pixels) decreases linearly as we move from left to right across the rendered image. The density of the medium is kept constant vertically and also along the camera rays. We then add fictitious matter, and increase its density vertically from top to bottom to study the cost and variance of individual methods. The analysis suggests to use

ratio tracking for media with large proportions of fictitious particles, e.g. around boundaries of the volume. Next-flight estimators improve the efficiency in the opposite case.

The next-flight estimators can be further improved by applying the decomposition concept and performing e.g. “residual next-flight ratio tracking”; this application remains to be investigated.

6. Sampling light transport paths

In this section, we discuss unidirectional path tracing that utilizes distance samplers from Section 4 for building individual segments. We also describe methods for performing various forms of next-event estimation, namely standard shadow rays (Section 6.2), two-segment shadow rays (Section 6.3), and three-segment shadow rays (Section 6.4), each of which draws distances from a different distribution and leverages one of the transmittance estimators described in Section 5. Lastly, we discuss approaches for handling spectrally varying collision coefficients (Section 6.5).

6.1. Unidirectional sampling

The best-known and most commonly used path-sampling technique is the one that starts from the camera or a light source and extends the path incrementally segment by segment. To determine the location of the next path vertex \mathbf{x}_{i+1} , a ray direction is first sam-

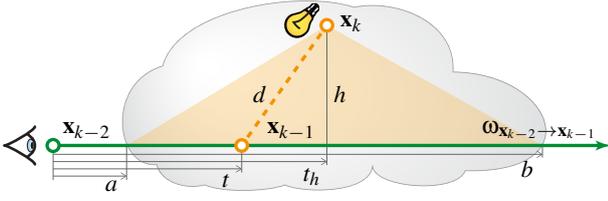


Figure 6: Equiangular sampling places points proportional to $1/d^2$, here limited to segment (a, b) .

pled with a distribution $p(\omega_{i+1})$, typically given by the medium phase function or by the surface BSDF at the location of that vertex. Then a propagation distance along the ray is sampled with $p(t_{i+1})$, e.g. using one of the sampling methods described in Section 4. The PDF of a path vertex is then conditioned on the previous path vertex and direction: $p(\mathbf{x}_{i+1} | \mathbf{x}_i, \omega_i) = p(\omega_{i+1} | \mathbf{x}_i, \omega_i) p(t_{i+1} | \mathbf{x}_i, \omega_{i+1}) G(\mathbf{x}_{i+1}, \mathbf{x}_i)$, where the geometry term is needed to convert from solid angle \times distance measure to volume or surface area measure. The PDF of the entire path $p(\bar{\mathbf{x}})$ is obtained by multiplying the sequence of conditional PDFs from the start of the path, with indices reversed if starting at the light source.

When computing a path-integral estimate of the form in Equation (14), the ratio $f_i(\bar{\mathbf{x}})/p(\bar{\mathbf{x}})$ is typically computed incrementally as consecutive vertices are sampled. This strategy is well-suited for the stochastic tracking-based distance sampling methods from Section 4.2, whose PDFs generally cannot be evaluated explicitly but cancel out (up to a weighting factor) with the transmittance term along the segment in the path contribution.

6.2. Luminaire sampling

A purely unidirectional sampling process can yield a non-zero contribution only when a light source (respectively the camera) is randomly hit. To increase the efficiency, the last vertex on the path can be “placed” on the a light source (camera) explicitly by sampling its area/volume. This technique is called *next-event estimation* [CCY67] and requires choosing a point on the light source (camera), casting a shadow ray, and computing the fractional visibility, e.g. using one of the methods described in Section 5.

In cases of light sources with finite extent, the sampling density of the point to which we cast the shadow ray can greatly impact the efficiency of the estimator. The naive approach of uniformly sampling the emissive area can significantly increase estimation variance. Several techniques have thus been developed for uniformly sampling solid angles of e.g. triangular [Arv95], rectangular [UFK13], disk and cylindrical [Gam16], elliptical [GUK*17], or polygonal [Arv01] luminaires.

While such importance sampling can greatly reduce variance, it may still remain high, especially with small light sources. In fact, for point lights the variance is unbounded as individual contributions tend to infinity when the last free vertex in the medium approaches (by random chance) the point light—the geometry term along the shadow ray tends to infinity in such cases. One way to ameliorate the high variance is to sample at least part of the path by taking into account the (sampled) position on the light source; we describe two such sampling recipes in the next two sections.

6.3. Equiangular sampling

Equiangular sampling is a technique for reducing variance of estimating directly in-scattered illumination along a ray. The technique chooses \mathbf{x}_k on the light first and then samples the second-last vertex \mathbf{x}_{k-1} along the ray with direction $\omega_{\mathbf{x}_{k-2} \rightarrow \mathbf{x}_{k-1}}$ with a PDF proportional to the inverse squared distance $d = \|\mathbf{x}_{k-1} - \mathbf{x}_k\|$, as illustrated in Figure 6. The first step is to parameterize distance d as a function of the ray parameter t :

$$d^2(t) = h^2 + (t - t_h)^2, \quad (53)$$

The PDF is then

$$p(t) = \frac{h}{(h^2 + (t - t_h)^2) \left(\tan^{-1} \left(\frac{b - t_h}{h} \right) - \tan^{-1} \left(\frac{a - t_h}{h} \right) \right)} \quad (54)$$

whose inverse CDF is

$$t(\xi) = t_h + h \tan \left((1 - \xi) \tan^{-1} \left(\frac{a - t_h}{h} \right) + \xi \tan^{-1} \left(\frac{b - t_h}{h} \right) \right), \quad (55)$$

allowing us to sample points along the ray proportionally to $1/d^2$ with $t \in [a, b]$.

This technique was originally developed in neutron transport under the name *once-more-collided-flux* estimator [KC77] and. It was later independently published in computer graphics by Kulla and Fajardo [KF12] who called it *equiangular* sampling because the angular distribution of directions from \mathbf{x}_k toward the sampled points is uniform.

6.4. Joint importance sampling

The generic pixel estimator in Equation (14) suggests that light transport paths should ideally be sampled from a joint vertex distribution that is proportional to the their measurement contribution. The practical difficulty in sampling from a prescribed joint distribution is that deriving the (conditional) distributions for the individual sampling decisions requires the successive marginalization of the corresponding variables one by one out of the joint PDF. This marginalization involves integration, which often cannot be performed analytically.

Most existing path-sampling techniques, including unidirectional sampling (Section 6.1), generate path vertices by instead prescribing the distributions for the individual sampling decisions. While these typically importance sample the local scattering and propagation events, there is no explicit control over the form of the final joint path distribution with this approach. The result is increased variance in the pixel estimator, especially in anisotropically scattering media and with light sources placed inside volumes.

For the special case of double scattering in media, Georgiev et al. [GKH*13] have found a way to sample from a prescribed joint distribution proportional to the product of geometry and phase function terms along a sequence of four path vertices—three segments. Given a point \mathbf{a} with an (optional) incident direction $\omega_{\mathbf{Ia}}$ as well as a ray $(\mathbf{d}, \omega_{\mathbf{dc}})$, they derive the conditional PDFs and corresponding techniques for sampling two intermediate connection vertices \mathbf{b} and \mathbf{c} . Figure 7 illustrates one factorization of the prescribed

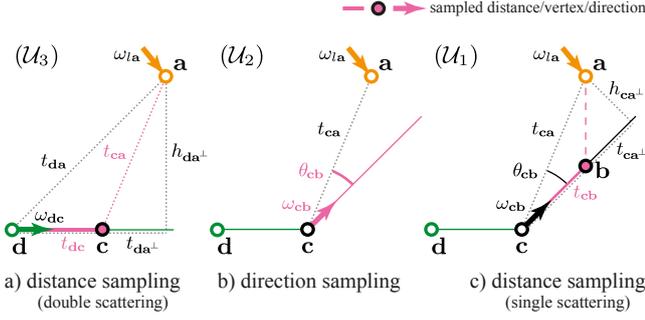


Figure 7: Given a point with an incident direction, $(\mathbf{a}, \omega_{\mathbf{a}})$, and a ray $(\mathbf{d}, \omega_{\mathbf{d}})$, joint importance sampling constructs a 2-vertex connection with a prescribed joint distribution which can be factorized into a series of distance, direction and distance sampling decisions. Figure reproduced from [GKH*13].

joint into a series of three sampling decisions, starting with a distance sampling along the ray. For the case of isotropically scattering media, where the medium phase function is constant, closed-form expressions for the sampling PDFs exist:

$$p(t_{\mathbf{cb}}) = \frac{t_{\mathbf{ca}} \sin \theta_{\mathbf{cb}}}{\pi - \theta_{\mathbf{cb}}} \frac{1}{h_{\mathbf{ca}\perp}^2 + (t_{\mathbf{ca}\perp} - t_{\mathbf{cb}})^2}, \quad (56)$$

$$p(\omega_{\mathbf{cb}} | t_{\mathbf{dc}}) = \frac{\pi - \theta_{\mathbf{cb}}}{\pi^3 \sin \theta_{\mathbf{cb}}}, \quad (57)$$

$$p(t_{\mathbf{dc}} | \omega_{\mathbf{cb}}, t_{\mathbf{dc}}) = \frac{C_{\mathbf{dc}}}{\sqrt{h_{\mathbf{da}\perp}^2 + (t_{\mathbf{da}\perp} - t_{\mathbf{dc}})^2}}, \quad (58)$$

where $C_{\mathbf{dc}}$ is a normalization constant. We list the PDFs in the order they are derived via successive marginalization, which is the opposite of the sampling order. Note that the PDF of $t_{\mathbf{cb}}$ is equivalent to the equiangular PDF from Section 6.3 with upper bound set to infinity. Also note that the PDF of $t_{\mathbf{dc}}$ is proportional to the *inverse distance* to \mathbf{a} , whereas the equiangular PDF of $t_{\mathbf{cb}}$ is proportional to the *inverse squared distance* to \mathbf{a} . Sampling from these PDFs is done using the inverses of their corresponding CDFs:

$$t_{\mathbf{cb}} = t_{\mathbf{ca}\perp} + h_{\mathbf{ca}\perp} \tan \left(\xi(\pi - \theta_{\mathbf{cb}}) + \theta_{\mathbf{cb}} - \frac{\pi}{2} \right), \quad (59)$$

$$\theta_{\mathbf{cb}} = \pi(1 - \sqrt{\xi_1}), \quad \phi_{\mathbf{cb}} = 2\pi\xi_2, \quad (60)$$

$$t_{\mathbf{dc}} = t_{\mathbf{da}} \sinh(\xi C_{\mathbf{dc}}) + t_{\mathbf{da}\perp} (1 - \cosh(\xi C_{\mathbf{dc}})). \quad (61)$$

For the case of anisotropically scattering media, with a non-constant phase function, no closed-form expressions for the sampling PDFs exist. To handle this case, Georgiev et al. [GKH*13] derive compact tabulations of these PDFs by exploiting various symmetries in their individual geometric configurations. Figure 8 demonstrates the benefit of using this technique for light-source connections in unidirectional path tracing.

6.5. Spectral tracking

So far, we ignored any wavelength dependency of all quantities. However, in the case of chromatic media, i.e. when the extinction varies with wavelength, there exist a few interesting variants of importance sampling. The straightforward way to handle such situations is to just trace a single wavelength (or one channel of a color

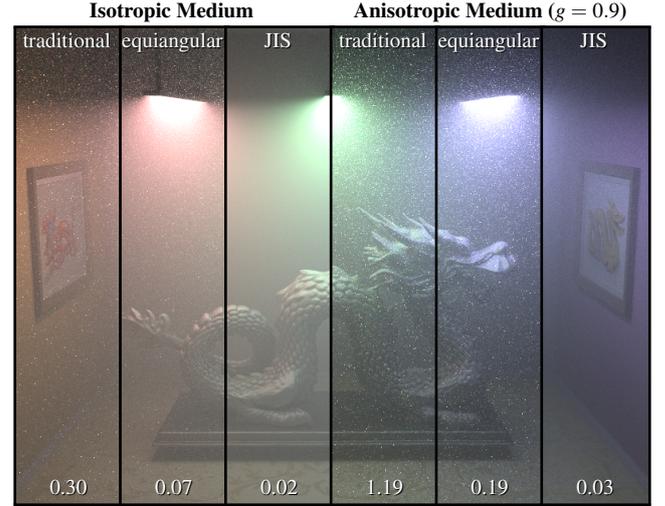


Figure 8: Single and double scattering in isotropic (left, 15 min.) and anisotropic (right, 30 min.) media rendered with unidirectional path tracing. Compared are light-source connections done via traditional transmittance-based distance sampling, equiangular distance sampling, and 3-random-decision joint importance sampling (JIS). RMS error is shown at the bottom. Figure reproduced from [GKH*13].

space in RGB rendering). Doing so, however, introduces chromatic noise and is wasteful because there is usually a large amount of correlation of the path contribution between different wavelengths. We could potentially save cost by calculating more wavelength contributions for every generated path.

6.5.1. Hero wavelength sampling

The hero wavelength sampling scheme [WND*14] (also see Radziszewski et al. [RBA09]) constructs a path with a single randomly sampled wavelength λ . To add color information, the spectral measurement contribution function $f(\bar{\mathbf{x}}, \lambda)$ of the resulting path is also evaluated for a stratified set of other wavelengths λ_i . The initial wavelength is called the hero wavelength because it determines the importance sampling. To arrive at an unbiased estimator for the pixel color I_j , the contribution of the path is accumulated for all wavelengths λ_i and weighted by the MIS weight using the balance heuristic:

$$\langle I_j \rangle = \frac{f(\bar{\mathbf{x}}, \lambda_i)}{p(\bar{\mathbf{x}}|\lambda_i)} \cdot \frac{p(\bar{\mathbf{x}}|\lambda_i)}{\sum_l p(\bar{\mathbf{x}}|\lambda_l)} = \frac{f(\bar{\mathbf{x}}, \lambda_i)}{\sum_l p(\bar{\mathbf{x}}|\lambda_l)}. \quad (62)$$

In this formula, we assume that $f(\bar{\mathbf{x}}, \lambda)$ already converts the spectral quantity to a linear tristimulus value ready to be accumulated in the frame buffer. This scheme is very general and not specialized for participating media in any way. The strength of this technique is that it will transparently perform optimal importance sampling for one wavelength and combine the estimators for the other wavelengths in a weighted way to reduce variance. Since it uses the MIS framework, it can be readily combined with a multitude of other estimators, as long as the PDF of hypothetically sampling $\bar{\mathbf{x}}$ with a different wavelength λ can be computed. The need to

compute these PDFs makes hero wavelength sampling relatively heavy weight (although SIMD instructions ameliorate some of this cost). In the context of participating media, this requirement is also the biggest limitation as null-collision algorithms cannot provide a noise-free estimate of the PDF, making the approach work best in conjunction with regular tracking.

6.5.2. Spectral tracking with null-collision methods

There are two main challenges of using standard delta tracking with spectrally-resolved extinction functions. First, we cannot use the hero-wavelength scheme since the free-path PDF cannot be evaluated. Second, (unweighted) delta tracking can simulate only one analog process at a time forcing us to estimate transport in each wavelength independently [EPG*13]. The cost of rendering chromatic media thus grows linearly with the number of wavelengths.

Weighted tracking, on the other hand, provides means to re-weight samples in case their distribution does not correspond to the analog process. As such, we can sample tentative collisions according to one wavelength while correctly estimating the transport in others via maintaining spectrally-resolved collision weights.

The main challenge of applying weighted tracking to spectrally-resolved media is balancing the cost and the variance. One could employ the hero-wavelength idea of choosing one wavelength to sample tentative collisions, but as soon as the extinction values in other wavelengths exceed the hero wavelength, the weights may grow uncontrollably and MIS cannot be used here to prevent that. Kutz et al. [KHLN17] thus propose to sample tentative collisions using $\bar{\mu}$ that ideally bounds *all* wavelengths. The authors also discuss several means for setting the collision probabilities; these are also shared across all wavelengths and directly impact the variance. Among the most practical are two schemes that incorporate the throughput (history) of the path and set the probabilities as:

$$P_{\star}(\mathbf{x}_j) = \text{reduce}(|\bar{w}(\bar{\mathbf{x}}, \lambda)\mu_{\star}(\mathbf{x}, \lambda)|)c \quad (63)$$

where \star represents absorption, scattering, or null-collision quantities, $\bar{w}(\bar{\mathbf{x}}, \lambda)$ is the product of collision weights up to vertex \mathbf{x}_j , and c is a normalization constant ensuring P_a , P_s , and P_n sum up to 1. The $\text{reduce}(\cdot)$ is to be substituted with $\text{max}(\cdot)$ or $\text{avg}(\cdot)$ to reduce the spectrally-resolved vectors to a single value. Taking the max produces images with less overall variance but leads to occasional fireflies. Using the avg prevents fireflies by ensuring that the path throughput $\bar{w}(\bar{\mathbf{x}}, \lambda)$ does not exceed a fixed value (the number of wavelengths) for any λ ; Kutz et al. [KHLN17] provide a proof in the supplementary material. The authors also show how to combine spectral and decomposition tracking.

7. Advanced methods

In this section, we review advanced methods, namely bidirectional path tracing (BDPT), many-light algorithms, density-estimation techniques and their unification with BDPT, radiance caching, most-probable-path approaches, Markov-chain techniques, zero-variance and stratification schemes, and we conclude by discussing sampling of emissive volumes.

7.1. Bidirectional path tracing

Bidirectional path tracing starts random walks to form transport paths both at the sensor and the light sources. In the presence of participating media, this proceeds just the same as in the vacuum case, only that propagation distances need to be sampled to simulate interactions with the medium and transmittance must be evaluated along deterministic connections.

Note that evaluating transmittance is symmetric as it does not depend on the tracing direction (from sensor or from light). The free-flight PDF derived from transmittance, however, is not symmetric: the normalization depends on the extinction coefficient at the target point $\mu_t(t)$. Thus, some care has to be taken when computing the probability density of complete paths for multiple importance sampling (MIS).

In the same context, MIS requires explicitly computing PDFs of individual samples and the hypothetical PDF of computing the same path by any other active technique. The value of the PDF $p(t) = \mu_t(t) \cdot T(\mathbf{x}, \mathbf{y})$ thus needs to be known and cannot be stochastically estimated, be it with an unbiased transmittance estimator. This is because the MIS weight—the PDF of the current sample divided by the sum of the PDFs of all possible techniques—contains division, which is not (in contrast to expectation) a linear operator.

It is possible to directly estimate the reciprocal of an integral instead [Boo07], which however is prone to high variance. Since MIS only needs weights that sum up to one to be correct, it is possible to instead use crude approximations of the free-flight PDF. This may however deteriorate the variance of the combined estimator, so one may prefer to use deterministic transmittance estimators in scenarios where an algorithm depends heavily on MIS.

7.2. Many-light methods

Many-light algorithms are a subset of bidirectional path tracing, where light subpaths are typically reused across multiple, rather short camera subpaths. The first instance of many-light rendering—instant radiosity [Kel97]—was developed to allow fast indirect-illumination computation on graphics hardware. The method and its derivatives precompute a shadow volume or a shadow map for each light-path vertex—referred to as a virtual point light (VPL)—utilizing it to quickly calculate visibility queries between the VPL and points seen by the camera. In a nutshell, many-light rendering reinterprets the problem of estimating global illumination as computing direct illumination (or single-scattering) from many virtual lights, which are generated via light tracing prior to actual rendering. Many extensions to the original algorithm were developed in the last two decades, most of which were surveyed by Dachsbacher et al. [DKH*13]. Here we focus only on those that address rendering in the presence of participating media.

7.2.1. Virtual point lights

One of the first applications of many-light rendering to scenes with volumes appeared in multidimensional lightcuts by Walter et al. [WABG06]. Raab et al. [RSK08] and Engelhardt et al. [ENSD12] then discuss, analyze, and propose solutions addressing the main issue of VPL methods: the splotty artifacts due to a

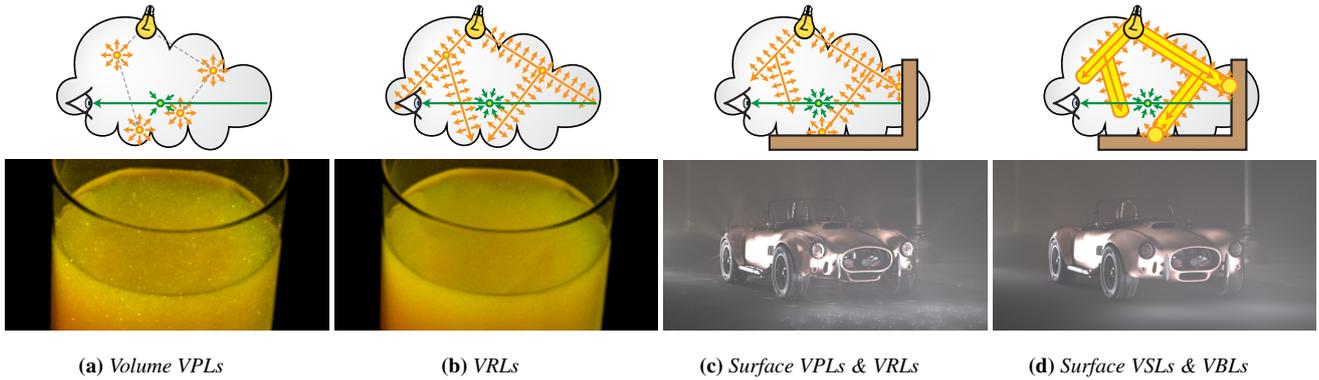


Figure 9: The two left columns demonstrate the conceptual and visual differences between virtual point lights (VPLs) and virtual ray lights (VRLs), used here to render multiple scattering inside orange juice. The right two columns compare complete solutions for rendering scenes with volumes and surfaces, one based on VPLs and VRLs and another based on virtual sphere lights (VSLs) and virtual beam lights (VBLs); only indirect illumination is shown to emphasize differences. Images reproduced from [NNDJ12b, NNDJ12a].

$1/d^2$ singularity that models the quadratic falloff of radiance emitted from a point light. While the impact of the singularity can be reduced by combining multiple path-construction strategies using MIS, as in BDPT, or avoided for shadow paths with multiple segments [GKH*13], illuminating a set of camera rays by a number of VPLs will always suffer from this problem manifesting itself as bright splotches in the image.

Raab et al. [RSK08] proposed to clamp the singularity and recover the lost transport via a path-tracing-based bias compensation [KK06] adapted to participating media. For volumes with rather isotropic profiles, Engelhardt et al. [ENSD12] observe that most energy lost due to clamping can be recovered by simulating two bounces of residual transport. The authors describe an approximate solution that is amenable to GPU acceleration yielding interactive frame rates.

Since many-light methods require repeatedly evaluating visibility to a finite collection of lights, shadow mapping techniques [Wil78] can provide substantial performance gains and hardware acceleration. In the context of participating media, several variants of deep shadow mapping [LV00] can be used to evaluate the fractional visibility. Salvi et al. [SVLL10] propose to adapt the number and placement of transmittance samples in each texel of the shadow map. Fourier opacity maps [JB10], transmittance function maps [DGMF11], and boundary-aware extinction mapping [GDML13] express the transmittance in the Fourier domain storing Fourier coefficients instead of the transmittance samples.

7.2.2. Virtual ray and beam lights

One approach to reduce the degree of the singularity, and thereby suppress the splotchy artifacts, is to distribute the energy *continuously* along the light path. Novák et al. [NNDJ12b] propose to turn each ray of the light path into a linear light—the so-called virtual ray light (VRL). The emission at any point along the VRL is defined by the radiance reaching the point along the light path, which is scaled by the scattering coefficient and directionally modulated by the phase function. Computing the in-scattered radiance along a camera ray involves numerically estimating a line-to-line integration problem. The authors propose to importance sample the double

integral using a PDF that is tabulated on the fly from the product of phase functions and the $1/d^2$ term. The integration can be further optimized by utilizing fully precomputed PDFs [GKH*13].

While the singularity is reduced, it is not removed completely. For surface rendering, Hašan et al. [HKWB09] proposed to “inflate” each virtual point light into a virtual sphere light (VSL), spreading its power over a disc that approximates the surfaces within the sphere. Novák et al. [NNDJ12a] proposed to apply this concept to ray lights, turning each path segment into a virtual beam light (VBL) with finite thickness. This avoids the singularity at the cost of introducing additional blurring. The transport is estimated by first sampling a location on the axis of the beam and then integrating over the surrounding spherical volume. While spreading the energy over a beam avoids the singularity, it may lead to overblurring of fine visual features. To that end, the authors propose to progressively shrink the beam radius as in progressive volumetric photon mapping [KZ11] to ensure the bias diminishes in the limit. The main concept and characteristic artifacts of a number of many-light primitives are illustrated in Figure 9.

7.2.3. Clustering of virtual lights

In order to address scalability, Walter et al. [WFA*05, WABG06, WKB12] proposed to organize VPLs into a tree hierarchy and find a suitable cut for each shading location. This reduces the number of evaluated contributions while still accounting for the energy of all VPLs. Recently, Yuksel and Yuksel [YY17] described a method where contributions of point lights are splatted onto a hierarchical grid yielding a scalable algorithm with better temporal coherence.

Frederickx et al. [FBD15] proposed to adapt the light-slice [OP11] method for clustering virtual ray lights. Another approach for clustering VRLs was described by Huo et al. [HWH*16] who cast the problem of estimating contributions as a matrix sampling and reconstruction task. Their adaptive scheme can efficiently reconstruct the matrix from a rather small number of sampled elements. All the aforementioned clustering methods significantly improve the scalability of many-light rendering.

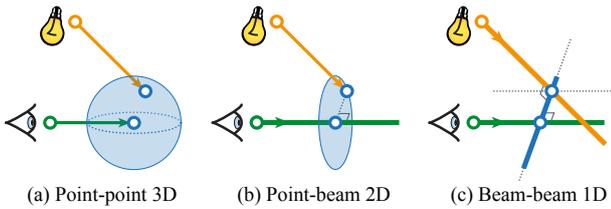


Figure 10: Different volumetric photon density estimators and their corresponding kernel (i.e. blur) dimensionalities.

7.3. Photon density estimation methods

Photon density estimation is another popular two-pass bidirectional rendering approach, where—just like in many-light methods—the cached light subpaths are reused across multiple camera subpaths. In contrast to many-light methods though, the energy carried by the vertices of those light subpaths (the so-called *photons*) is used to estimate the density of inscattered flux (i.e. the outgoing radiance) around points seen from the camera. The basic approach was first proposed by Arvo [Arv86], but was made practical for rendering of complex surfaces by Jensen [Jen96] who proposed storing all photons in a search-efficient kd-tree data structure and coined the term *photon mapping*. Hachisuka et al. [HOJ08] showed a photon-mapping formulation that allows for reducing the bias by progressively accumulating more photons in the radiance estimate and shrinking the reconstruction kernel.

This approach has proved very useful in practice thanks to a key property of density estimation: it regularizes, i.e. blurs, the light path space around the estimation location. While the blurring introduces bias, it also makes it easier to capture lighting effects that would otherwise be difficult or even impossible to sample using unbiased techniques such as BDPT. Prime examples of such effects are caustics seen directly or through reflection or refraction.

7.3.1. Photon points

Jensen and Christensen first demonstrated the use of photon maps [JC98] to render volumetric light-transport effects. Their algorithm steps along every camera ray in a medium and collects the photons around each ray-marching location. This “point-point” query is illustrated in Figure 10a. It is particularly well suited for rendering volumetric caustics, but its main disadvantage is that it is difficult to find a good ray-marching step size. A too large step size yields a noisy result, and a too small step size induces many photon queries along the ray, which can slow down rendering significantly.

To address the deficiency of ray marching, Jarosz et al. [JZJ08a] proposed a volumetric radiance estimate that gathers photons along the entire camera ray and computes their contributions without point sampling. This so-called *beam radiance estimate* can render scenes with participating media with significantly less noise than conventional photon mapping. The “point-beam” query is illustrated in Figure 10b. This work revealed that the approach of using point samples in participating media is an unnecessary legacy from surface rendering.

Jakob et al. [JRJ11] investigated the use of *parametric density estimation* to represent and reconstruct the spatial distribution of radiance in a medium. They still shoot photons from the light sources

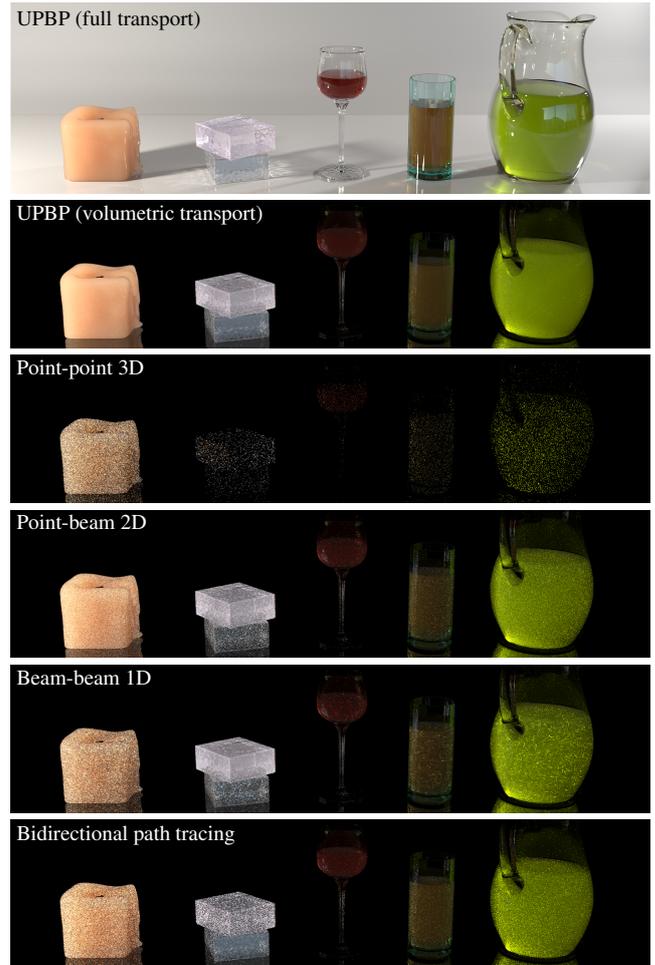


Figure 11: Equal-time comparison of UPBP against BDPT and the individual density estimators illustrated in Figure 10, showing that the combined algorithm can be significantly more efficient than each of its components alone. Images reproduced from [KGH* 14].

but only use them to fit the parameters of a hierarchical Gaussian mixture model and then immediately discard them. The Gaussian model allows for anisotropic reconstruction and a level-of-detail representation that can be about three orders of magnitude more compact than a photon map of the same quality. The Gaussian parameters can be fit progressively as photons are shot and discarded in a way that maintain temporal coherence in animations.

7.3.2. Photon beams

Recognizing the potential of beam gathering, Jarosz et al. [JNSJ11] devised a generalized theory of volumetric density estimation using beams along camera rays, light rays, or both. Sun et al. [SZLG10] concurrently proposed a specialized method for rendering caustics and single scattering that corresponds to one such estimator. Photon beams supersede points as fundamental entities in volumetric light transport and provide a more efficient intermediate representation of lighting. They can significantly improve quality over standard volumetric photon mapping and are also amenable to artistic ma-

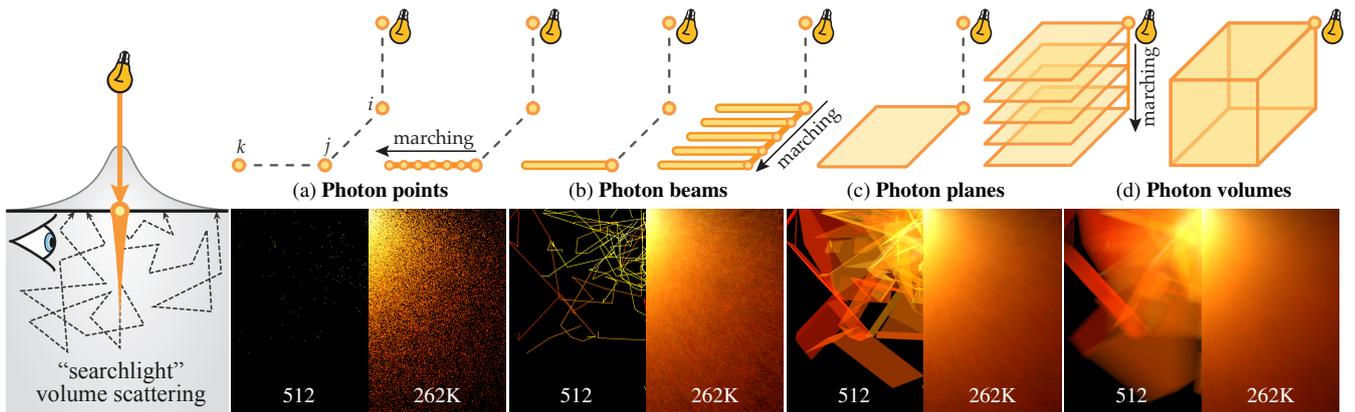


Figure 12: A comparison between photon points, beams, planes, and volumes. Bitterli and Jarosz [BJ17] form the higher-dimensional estimators by computing the limit process of “marching” photons along preceding light path segments, which allows to progressively reduce variance and bias. The canonical renderings in the bottom row use these successive estimators (with two sample counts, split vertically) on a searchlight problem setup, confirming that higher-order samples can significantly improve quality. Figure reproduced from [BJ17].

nipulation [NJS*11]. Photon beams reduce both the variance and the bias (blur), enabling the rendering of extremely sharp details such as volumetric caustics using much less photon data than previous point-based methods. The least amount of blur is achieved by a “beam-beam” estimator that uses a 1D reconstruction kernel [JNSJ11], as illustrated in Figure 10c.

To make photon beams converge to the correct solution with a bounded memory footprint, Jarosz et al. [JNT*11] devised a progressive formulation akin to progressive photon mapping for surfaces [HOJ08, KZ11]. They also proposed an efficient way to estimate the transmittance along beams in heterogeneous media at multiple locations, each corresponding to a different radiance estimation query. Along each beam they store several free-flight distances sampled using delta tracking (see Section 4.2.1). The fraction of distances that are beyond a given location provides an unbiased multi-sample track-length estimate of the transmittance up to that location (see Section 5.2.2). The longest of these distances effectively determines the length of such a “short” beam. Belcour et al. [BBS14] improved visual convergence of progressive photon beams by analyzing absorption and scattering in the Fourier domain and identifying situations in which the radius can stop decreasing.

Gathering multi-scattered contributions from photon beams was also accomplished using the diffusion approximation. Habel et al. [HCJ13] sample points on a beam entering a translucent material and used the dipole model [JMLH01] to approximate the sub-surface transport to surface points. d’Eon [d’E14] extended their approach by using the dipole model to connect to refracted camera paths, making the BSSRDF model reciprocal and more accurate.

7.4. Combining density estimation and path tracing

Bidirectional path tracing (Section 7.1) is a more robust method than unidirectional path tracing thanks to its effective combination of the strengths of various path-sampling techniques into one unified algorithm via MIS. Similarly, the various volumetric density

estimators based on points and beams have their respective pros and cons. Recently, Křivánek et al. [KGH*14] analyzed the relative efficiency of these estimators and showed that density estimation based on point primitives yields lower variance in dense media, while sparse media are better handled by estimators based on beams. They also reformulated beam-based radiance estimators [JNSJ11] as path-integral estimators, which enables their MIS combination also with unbiased path-sampling techniques from bidirectional path tracing. The resulting algorithm, *unified points, beams, and paths* (UPBP), provides a more robust solution that is able to render scenes with different kinds of media, where the individual estimators may each fail in complementary ways (see Figure 11).

7.4.1. Photon planes, volumes, and beyond

The 1D photon beam sampling primitive can bring a significant improvement over 0D photon points by reducing variance as well as bias, though not eliminating it completely. Most recently, Bitterli and Jarosz [BJ17] presented a general theory of volumetric light-transport simulation using higher-dimensional samples, including 2D photon “planes” and 3D photon “volumes”; see Figure 12. They extended Křivánek et al.’s [KGH*14] efficiency analysis to reveal that when photon beams already improve over points, the higher-dimensional samples provide even greater benefits.

Bitterli and Jarosz [BJ17] also show that already the 2D photon planes allow for a fully unbiased formulation of volumetric density estimation, though in practice adding a small amount of blur helps reduce variance. The intuition why blur is no longer strictly required is that the probability of an infinitely thin camera ray intersecting an infinitely thin random photon plane is non-zero. In contrast, the probability of two infinitely thin rays (beams) to intersect is zero, so some form of blur is required for a non-zero contribution.

7.5. Volumetric radiance caching

Both many-light algorithms and density-estimation approaches reuse light subpaths to increase efficiency. Jarosz et al. [JDZJ08]



Figure 13: A scene with a fog rendered using 175K cache points and 27M cache queries with volumetric radiance caching [JDZJ08]. The insets provide an equal-time comparison to path tracing. Images reproduced from [JDZJ08].

propose a different method for caching and reusing computation inspired by irradiance and radiance caching for surfaces [WRC88, KGPB05]. Their method sparsely samples and caches radiance estimates in an octree. The estimates are then extrapolated using gradients and reused for nearby locations visited by camera paths. The key to the approach is accurately computing gradients of the radiance (for high-quality extrapolation) and estimating the induced extrapolation error (to know when to stop extrapolating). Jarosz et al. [JDZJ08] first derived gradients (and the corresponding error) while ignoring changes in occlusion, and later [JZJ08b] showed how to account for occlusion gradients at surfaces in the presence of media. Most recently, Marco et al. [MJJG18] derived a second-order, occlusion-aware approach inspired by recent Hessian-based irradiance caching approaches for surfaces [JSKJ12, SJJ12]. Being a view-driven method, such volumetric radiance caching approaches can handle large scenes well, outperforming path tracing and also some of the density estimation methods. Figure 13 provides a comparison to a vanilla unidirectional path tracer.

7.6. Most-probable-path methods

For dense and highly forward-scattering media, transport paths can be characterized by many collision events and low deviation of the outgoing direction from the incoming. Ignoring backward scattering, the resulting paths are similar to continuous curves with low curvature. These observations motivate a family of methods based on so-called *most-probable paths*. These techniques were introduced to graphics by Premoze et al. [PAS03, PAT*04] based on theory initially developed by Tessendorf in a string of publications [Tes87, Tes09]. Frederickx and Dutré [FD17] provide an excellent summary of the derivation of these methods in their supplemental material. Note that in this context, *path integral* refers to the Feynman integral [FH65], not to what we describe in Equation (7).

While these approaches are mostly used as a closed-form approximation in the sense of a dipole model, they can be used to construct transport paths. Weber et al. [WHD17] presented a simplified multi-vertex next-event estimation scheme based on these ideas. They use a straight line as the point of the expansion instead of the most-probable path. This approach is motivated by the observation that for Henyey-Greenstein phase functions the penalty in path-measurement contribution incurred by one single kink is much lower than the penalty for uniform curvature spread out over many

scattering events. That is, the phase function has a much sharper falloff near $\cos\theta = 1$ than in the rest of the distribution.

This is also analyzed in appendix L.1 in [FD17] and is even more true for multi-lobe approximations of phase functions predicted by Lorenz-Mie theory, which include more back-scattering contribution. In that sense the most-probable-path theory is an approximation that neglects the back scattering effects present in scattering regimes governed by Mie theory.

7.7. Markov-chain methods

Metropolis light transport (MLT) [Vea97] can be used to render participating media [PKK00]. In general, to compute the Metropolis-Hastings acceptance probability the measurement-contribution function needs to be evaluated exactly. This means that stochastic approximations utilizing null collisions to evaluate transmittance can be problematic. The Metropolis-Hastings scheme can still be used if the additional random variates needed to trace through the volume are explicitly added to the state of the Markov chain. This is similar to what has been shown by Schüssler et al. [SHHD17] for stochastic multiple scattering inside microfacet surfaces in their supplemental materials. There are approaches to analyze the convergence property of the Markov chain in this case [KBSU15]. As pointed out by Raab et al. [RSK08], running the Markov chain in primary sample space [KSKAC02], i.e. on the random numbers, circumvents this problem: the contribution of a sample given all the random numbers is deterministic and can be evaluated. This formulation, however, may require an unbounded number of random numbers for every path segment that passes through a medium, and can lead to rippling effects inhibiting the efficiency of local exploration. This latter issue has recently been investigated for the case of surface transport [Pan17, BJNJ17].

7.8. Zero-variance random walks

Zero-variance random-walk theory is an intriguing concept: it analyzes the possibility to create random walks without variance. This means that at every scattering point, an outgoing direction as well as a distance to the next scattering event has to be perfectly importance sampled by all terms of the measurement equation: the product of phase function, transmittance, and incoming importance (or radiance, depending on the direction of the random walk). This has been explored by Dwivedi [Dwi82b, Dwi82a] and Booth [Boo87]. Hoogenboom [Hoo08] provides a good overview.

In practice, it is hard to obtain a good estimate of the incoming-radiance field, but in simple canonical cases a closed-form approximation can be used to guide importance sampling and reduce variance. This has been explored for dense, isotropic, highly scattering media: Křivánek and d'Eon [Kd14] locally fitted a slab to the bounding geometry (see Figure 14, left). They demonstrated that this can substantially improve the variance of path traced subsurface scattering such as in skin, by quickly guiding random walks back to the surface where light contribution can be picked up.

Meng et al. [MHD16] later extended this idea for cases where quickly escaping back through the entered surface is not a good

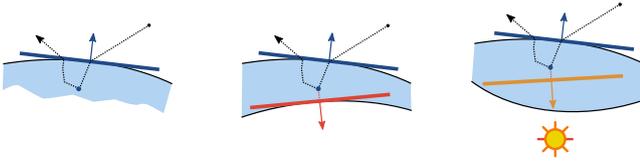


Figure 14: Illustration of random-walk-biasing schemes inspired by Dwivedi. Left: fitting a local slab to escape the volume close to the point of entry works well for thick volumes. Center: finding the closest surface point at every internal scattering point and biasing towards the surface point can help escaping thin geometric features. Right: explicitly biasing towards a light source improves backlit scenarios. Figure reproduced from [MHD16].

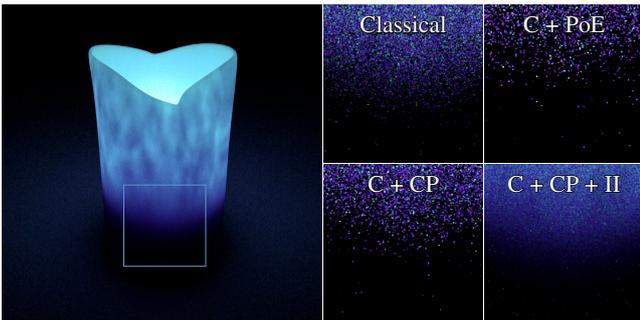


Figure 15: Classical random walks (top left inset) can be noisy for dense media like candle wax. Biasing the random walk towards the point of entry (PoE) [Kd14] increases variance even more in this case, since most of the light comes from the inside. The closest point heuristic (CP) does not improve matters much as the geometry is not thin enough. Illumination biasing (II, bottom right inset) has been designed for such cases. Images reproduced from [MHD16].

idea. This is the case for thin features of geometry such as ears, especially in backlit scenarios. They devised two more biasing methods to take geometric features into account: one method searches for the closest point to the boundary at every scattering vertex (cf. Figure 14, center). This increases the chances to escape thin geometry. The second method explicitly takes incident illumination into account, specifically improving the variance of the random walk for backlit cases (cf. Figure 14, right). This last case would actually suffer from *increased* variance when using the biasing method by Krivánek and d'Eon [Kd14], as can be seen in Figure 15.

7.9. Emissive media

Lighting setups including emissive media that subtend a very small solid angle from most shading points, such as candle lights, can cause high variance. A small solid angle means that the chances to intersect these light sources accidentally are very low. It is thus necessary to include some kind of next-event estimation, explicitly sampling locations inside potentially small and thin flames. This has been done for voxel grids [VH13]. Problems also arise when distances are sampled proportional to transmittance as it is very unlikely to create a collision in extremely thin media such as candle flames. This is true even for null-collision methods that generate additional tentative collisions. Simon et al. [SHZD17] propose to

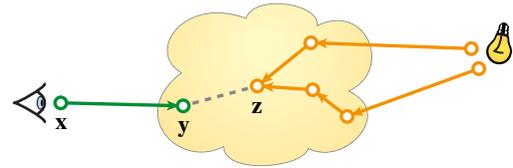


Figure 16: A difficulty with BDPT and emissive media. Accounting for the emission at the interior path vertex z requires evaluating a marginal PDF, which involves integration over all possible light subpaths arriving at z (colored in orange).

use an estimator that performs line integration: they utilize regular tracking to deterministically pick up all emission between two path vertices. The authors observe that combining such an estimator with next-event estimation deteriorates performance and address the issues by modifying the next-event estimation accordingly: they start by sampling a path vertex inside an emissive volume, but then use only the direction to this point and independently sample a free-flight distance. Then, an estimator is constructed which picks up emission along all of the resulting path segment. For this, the PDF needs to be evaluated as a marginal integral over the half-open line segment. This is sped up by using a coarse voxelization. Next-event estimation performed this way is able to construct paths with end points outside the emissive medium, for instance on non-emissive surface geometry. This crucial property makes sure an MIS combination with BSDF or phase function sampling stays effective.

When using bidirectional path tracing (BDPT), accounting for emission along the interior of a path is even more difficult. Figure 16 illustrates the problem, where a light transport path connecting the light source and the camera passes through an emissive participating medium. Accounting for the emission at vertex z requires constructing an estimator for the path $\bar{x} = (x, y, z)$, which in turn involves computing its PDF. This can be difficult because that PDF includes a marginalization: we need to account for all possible ways to sample z from a light source, i.e. to integrate the PDFs for sampling light subpaths of any length connecting a light source and z . The figure shows two such light subpaths. This is not an issue when emission is only considered at the end points of paths.

8. Acceleration data structures

Since data access usually dominates the overall render time, acceleration data structures are key for achieving good performance when rendering heterogeneous volumes. An L2 cache reference is 14–25× faster than a DDR memory reference [Lev09]. This factor can almost directly be observed as total render time difference when comparing algorithms with good caching behaviour to ones with poor access patterns. As volume data can quickly grow into the hundreds of gigabytes for production scenarios, this problem is aggravated by out of core memory accesses to disk.

In order to store spatially-varying optical properties, two open-source libraries—Field3D [WZC*10] and OpenVDB [Mus13]—have gained popularity in the movie industry. Adapting the resolution, or using a coarser representation of the field, can further improve performance when the voxel detail is below the image resolution, or when tracing high-order bounces.

Accelerators for null-collision methods. The performance of stochastic trackers depends heavily on how tightly majorants (and minorants) bound the extinction function. The majorants directly impact the expected number of null collisions and they are typically precomputed and stored in a separate data structure, e.g. super voxels [SKTM11, NSJ14], at a resolution coarser than the actual volume. Adaptive variants such as kd-trees [YIC*10, YIC*11] or octrees [KHLN17] allow one to adjust the memory cost locally and can yield substantial performance benefits. Yue et al. [YIC*11] provide guidance for optimally subdividing space to minimize the number of null collisions.

In order to sample tentative collisions, the tracking algorithm steps through the structure that stores the majorants until it reaches a sampled value of (majorant) optical thickness [SKTM11], or samples a tentative collision in each voxel independently proceeding to the next one if the collision occurs outside of the current voxel [YIC*10]. Once the tentative collision is found, the algorithm potentially starts traversing another hierarchical data structure—the one that holds extinction values—at the root node. This can incur a large memory-access overhead. One may thus prefer to use a single hierarchical accelerator, e.g. an octree, storing the fine optical properties in leaves and filtered values and majorants in interior nodes.

Accelerators for regular tracking. Regular tracking may be a viable alternative to stochastic samplers in certain situations, especially when adaptive memory structures, such as nested grids [SHZD17], are used. The structure should allow predictable memory accesses of consecutive voxels, which in turn enables the cache to amortize the memory-access cost over many extinction lookups. Another important aspect to consider is whether an interpolation scheme is required. Nested grids typically provide a fast access to immediate neighbors due to their rigid structure. If high-quality (high-order) interpolation schemes are desired, regular tracking at coarse levels may become infeasible and it may be more efficient to track through the finer resolution instead.

The performance of individual trackers depends heavily on the characteristics of the medium and the used storage. In general, null collision-based trackers are fast if the majorant is tight for a long distance along the ray, and many fine-detail voxels can thus be skipped. A well-designed adaptive data structure is key for achieving good performance with regular tracking in such scenarios, and its noise-free nature may provide a competitive advantage.

Motion blur. To simulate motion blur in case of time-varying volumes, the particle density may be transported on the fly to match the current time by employing Eulerian motion blur [KK07]. This can be combined with stochastic trackers or ray marching since these only require a point query. Regular tracking works better when the data is stored in a 4D grid (space and time) [Wre16]. Majorants for stochastic tracking will need to be stored either conservatively over the whole shutter interval, or in a time-resolved manner, too.

9. Remaining challenges and open problems

Simulating light transport in participating media efficiently is a challenging problem. Individual methods surveyed in this report are typically well-suited for particular classes of problems,

and while some effort has been made in unifying different approaches [KGH*14], having a solution that handles vastly different scenarios without significant computational overhead is still an open problem. Currently, combinations that employ MIS typically induce an extra cost, and an intimate technical knowledge of all existing techniques is required to efficiently render a specific scene.

Null-collision methods and MIS. The integration of null-collision methods into MIS-based estimators, in particular, is currently very difficult. The rejection-sampling nature of these methods makes it currently impossible to evaluate the free-path PDF exactly, and one needs to compute the MIS weight using deterministic approximations that incur suboptimal weighting and computational overhead. Another open question, specific to null-collision transmittance estimators, is how to choose the optimal tracker, or when to switch from one to the other. Figure 5 studies the performance of these estimators in a canonical setting only, and devising a robust switching criterion for general use remains an open problem.

Machine learning. Machine learning, in general, has a great potential of accelerating light-transport simulations. The vast cost of accessing voxelized data and the high-albedo nature of certain volumetrics (e.g. clouds) make it challenging for MC methods to deliver the results at tractable costs. Importance sampling such high-dimensional spaces is hard, especially when peaky forward-scattering phase functions are used. The high expense can be mitigated by incorporating various aggregators [MWM07, MPH*15, MPG*16], diffusion approximations [JMLH01, dl11, HCJ13, FHK14, KPS*14], or deep learning [KMM*17], but this always comes at the cost of sacrificing unbiasedness. Applying machine learning at the level of distance and directional samples, as has been proposed for surface rendering [VKv*14, VK16, HEV*16, MGN17], or other forms of path guiding could potentially provide significant practical benefits.

Joint handling of surfaces and volumes. Producing a final render usually means including volumes and surfaces in the same scene. This is already challenging from an asset creation standpoint, since often volume data sets come from simulations while surfaces are created in a different workflow. The same separation can be observed in specialized algorithms tailored either for volumes or for surfaces. This can go as far as using completely different renderers for the two tasks, leading to problems when the two results are to be combined. Thus, a unified scene representation incorporating both surfaces and volumes in the same context might lend itself better towards adaptive level of detail.

Generalizations. Generalizing some of the aforementioned methods to anisotropic media still remains to be investigated, and very little research (in graphics) has so far been devoted to discrete random media [MWM07, MPH*15, MPG*16] or media where the positions of scatterers are not independent, but correlated. Handling such scenarios will require a generalized form of the classical RTE. Transient rendering is another area of transport simulation that requires resolving light paths in time. Methods for constructing paths that attain a specific total length (duration) currently exist only for two-segment paths [JMM*14]. Building a time-constrained light trajectory with multiple bounces remains to be developed.

Towards automatic algorithms. While the adoption of Monte Carlo methods for (volume) rendering has improved artistic workflow, physical accuracy, and usability, we are still far from the goal of a fully automatic rendering process. In a way, the state of the art in this field is much like that of photography in the early 1900s. Then, a photographer had to be a master technician with knowledge of physics, chemistry and engineering to adjust camera parameters based on the particulars of each captured scene. Consumer smartphone cameras now make these decisions automatically for the typical photographer, but we are a long way from that goal for Monte Carlo rendering.

10. Conclusion

We have reviewed and summarized components and complete algorithms for physically-based rendering of scenes containing participating media. We focused on methods that employ various forms of Monte Carlo integration. These provide practical solutions to transport problems that arise in visual effects, feature animation, architectural design, and product visualization, and are currently heavily employed for rendering images in the corresponding industries. Our report can serve as a starting point for practitioners to get quickly oriented in the field. It also provides a reference for researchers and links for those interested in developments in nuclear sciences, neutron transport, and particle physics. Further connections can be found in the book by d'Eon [d'E16], which covers closed-form solutions and complements the overview of numerical recipes provided here.

Acknowledgments

We thank Jaroslav Křivánek for insightful comments, Peter Kutz for tracing down many of the early publications presenting the concept of delta tracking, and anonymous reviewers for suggesting numerous improvements to the overall structure and exposition.

References

- [ABW14] AMENT M., BERGMANN C., WEISKOPF D.: Refractive radiative transfer equation. *ACM TOG* 33, 2 (Apr. 2014), 17:1–17:22. 2
- [AK90] ARVO J., KIRK D.: Particle transport and image synthesis. *Computer Graphics (Proc. SIGGRAPH)* 24, 4 (Sept. 1990), 63–66. 11, 12
- [Arv86] ARVO J.: Backward ray tracing. In *ACM SIGGRAPH '86 Course Notes - Developments in Ray Tracing* (1986), pp. 259–263. 18
- [Arv95] ARVO J.: Stratified sampling of spherical triangles. In *Proc. SIGGRAPH 95* (New York, NY, USA, 1995), ACM, pp. 437–438. 14
- [Arv01] ARVO J.: Stratified sampling of 2-manifolds. In *Proc. SIGGRAPH 01 Courses* (Aug. 2001). 14
- [AW87] AMANATIDES J., WOO A.: A fast voxel traversal algorithm for ray tracing. In *Eurographics '87* (1987), pp. 3–10. 5
- [BBS14] BELCOUR L., BALA K., SOLER C.: A local frequency analysis of light scattering and absorption. *ACM TOG* 33, 5 (Sept. 2014), 163:1–163:17. 19
- [Ber63] BERTINI H. W.: *Monte Carlo simulations on intranuclear cascades*. Tech. Rep. ORNL–3383, Oak Ridge National Laboratory, Oak Ridge, TN, USA, Apr. 1963. 6
- [BJ17] BITTERLI B., JAROSZ W.: Beyond points and beams: Higher-dimensional photon samples for volumetric light transport. *ACM TOG (Proc. SIGGRAPH)* 36, 4 (July 2017). 9, 19
- [BJN17] BITTERLI B., JAKOB W., NOVÁK J., JAROSZ W.: Reversible jump Metropolis light transport using inverse mappings. *ACM TOG* 37, 1 (Oct. 2017). 20
- [BM58] BUTCHER J. C., MESSEL H.: Electron number distribution in electron-photon showers. *Phys. Rev.* 112 (Dec. 1958), 2096–2106. 6
- [BM60] BUTCHER J. C., MESSEL H.: Electron number distribution in electron-photon showers in air and aluminium absorbers. *Nuclear Physics* 20 (1960), 15–128. 6
- [BM01] BROWN F. B., MARTIN W. R.: *Monte Carlo particle transport in media with exponentially varying time-dependent cross-sections*. Tech. Rep. LA-UR-01-0675, Los Alamos National Laboratory, 2001. 5
- [BM03] BROWN F. B., MARTIN W. R.: Direct sampling of Monte Carlo flight paths in media with continuously varying cross-sections. In *Proc. of ANS Mathematics & Computation Topical Meeting* (April 2003), pp. 6–11. 5
- [Boo87] BOOTH T. E.: Generalized zero-variance solutions and intelligent random numbers. In *Proceedings of the 19th Conference on Winter Simulation* (1987), WSC '87, pp. 445–451. 20
- [Boo07] BOOTH T. E.: Unbiased Monte Carlo estimation of the reciprocal of an integral. *Nuclear Science and Engineering* 156, 3 (2007), 403–407. 16
- [CCT72] CARTER L. L., CASHWELL E. D., TAYLOR W. M.: Monte Carlo sampling with continuously varying cross sections along flight paths. *Nuclear Science and Engineering* 48, 4 (1972), 403–411. 8
- [CCY67] COVEYOU R. R., CAIN V. R., YOST K. J.: Adjoint and importance in Monte Carlo application. *Nuclear Science and Engineering* 27, 2 (1967), 219–234. 14
- [Cha60] CHANDRASEKHAR S.: *Radiative transfer*. Dover Publications, 1960. 3
- [CJ16] CHRISTENSEN P. H., JAROSZ W.: The path to path-traced movies. *Foundations and Trends in Computer Graphics and Vision* 10, 2 (October 2016), 103–175. 1
- [Col68] COLEMAN W. A.: Mathematical verification of a certain Monte Carlo sampling technique and applications of the technique to radiation transport problems. *Nuclear Science and Engineering* 32, 1 (Apr. 1968), 76–81. 6, 7, 8
- [CPP*05] CEREZO E., PÉREZ F., PUEYO X., SERON F. J., SILLION F. X.: A survey on participating media rendering techniques. *The Visual Computer* 21, 5 (2005), 303–328. 1
- [Cra78] CRAMER S. N.: Application of the fictitious scattering radiation transport model for deep-penetration Monte Carlo calculations. *Nuclear Science and Engineering* 65, 2 (1978), 237–253. 7, 8, 11, 12
- [d'E14] D'EON E.: A dual-beam 3d searchlight BSSRDF. In *ACM SIGGRAPH 2014 Talks* (New York, NY, USA, 2014), ACM, pp. 65:1–65:1. 19
- [d'E16] D'EON E.: *A Hitchhiker's Guide to Multiple Scattering*. 2016. 23
- [DGMF11] DELALANDRE C., GAUTRON P., MARVIE J.-E., FRANÇOIS G.: Transmittance function mapping. In *Proc. Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2011), ACM, pp. 31–38. 17
- [dI11] D'EON E., IRVING G.: A quantized-diffusion model for rendering translucent materials. *ACM TOG (Proc. SIGGRAPH)* 30, 4 (July 2011), 56:1–56:14. 22
- [DKH*13] DACHSBACHER C., KŘIVÁNEK J., HAŠAN M., ARBREE A., WALTER B., NOVÁK J.: Scalable realistic rendering with many-light methods. *CGF* 33, 1 (2013), 88–104. 16
- [Dwi82a] DWIVEDI S. R.: A new importance biasing scheme for deep-penetration Monte Carlo. *Am. nucl. Energy* 9 (1982), 359–368. 20
- [Dwi82b] DWIVEDI S. R.: Zero variance biasing schemes for Monte Carlo calculations of neutron and radiation transport problems. *Nuclear Science and Engineering* 80, 1 (1982), 172–178. 20

- [ENSD12] ENGELHARDT T., NOVÁK J., SCHMIDT T.-W., DACHS-BACHER C.: Approximate bias compensation for rendering scenes with heterogeneous participating media. *CGF (Proc. Pacific Graphics)* 31, 7 (Sept. 2012), 2145–2154. [16](#), [17](#)
- [EPG*13] EYMET V., POITOU D., GALTIER M., HAFI M. E., TERRÉE G., FOURNIER R.: 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 (Apr. 2013), 145–157. [6](#), [7](#), [8](#), [16](#)
- [FBD15] FREDERICKX R., BARTELS P., DUTRÉ P.: Adaptive lightslice for virtual ray lights. In *Eurographics 2015 - Short Papers* (2015), The Eurographics Association, pp. 61–64. [17](#)
- [FD17] FREDERICKX R., DUTRÉ P.: A forward scattering dipole model from a functional integral approximation. *ACM TOG* 36, 4 (July 2017). [20](#)
- [FH65] FEYNMAN R., HIBBS A.: *Quantum mechanics and path integrals*. International series in pure and applied physics. McGraw-Hill, 1965. [20](#)
- [FHF*17] FASCIONE L., HANIKA J., FAJARDO M., CHRISTENSEN P., BURLEY B., GREEN B., PIEKÉ R., KULLA C., HERY C., VILLEMEN R., HECKENBERG D., MAZZONE A.: SIGGRAPH 2017 course notes: Path tracing in production (parts 1 and 2). In *SIGGRAPH Courses* (Aug. 2017). [1](#)
- [FHK14] FRISVAD J. R., HACHISUKA T., KJELSDEN T. K.: Directional dipole model for subsurface scattering. *ACM TOG* 34, 1 (Dec. 2014), 5:1–5:12. [22](#)
- [Gam16] GAMITO M. N.: Solid angle sampling of disk and cylinder lights. *CGF (Proc. EGSR)* 35, 4 (2016), 25–36. [14](#)
- [GBC*13] GALTIER M., BLANCO S., CALIOT C., COUSTET C., DAUCHET J., HAFI M. E., EYMET V., FOURNIER R., GAUTRAIS J., KHUONG A., PIAUD B., TERRÉE G.: Integral formulation of null-collision Monte Carlo algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer* 125 (Apr. 2013), 57–68. [6](#), [7](#), [8](#), [13](#)
- [GBD*16] GALTIER M., BLANCO S., DAUCHET J., EL HAFI M., EYMET V., FOURNIER R., ROGER M., SPIESSER C., TERRÉE G.: Radiative transfer and spectroscopic databases: A line-sampling Monte Carlo approach. *Journal of Quantitative Spectroscopy and Radiative Transfer* 172 (Mar. 2016), 83–97. [8](#)
- [GDML13] GAUTRON P., DELALANDRE C., MARVIE J.-E., LECOCQ P.: Boundary-aware extinction mapping. *CGF* 32, 7 (2013), 305–314. [17](#)
- [GKH*13] GEORGIEV I., KRÍVÁNEK J., HACHISUKA T., NOWROUZEZHAI D., JAROSZ W.: Joint importance sampling of low-order volumetric scattering. *ACM TOG (Proc. SIGGRAPH Asia)* 32, 6 (Nov. 2013), 164:1–164:14. [9](#), [14](#), [15](#), [17](#)
- [GUK*17] GUILLÉN I., UREÑA C., KING A., FAJARDO M., GEORGIEV I., LÓPEZ-MORENO J., JARABO A.: Area-preserving parameterizations for spherical ellipses. *CGF (Proc. EGSR)* 36, 4 (2017), 179–187. [14](#)
- [HCJ13] HABEL R., CHRISTENSEN P. H., JAROSZ W.: Photon beam diffusion: a hybrid Monte Carlo method for subsurface scattering. *CGF (Proc. EGSR)* 32, 4 (June 2013). [19](#), [22](#)
- [HEV*16] HERHOLZ S., ELEK O., VORBA J., LENSCH H., KRÍVÁNEK J.: Product Importance Sampling for Light Transport Path Guiding. *CGF (Proc. EGSR)* (2016). [22](#)
- [HG41] HENYEV L. G., GREENSTEIN J. L.: Diffuse radiation in the galaxy. *Astrophysical Journal* 93 (Jan. 1941), 70–83. [2](#)
- [HKWB09] HAŠAN M., KRÍVÁNEK J., WALTER B., BALA K.: Virtual spherical lights for many-light rendering of glossy scenes. *ACM TOG (Proc. SIGGRAPH Asia)* 28, 5 (Dec. 2009), 1–6. [17](#)
- [HOJ08] HACHISUKA T., OGAKI S., JENSEN H. W.: Progressive photon mapping. *ACM TOG (Proc. SIGGRAPH Asia)* 27, 5 (Dec. 2008), 130:1–130:8. [18](#), [19](#)
- [Hoo08] HOOGENBOOM J. E.: Zero-variance Monte Carlo schemes revisited. *Nucl. Sci. Eng.* 160, 1 (2008), 1–22. [20](#)
- [HTAT*06] HUBERT-TREMBLAY V., ARCHAMBAULT L., TUBIC D., ROY R., BEAULIEU L.: Octree indexing of DICOM images for voxel number reduction and improvement of Monte Carlo simulation computing efficiency. *Medical Physics* 33, 8 (2006), 2819–2831. [5](#)
- [HWH*16] HUO Y., WANG R., HU T., HUA W., BAO H.: Adaptive matrix column sampling and completion for rendering participating media. *ACM TOG (Proc. SIGGRAPH Asia)* 35, 6 (Nov. 2016), 167:1–167:11. [17](#)
- [JA18] JARABO A., ARELLANO V.: Bidirectional rendering of vector light transport. *CGF* (2018), n/a–n/a. awaiting publication. [2](#), [4](#)
- [JAM*10] JAKOB W., ARBREE A., MOON J. T., BALA K., MARSCHNER S.: A radiative transfer framework for rendering materials with anisotropic structure. *ACM TOG* 29, 4 (July 2010), 53:1–53:13. [2](#)
- [JB10] JANSEN J., BAVOIL L.: Fourier opacity mapping. In *Proc. Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2010), ACM, pp. 165–172. [17](#)
- [JC98] JENSEN H. W., CHRISTENSEN P. H.: Efficient simulation of light transport in scenes with participating media using photon maps. In *Proc. SIGGRAPH 98* (New York, NY, USA, 1998), Annual Conference Series, ACM, pp. 311–320. [11](#), [18](#)
- [JDZJ08] JAROSZ W., DONNER C., ZWICKER M., JENSEN H. W.: Radiance caching for participating media. *ACM TOG* 27, 1 (Mar. 2008), 7:1–7:11. [19](#), [20](#)
- [Jen06] JENSEN J. L. W. V.: Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Math.* 30 (1906), 175–193. [10](#)
- [Jen96] JENSEN H. W.: Global illumination using photon maps. In *Proc. Eurographics Workshop on Rendering Techniques* (London, UK, Aug. 1996), Springer-Verlag, pp. 21–30. [11](#), [18](#)
- [JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Proc. SIGGRAPH 01* (New York, NY, USA, 2001), Annual Conference Series, ACM, pp. 511–518. [19](#), [22](#)
- [JMM*14] JARABO A., MARCO J., MUÑOZ A., BUISAN R., JAROSZ W., GUTIERREZ D.: A framework for transient rendering. *ACM TOG (Proc. SIGGRAPH Asia)* 33, 6 (2014). [22](#)
- [JNSJ11] JAROSZ W., NOWROUZEZHAI D., SADEGHI I., JENSEN H. W.: A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM TOG* 30, 1 (Feb. 2011), 5:1–5:19. [18](#), [19](#)
- [JNT*11] JAROSZ W., NOWROUZEZHAI D., THOMAS R., SLOAN P.-P., ZWICKER M.: Progressive photon beams. *ACM TOG (Proc. SIGGRAPH Asia)* 30, 6 (Dec. 2011), 181:1–181:12. [11](#), [19](#)
- [JRJ11] JAKOB W., REGG C., JAROSZ W.: Progressive expectation-maximization for hierarchical volumetric photon mapping. *Computer Graphics Forum (Proceedings of EGSR)* 30, 4 (June 2011). [18](#)
- [JSKJ12] JAROSZ W., SCHÖNEFELD V., KOBELT L., JENSEN H. W.: Theory, analysis and applications of 2d global illumination. *ACM TOG* 31, 5 (Sept. 2012), 125:1–125:21. [20](#)
- [JZJ08a] JAROSZ W., ZWICKER M., JENSEN H. W.: The beam radiance estimate for volumetric photon mapping. *CGF (Proc. Eurographics)* 27, 2 (Apr. 2008), 557–566. [18](#)
- [JZJ08b] JAROSZ W., ZWICKER M., JENSEN H. W.: Irradiance gradients in the presence of participating media and occlusions. *CGF (Proc. EGSR)* 27, 4 (June 2008), 1087–1096. [20](#)
- [Kaj86] KAJIYA J. T.: The rendering equation. *Computer Graphics (Proc. SIGGRAPH)* (1986), 143–150. [3](#)
- [KBSU15] KRONANDER J., B. SCHÖN T., UNGER J.: Pseudo-marginal Metropolis light transport. In *Proceedings of SIGGRAPH ASIA Technical Briefs* (2015). [20](#)
- [KC77] KALLI H. J., CASHWELL E. D.: *Evaluation of three Monte Carlo estimation schemes for flux at a point*. Tech. Rep. LA–6865-MS, Los Alamos Scientific Lab., 1977. [14](#)

- [Kd14] KRÍVÁNEK J., D'EON E.: A zero-variance-based sampling scheme for Monte Carlo subsurface scattering. In *ACM SIGGRAPH 2014 Talks* (New York, NY, USA, 2014), ACM, pp. 66:1–66:1. 20, 21
- [Kel97] KELLER A.: Instant radiosity. In *Proc. SIGGRAPH 97* (New York, NY, USA, 1997), Annual Conference Series, ACM Press/Addison-Wesley Publishing Co., pp. 49–56. 16
- [KF12] KULLA C., FAJARDO M.: Importance sampling techniques for path tracing in participating media. *CGF (Proc. EGSR) 31*, 4 (June 2012), 1519–1528. 9, 14
- [KGH*14] KRÍVÁNEK J., GEORGIEV I., HACHISUKA T., VÉVODA P., ŠÍK M., NOWROUZEZAHRAI D., JAROSZ W.: Unifying points, beams, and paths in volumetric light transport simulation. *ACM TOG (Proc. SIGGRAPH) 33*, 4 (2014), 103:1–103:13. 11, 18, 19, 22
- [KGPB05] KRÍVÁNEK J., GAUTRON P., PATTANAİK S., BOUATOUCH K.: Radiance caching for efficient global illumination computation. *IEEE TVCG 11*, 5 (Sept.-Oct. 2005), 550–561. 20
- [KHLN17] KUTZ P., HABEL R., LI Y. K., NOVÁK J.: Spectral and de-composition tracking for rendering heterogeneous volumes. *ACM TOG (Proc. SIGGRAPH) 36*, 4 (July 2017). 7, 8, 9, 11, 13, 16, 22
- [KK06] KOLLIG T., KELLER A.: Illumination in the presence of weak singularities. In *Monte Carlo and Quasi-Monte Carlo Methods 2004*. Springer, 2006, pp. 245–257. 17
- [KK07] KIM D., KO H.-S.: Eulerian motion blur. In *Eurographics Workshop on Natural Phenomena* (2007). 22
- [KMM*17] KALLWEIT S., MÜLLER T., MCWILLIAMS B., GROSS M., NOVÁK J.: Deep scattering: Rendering atmospheric clouds with radiance-predicting neural networks. *ACM TOG (Proc. SIGGRAPH Asia) 36*, 6 (Nov. 2017). 22
- [KPS*14] KOERNER D., PORTSMOUTH J., SADLO F., ERTL T., EBERHARDT B.: Flux-limited diffusion for multiple scattering in participating media. *CGF 33*, 6 (Sept. 2014), 178–189. 22
- [KSKAC02] KELEMEN C., SZIRMAY-KALOS L., ANTAL G., CSOKA F.: A simple and robust mutation strategy for the Metropolis light transport algorithm. *CGF 21*, 3 (2002), 531–540. 20
- [KZ11] KNAUS C., ZWICKER M.: Progressive photon mapping: a probabilistic approach. *ACM TOG 30*, 3 (May 2011), 25:1–25:13. 17, 19
- [Lam60] LAMBERT J. H.: *Photometria: Sive de Mensura et Gradibus Luminis, Colorum et Umbrae*. 1760. 3
- [Lep10] LEPPÄNEN J.: 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. 5
- [Lev09] LEVINTHAL D.: Performance analysis guide for Intel Core i7 processor and Intel Xeon 5500 processors. https://software.intel.com/sites/products/collateral/hpc/vtune/performance_analysis_guide.pdf, 2009. 21
- [LK91] LUX I., KOBLINGER L.: *Monte Carlo particle transport methods: neutron and photon calculations*. CRC Press, 1991. 10
- [Lor90] LORENZ L.: Lysbevægelsen i og uden for en af plane lysbølger belyst kugle. In *Det Kongelige Danske Videnskabernes Selskabs Skrifter (trykt utg.): Naturvidenskabelig og Mathematisk Afdeling* (1890). 2
- [LV00] LOKOVIC T., VEACH E.: Deep shadow maps. In *Proc. SIGGRAPH 00* (New York, NY, USA, 2000), Annual Conference Series, ACM Press/Addison-Wesley Publishing Co., pp. 385–392. 17
- [MGN17] MÜLLER T., GROSS M., NOVÁK J.: Practical path guiding for efficient light-transport simulation. *CGF (Proc. EGSR) 36*, 4 (June 2017), 91–100. 22
- [MHD16] MENG J., HANIKA J., DACHSBACHER C.: Improving the Divedi sampling scheme. *Computer Graphics Forum (Proceedings of Eurographics Symposium on Rendering) 35*, 4 (June 2016), 37–44. 20, 21
- [Mie08] MIE G.: Beiträge zur Optik trüber Medien, speziell kolloidaler Metallösungen. *Annalen der Physik 330* (1908), 377–445. 2
- [Mil67] MILLER L. B.: *Monte Carlo analysis of reactivity coefficients in fast reactors; general theory and applications*. Tech. Rep. ANL-7307, Argonne National Laboratory, Argonne, IL, USA, Mar. 1967. 6, 7, 8
- [MJG18] MARCO J., JARABO A., JAROSZ W., GUTIERREZ D.: Second-order occlusion-aware volumetric radiance caching. *ACM TOG 37*, 2 (Feb. 2018). 20
- [MK15] MORGAN L. W. G., KOTLYAR D.: Weighted-delta-tracking for Monte Carlo particle transport. *Annals of Nuclear Energy 85* (2015), 1184–1188. 8
- [MPG*16] MÜLLER T., PAPAS M., GROSS M., JAROSZ W., NOVÁK J.: Efficient rendering of heterogeneous polydisperse granular media. *ACM TOG (Proc. SIGGRAPH Asia) 35*, 6 (2016), 168:1–168:14. 22
- [MPH*15] MENG J., PAPAS M., HABEL R., DACHSBACHER C., MARSCHNER S., GROSS M., JAROSZ W.: Multi-scale modeling and rendering of granular materials. *ACM TOG (Proc. SIGGRAPH) 34*, 4 (July 2015), 49:1–49:13. 22
- [Muñ14] MUÑOZ A.: Higher order ray marching. *CGF 33*, 8 (2014), 167–176. 6
- [Mus13] MUSETH K.: VDB: High-resolution sparse volumes with dynamic topology. *ACM TOG 32*, 3 (July 2013), 27:1–27:22. 21
- [MWM07] MOON J. T., WALTER B., MARSCHNER S. R.: Rendering discrete random media using precomputed scattering solutions. In *Proc. EGSR* (June 2007), Eurographics Association, pp. 231–242. 22
- [NJS*11] NOWROUZEZAHRAI D., JOHNSON J., SELLE A., LACEWELL D., KASCHALK M., JAROSZ W.: A programmable system for artistic volumetric lighting. *ACM Transactions on Graphics (Proceedings of SIGGRAPH) 30*, 4 (August 2011), 29:1–29:8. 19
- [NNDJ12a] NOVÁK J., NOWROUZEZAHRAI D., DACHSBACHER C., JAROSZ W.: Progressive virtual beam lights. *CGF (Proc. EGSR) 31*, 4 (June 2012), 1407–1413. 9, 17
- [NNDJ12b] NOVÁK J., NOWROUZEZAHRAI D., DACHSBACHER C., JAROSZ W.: Virtual ray lights for rendering scenes with participating media. *ACM TOG (Proc. SIGGRAPH) 31*, 4 (July 2012), 60:1–60:11. 9, 17
- [NSJ14] NOVÁK J., SELLE A., JAROSZ W.: Residual ratio tracking for estimating attenuation in participating media. *ACM TOG (Proc. SIGGRAPH Asia) 33*, 6 (Nov. 2014), 179:1–179:11. 7, 9, 11, 12, 13, 22
- [OP11] OU J., PELLACINI F.: LightSlice: matrix slice sampling for the many-lights problem. *ACM TOG (Proc. SIGGRAPH Asia) 30*, 6 (Dec. 2011), 179:1–179:8. 17
- [Pan17] PANTALEONI J.: Charted Metropolis light transport. *ACM TOG (Proc. SIGGRAPH) 36*, 4 (2017), 75:1–75:15. 20
- [PAS03] PREMOŽE S., ASHIKHMIN M., SHIRLEY P.: Path integration for light transport in volumes. In *Proc. EGSR* (Aire-la-Ville, Switzerland, 2003), Eurographics Association, pp. 52–63. 20
- [PAT*04] PREMOŽE S., ASHIKHMIN M., TESSENDORF J., RAMAMOORTHY R., NAYAR S.: Practical rendering of multiple scattering effects in participating media. In *Proc. EGSR* (Aire-la-Ville, Switzerland, 2004), Eurographics Association, pp. 363–374. 20
- [Peg16] PEGORARO V.: *Handbook of Digital Image Synthesis: Scientific Foundations of Rendering*. CRC Press, 2016. 2
- [PH89] PERLIN K. H., HOFFERT E. M.: Hypertexture. *Computer Graphics (Proc. SIGGRAPH) 23*, 3 (July 1989), 253–262. 5
- [PKK00] PAULY M., KOLLIG T., KELLER A.: Metropolis light transport for participating media. In *Proc. Eurographics Workshop on Rendering Techniques* (London, UK, 2000), Springer-Verlag, pp. 11–22. 3, 10, 20
- [RAMN12] RAMAMOORTHY R., ANDERSON J., MEYER M., NOWROUZEZAHRAI D.: A theory of Monte Carlo visibility sampling. *ACM TOG 31*, 5 (2012). 10

- [Ray71] RAYLEIGH J. W. S. L.: On the scattering of light by small particles. *Philosophical Magazine* 64 (1871), 447–454. 2
- [RBA09] RADZISZEWSKI M., BORYCZKO K., ALDA W.: An improved technique for full spectral rendering. *Journal of WSCG* 17 (2009), 9–16. 15
- [RSK08] RAAB M., SEIBERT D., KELLER A.: Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer, 2008, pp. 591–606. 6, 7, 16, 17, 20
- [SBB*99] SUTTON T. M., BROWN F. B., BISCHOFF F. G., MACMILLAN D. B., ELLIS C. L., WARD J. T., BALLINGER C. T., KELLY D. J., SCHINDLER L.: *The physical models and statistical procedures used in the RACER Monte Carlo code*. Tech. Rep. KAPL-4840, Knolls Atomic Power Laboratory, Niskayuna, NY, USA, July 1999. 5
- [SG69] SPANIER J., GELBARD E. M.: *Monte Carlo principles and neutron transport problems*. Addison-Wesley series in computer science and information processing. Addison-Wesley Pub. Co., 1969. 8, 10, 11
- [SHHD17] SCHÜSSLER V., HEITZ E., HANIKA J., DACHSBACHER C.: Microfacet-based normal mapping for robust Monte Carlo path tracing. *ACM TOG (Proc. SIGGRAPH Asia)* 36, 6 (Nov. 2017), 205:1–205:12. 20
- [SHZD17] SIMON F., HANIKA J., ZIRR T., DACHSBACHER C.: Line integration for rendering heterogeneous emissive volumes. *CGF (Proc. EGSR)* 36, 4 (June 2017). 21, 22
- [SJJ12] SCHWARZHAUPT J., JENSEN H. W., JAROSZ W.: Practical Hessian-based error control for irradiance caching. *ACM TOG (Proc. SIGGRAPH Asia)* 31, 6 (Nov. 2012). 20
- [SKGM*17] SZIRMAY-KALOS L., GEORGIEV I., MAGDICS M., MOLNÁR B., LÉGRÁDY D.: Unbiased estimators to render procedurally generated inhomogeneous participating media. *CGF (Proc. Eurographics)* 36, 2 (2017). 7, 8, 12
- [SKTM11] SZIRMAY-KALOS L., TÓTH B., MAGDICS M.: Free path sampling in high resolution inhomogeneous participating media. *CGF* 30, 1 (2011), 85–97. 11, 22
- [Sku68] SKULLERUD H. R.: 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. 6
- [Spa66] SPANIER J.: Two pairs of families of estimators for transport problems. *SIAM Journal on Applied Mathematics* 14, 4 (July 1966), 702–713. 9, 10, 11
- [SSH73] SHAMSUNDAR N., SPARROW E. M., HEINISCH R. P.: Monte Carlo radiation solutions—effect of energy partitioning and number of rays. *International Journal of Heat and Mass Transfer* 16, 3 (1973), 690–694. 8
- [Ste66] STEEN N. M.: *A simple method to improve the efficiency of the Σ_a/Σ_t estimator in certain Monte Carlo programs*. Tech. Rep. WAPD-TM-609, Bettis Atomic Power Lab., Pittsburgh, PA, USA, Jan. 1966. 8
- [SVLL10] SALVI M., VIDIMCE K., LAURITZEN A., LEFOHN A.: Adaptive volumetric shadow maps. *CGF (Proc. EGSR)* 29, 4 (2010), 1289–1296. 17
- [SZLG10] SUN X., ZHOU K., LIN S., GUO B.: Line space gathering for single scattering in large scenes. *ACM TOG (Proc. SIGGRAPH)* 29, 4 (July 2010), 54:1–54:8. 18
- [Tes87] TESSENDORF J.: Radiative transfer as a sum over paths. *Phys. Rev. A* 35 (Jan. 1987), 872–878. 20
- [Tes09] TESSENDORF J.: Numerical integration of the Feynman path integral for radiative transport. In *International Conference on Mathematics, Computational Methods & Reactor Physics* (2009). 20
- [UFK13] UREÑA C., FAJARDO M., KING A.: An area-preserving parametrization for spherical rectangles. *CGF (Proc. EGSR)* 32, 4 (2013), 59–66. 14
- [Vea97] VEACH E.: *Robust Monte Carlo methods for light transport simulation*. PhD thesis, Stanford University, Stanford, CA, USA, Dec. 1997. 3, 4, 7, 20
- [VH13] VILLEMEN R., HERY C.: Practical illumination from flames. *Journal of Computer Graphics Techniques* 2, 2 (2013). 21
- [VK16] VORBA J., KŘIVÁNEK J.: Adjoint-driven Russian roulette and splitting in light transport simulation. *ACM TOG* 35, 4 (jul 2016). 22
- [VKv*14] VORBA J., KARLÍK O., ŠÍK M., RITSCHER T., KŘIVÁNEK J.: On-line learning of parametric mixture models for light transport simulation. *ACM TOG (Proc. SIGGRAPH)* 33, 4 (Aug. 2014). 22
- [vN51] VON NEUMANN J.: Various techniques used in connection with random digits. *Journal of Research of the National Bureau of Standards, Appl. Math. Series* 12 (1951), 36–38. 6
- [WABG06] WALTER B., ARBREE A., BALA K., GREENBERG D. P.: Multidimensional lightcuts. *ACM TOG (Proc. SIGGRAPH)* 25, 3 (July 2006), 1081–1088. 16, 17
- [WFA*05] WALTER B., FERNANDEZ S., ARBREE A., BALA K., DONIKIAN M., GREENBERG D. P.: Lightcuts: a scalable approach to illumination. *ACM TOG (Proc. SIGGRAPH)* 24, 3 (July 2005), 1098–1107. 17
- [WHD17] WEBER P., HANIKA J., DACHSBACHER C.: Multiple vertex next event estimation for lighting in dense, forward-scattering media. *CGF (Proc. Eurographics)* 36 (Apr. 2017), 21–30. 20
- [Wil78] WILLIAMS L.: Casting curved shadows on curved surfaces. *Computer Graphics (Proc. SIGGRAPH)* 12, 3 (Aug. 1978), 270–274. 17
- [WKB12] WALTER B., KHUNGURN P., BALA K.: Bidirectional lightcuts. *ACM TOG (Proc. SIGGRAPH)* 31, 4 (July 2012), 59:1–59:11. 17
- [WMHL65] WOODCOCK E. R., MURPHY T., HEMMINGS P. J., LONGWORTH T. C.: 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* (1965), Argonne National Laboratory. 6
- [WND*14] WILKIE A., NAWAZ S., DROSKE M., WEIDLICH A., HANIKA J.: Hero wavelength spectral sampling. *CGF (Proc. EGSR)* 33, 4 (June 2014), 123–131. 15
- [WRC88] WARD G. J., RUBINSTEIN F. M., CLEAR R. D.: A ray tracing solution for diffuse interreflection. *Computer Graphics (Proc. SIGGRAPH)* 22, 4 (June 1988), 85–92. 20
- [Wre16] WRENNINGE M.: Efficient rendering of volumetric motion blur using temporally unstructured volumes. *Journal of Computer Graphics Techniques (JCGT)* 5, 1 (January 2016), 1–34. 22
- [WZC*10] WRENNINGE M., ZAFAR N. B., CLIFFORD J., GRAHAM G., PENNEY D., KONTKANEN J., TESSENDORF J., CLINTON A.: SIGGRAPH 2010 course notes: Volumetric methods in visual effects. In *SIGGRAPH Courses* (July 2010). 21
- [YIC*10] YUE Y., IWASAKI K., CHEN B.-Y., DOBASHI Y., NISHITA T.: Unbiased, adaptive stochastic sampling for rendering inhomogeneous participating media. *ACM TOG (Proc. SIGGRAPH Asia)* 29, 6 (Dec. 2010), 177:1–177:8. 22
- [YIC*11] YUE Y., IWASAKI K., CHEN B., DOBASHI Y., NISHITA T.: Toward optimal space partitioning for unbiased, adaptive free path sampling of inhomogeneous participating media. *CGF (Proc. Pacific Graphics)* 30, 7 (2011), 1911–1919. 22
- [YY17] YUKSEL C., YUKSEL C.: Lighting grid hierarchy for self-illuminating explosions. *ACM TOG (Proc. SIGGRAPH)* 36, 4 (2017), 110:1–110:10. 17
- [ZCB61] ZERBY C. D., CURTIS R. B., BERTINI H. W.: *The relativistic doppler problem*. Tech. Rep. ORNL-61-7-20, Oak Ridge National Laboratory, Oak Ridge, TN, USA, July 1961. 6