

Joint Importance Sampling of Low-Order Volumetric Scattering

Iliyan Georgiev^{1,2} Jaroslav Křivánek³ Toshiya Hachisuka⁴ Derek Nowrouzezahrai^{1,5} Wojciech Jarosz¹

¹Disney Research Zürich ²Saarland University ³Charles University, Prague ⁴Aarhus University ⁵Université de Montréal

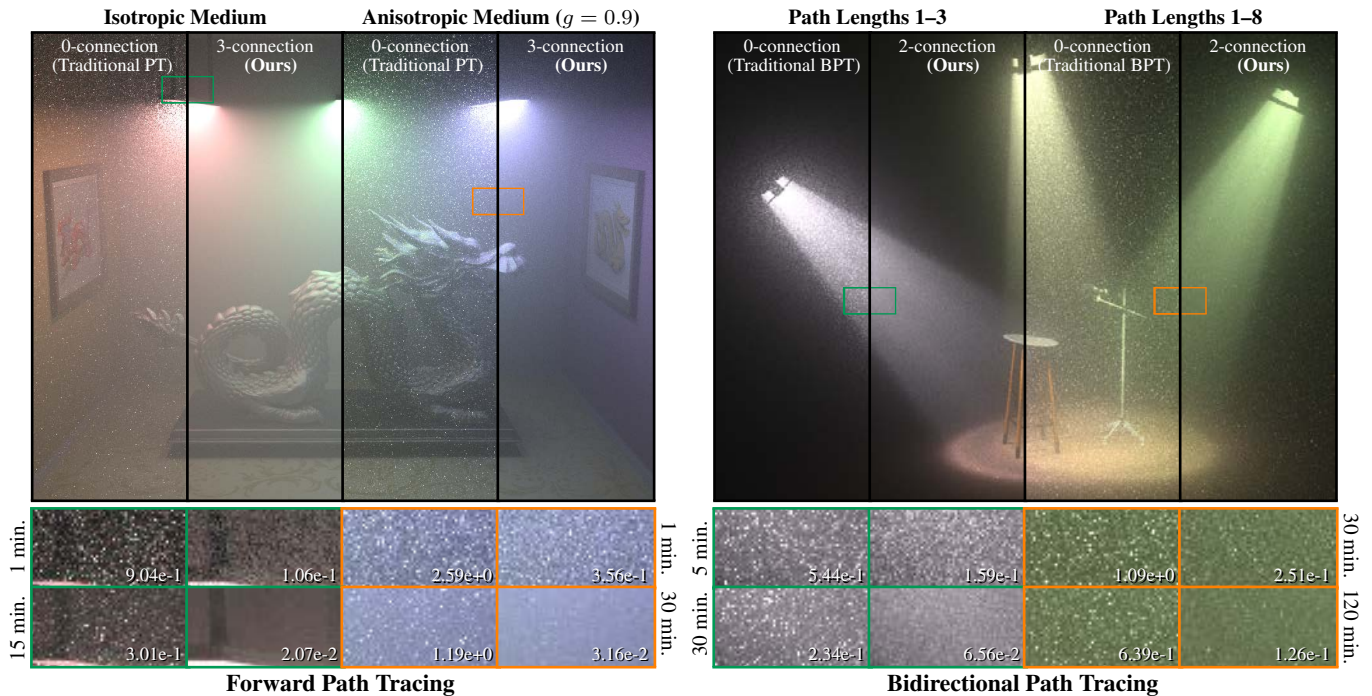


Figure 1: We present novel importance sampling techniques for constructing paths in participating media and apply them to unidirectional and bidirectional path tracing. By generalizing traditional shadow connections (0-connection) to longer, importance sampled connection subpaths (2- and 3-connection) we obtain $5\times$ to $37\times$ reduction in RMS error (see zoom-ins), corresponding to $25\times$ to $1444\times$ reduction in render time.

Abstract

Central to all Monte Carlo-based rendering algorithms is the construction of light transport paths from the light sources to the eye. Existing rendering approaches sample path vertices *incrementally* when constructing these light transport paths. The resulting probability density is thus a product of the *conditional* densities of each local sampling step, constructed without explicit control over the form of the final *joint* distribution of the complete path. We analyze why current incremental construction schemes often lead to high variance in the presence of participating media, and reveal that such approaches are an unnecessary legacy inherited from traditional surface-based rendering algorithms. We devise *joint importance sampling* of path vertices in participating media to construct paths that explicitly account for the product of all scattering and geometry terms along a sequence of vertices instead of just locally at a single vertex. This leads to a number of practical importance sampling routines to explicitly construct single- and double-scattering subpaths in anisotropically-scattering media. We demonstrate the benefit of our new sampling techniques, integrating them into several path-based rendering algorithms such as path tracing, bidirectional path tracing, and many-light methods. We also use our sampling routines to generalize deterministic shadow connections to *connection subpaths* consisting of two or three random decisions, to efficiently simulate higher-order multiple scattering. Our algorithms significantly reduce noise and increase performance in renderings with both isotropic and highly anisotropic, low-order scattering.

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

Keywords: global illumination, participating media, importance sampling, path tracing, bidirectional path tracing, many-light rendering methods

Links: [DL](#) [PDF](#) [WEB](#)



©ACM, 2013. This is the authors' version of the work (revision 2, last updated: 5 Jan 2017). It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in ACM Transactions on Graphics, 32, 6, November 2013. doi.acm.org/10.1145/2508363.2508411

1 Introduction

Faithfully simulating light transport in participating media is a challenging problem with important applications across many diverse fields, such as medical imaging and nuclear physics, as well as computer graphics. Over the last three decades, the computer graphics community has adapted and developed a sizable arsenal of methods for rendering participating media. Though no single method outperforms all others in every situation, we now have efficient approximations for rendering specific scattering scenarios, including: general Monte Carlo integration [Kajiya 1986; Pauly et al. 2000], (semi-)analytic approximations [Sun et al. 2005; Pegoraro et al. 2009; Walter et al. 2009], the diffusion approximation [Stam 1995; Jensen et al. 2001; D’Eon and Irving 2011], and the Feynman path integral approximation [Tessendorf 1987; Premože et al. 2003; Premože et al. 2004]. Though these methods are each well-suited for specific scattering scenarios, a common limitation of all these methods is their inefficiency or inaccuracy in handling participating media with highly anisotropic scattering and moderate albedo. We fill this important gap and improve general Monte Carlo integration techniques such as path tracing, which are typically slow for this class of media, by developing a set of importance sampling techniques that significantly improve rendering performance.

Path-sampling approaches traditionally construct light transport paths by sequentially sampling local scattering events. Unfortunately, this incremental construction sacrifices a global view of the path’s contribution, and can lead to significant variance in the final estimator. In fact, though not widely known in computer graphics, it has long been known in related fields [Kalos 1963] that methods like volumetric path tracing with explicit shadow connections suffer from *infinite* variance when participating media surrounds light sources.

In the face of these challenges, we propose *joint path importance sampling* to construct paths in participating media while accounting for the product of anisotropic phase functions and geometric terms across sequences of path vertices. Though importance sampling the BRDF, the lighting, or their product has been investigated for single-bounce surface illumination, generalizations to multiple bounces (a sequence of vertices) are challenging due to the curse of dimensionality. Extensions to participating media additionally require sampling a propagation distance, which increases the dimensionality even further. We instead exploit this extra dimension as a degree of freedom, along with symmetries uniquely present in media light transport, to make joint importance sampling of path vertices practical. We develop several novel path importance sampling techniques for participating media, introducing *connection subpaths* to join a light and the eye with two additional intermediate vertices while accounting for the product of phase functions and geometry terms.

Contributions and overview. In summary, we present the following theoretical and practical contributions:

- We express previous sampling methods for participating media as general path sampling and connection strategies in a unified framework under the path integral formulation (Section 3).
- We derive analytic formulas for *exact* importance sampling of the product of geometry terms for a sequence of vertices (three path segments) joining two given subpaths (Section 5).
- We exploit unique symmetries in volumetric light transport to build compact tabulations for importance sampling the product of geometry and anisotropic phase function terms across a sequence of vertices joining two given subpaths (Section 6).
- We demonstrate the versatility of our new sampling techniques, applying them to several rendering algorithms—such as path tracing, bidirectional path tracing, and many-light rendering (Section 7)—significantly improving rendering quality by more accurately modeling transport along longer light subpaths.

2 Related Work

Central to all Monte Carlo-based rendering algorithms is the task of sampling light transport paths that connect light sources to the eye. Veach [1997] formalized this idea with the path integral formulation of light transport for surfaces, and Pauly et al. [2000] extended it to participating media. The path integrand can be factorized into the product of several terms, such as the geometry, transmittance, and scattering terms. For Monte Carlo integration, one ideally wishes to sample paths with a probability density function (PDF) proportional to that integrand. In prior work, subpaths are typically constructed incrementally, vertex-by-vertex, followed by a deterministic subpath connection. This yields path PDFs that do not account for some high-variance terms. We introduce a new set of importance sampling techniques for constructing *connection subpaths* which join light and eye subpaths with two additional intermediate vertices (three segments). In contrast to deterministic connections, our approach accounts for the product of several geometry and scattering terms along the connection subpath, leading to significant variance reduction.

Surface rendering. Various techniques can be employed atop basic importance sampling. Product sampling [Burke et al. 2005; Clarberg et al. 2005; Clarberg and Akenine-Möller 2008; Tsai et al. 2008; Jarosz et al. 2009] can account for the product of BRDF and lighting, but only for a single bounce of illumination, and with high memory and computation costs. Multiple importance sampling (MIS) [Veach and Guibas 1995] can combine several PDFs—such as for BRDF or light importance sampling [Pharr and Humphreys 2010]—while preserving the strengths of each individual technique. Bidirectional path tracing employs MIS to combine sampling techniques based on connecting two independent subpaths, each constructed incrementally. However, the resulting combined joint path PDF remains a simple linear combination of the PDFs of the individual techniques, as opposed to the product of all path integrand terms. While Metropolis light transport and its variants [Veach and Guibas 1997; Pauly et al. 2000; Cline et al. 2005; Jakob and Marschner 2012b] do sample light paths with density proportional to the full integrand, they usually suffer from sample correlation and are in most practical cases outperformed by classic Monte Carlo approaches based on independent samples. Our sampling techniques are also compatible with MIS and, since they directly importance sample according to transport along a sequence of path vertices, improve efficiency over the combination of existing sampling techniques.

Volume rendering. Incremental, vertex-by-vertex path construction in participating media traditionally proceeds by sampling the direction of the next path segment proportionally to the phase function, and choosing the propagation distance along the segment proportionally to transmittance [Woodcock et al. 1965; Raab et al. 2008]. Kulla and Fajardo [2012] observed that the geometry term is often responsible for more variance than the transmittance term and proposed equiangular sampling to choose the propagation distance proportionally to the geometry term when computing single scattering in isotropic media. Recent extensions to density estimation [Jensen 1996; Jensen and Christensen 1998] have proposed using lines, instead of vertices, as the eye [Jarosz et al. 2008] and light [Sun et al. 2010; Jarosz et al. 2011] path building block, leading to significant improvements. Inspired by this line of work, Novák et al. [2012] analyzed many-light rendering in anisotropically-scattering media and showed that such line-based path construction can probably reduce singularities in many-light rendering (see Křivánek et al. [2012] for a comprehensive survey of many-light techniques). We show that the methods of Kulla and Fajardo [2012] and Novák et al. [2012] can be interpreted as general path sampling techniques and employed for path construction in bidirectional path sampling algorithms using both *segments* (lines) and *vertices*. We unify and formalize these approaches with joint importance sampling, and propose new joint path sampling techniques for more substantial variance reduction.

Neutron transport & radiative transfer. Many of these sophisticated participating media rendering techniques have been adapted (and sometimes unknowingly reinvented) from the neutron transport and radiative transfer literatures, which have a much longer history of investigating these problems [Chandrasekar 1960; Spanier and Gelbard 1969]. Density estimation in volumetric photon mapping uses a *collision estimator*, and the path segment improvements [Jarosz et al. 2011; Novák et al. 2012] share similarity with, and extend upon, *track length estimators* [Spanier 1966]. Volumetric path tracing with shadow connections corresponds to the *uncollided flux* (UCF) estimator used in neutron transport simulations since the early 1960s.

Kalos [1963] proved that the UCF estimator has infinite variance (due to a $1/r^2$ singularity in the geometry term) and, even worse, its use leads to an abysmal $1/\sqrt[3]{N}$ convergence rate [Kalos 1963; Dubi et al. 1982]. To remedy this problem, he proposed the *once collided flux* (OCF) estimator—which samples an extra path vertex when forming the connection—and proved that this strategy reduces the order of the singularity to $1/r$ and recovers the more favorable $1/\sqrt{N}$ convergence rate. Since this seminal work, many researchers have suggested improvements or simplifications to the OCF [Steinberg and Kalos 1971; Kalli and Cashwell 1977; Rief et al. 1984; Raab and Beikert 1999]. Notable examples include the use of an equiangular approach (similar to Kulla and Fajardo’s [2012]) to sample the extra vertex [Kalli and Cashwell 1977; Rief et al. 1984]. Much like Kalos’ OCF estimator and its extensions, we sample additional path vertices to reduce singularities; however, we go a step further by considering not just one but two vertices, while also accounting for all geometry terms and anisotropic phase functions along the connecting subpath.

3 Background

To position our contribution, we present a unified formulation of several existing rendering methods as different sampling techniques under the path integral framework [Veach 1997]. This will allow us to discuss interrelations, identify respective weaknesses, and motivate our new approaches with respect to existing work. We first briefly review the extended path integral framework for scenes with participating media [Jakob and Marschner 2012a], before discussing and comparing current techniques for sampling and connecting paths.

3.1 Path Integral Framework


Image pixel intensity I is computed as an integral over the space Ω of light transport paths: $I = \int_{\Omega} f(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}})$. Here, $\bar{\mathbf{x}} = \mathbf{x}_0, \dots, \mathbf{x}_k$ is a length- k path with $k \geq 1$ segments and $k + 1$ vertices, where the first vertex \mathbf{x}_0 is on a light source, the last vertex \mathbf{x}_k is on the eye lens, and $\mathbf{x}_1, \dots, \mathbf{x}_{k-1}$ are scattering points on surfaces and in media. The differential measure on the path space $d\mu(\bar{\mathbf{x}})$ corresponds to area integration for surface vertices and volume integration for medium vertices. The measurement contribution function $f(\bar{\mathbf{x}})$ is the product of the geometry throughput $G(\bar{\mathbf{x}})$, scattering throughput $\rho(\bar{\mathbf{x}})$, path transmittance $T_r(\bar{\mathbf{x}})$, and path visibility $V(\bar{\mathbf{x}})$:

$$f(\bar{\mathbf{x}}) = G(\bar{\mathbf{x}})\rho(\bar{\mathbf{x}})T_r(\bar{\mathbf{x}})V(\bar{\mathbf{x}}), \quad (1)$$

where

$$\begin{aligned} G(\bar{\mathbf{x}}) &= \prod_{i=0}^{k-1} G(\mathbf{x}_i, \mathbf{x}_{i+1}) & \rho(\bar{\mathbf{x}}) &= \prod_{i=0}^k \rho(\mathbf{x}_i) \\ T_r(\bar{\mathbf{x}}) &= \prod_{i=0}^{k-1} T_r(\mathbf{x}_i, \mathbf{x}_{i+1}) & V(\bar{\mathbf{x}}) &= \prod_{i=0}^{k-1} V(\mathbf{x}_i, \mathbf{x}_{i+1}) \end{aligned}$$

We illustrate these terms in the diagram below.



We overload G , T_r and V to apply to both entire paths and individual segments (or vertices, for ρ). The geometry term for a segment $\mathbf{x}\mathbf{y}$ is given by $G(\mathbf{x}, \mathbf{y}) = \frac{D(\mathbf{x} \rightarrow \mathbf{y})D(\mathbf{y} \rightarrow \mathbf{x})}{\|\mathbf{x} - \mathbf{y}\|^2}$, where $D(\mathbf{x} \rightarrow \mathbf{y}) = |n_{\mathbf{x}} \cdot \omega_{\mathbf{x}\mathbf{y}}|$ if \mathbf{x} is on a surface and $D(\mathbf{x} \rightarrow \mathbf{y}) = 1$ if \mathbf{x} is in a medium, and likewise for $D(\mathbf{y} \rightarrow \mathbf{x})$. The transmittance of segment $\mathbf{x}\mathbf{y}$ is given by $T_r(\mathbf{x}, \mathbf{y}) = \exp\left(-\int_0^{\|\mathbf{x}-\mathbf{y}\|} \sigma_t(\mathbf{x} + t\omega_{\mathbf{x}\mathbf{y}}) dt\right)$ and $V(\mathbf{x}, \mathbf{y})$ is its binary visibility. We define the scattering term $\rho(\mathbf{x}_i)$ as

$$\rho(\mathbf{x}_i) = \begin{cases} L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1) & \text{if } i = 0, \\ W_e(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) & \text{if } i = k, \\ \rho_s(\mathbf{x}_{i-1} \rightarrow \mathbf{x}_i \rightarrow \mathbf{x}_{i+1}) & \text{if } \mathbf{x}_i \text{ on surface, and} \\ \rho_p(\mathbf{x}_{i-1} \rightarrow \mathbf{x}_i \rightarrow \mathbf{x}_{i+1})\sigma_s(\mathbf{x}_i) & \text{if } \mathbf{x}_i \text{ in medium,} \end{cases} \quad (2)$$

where ρ_s and ρ_p are the BSDF and phase function, and we treat the emission L_e and sensor importance W_e as special scattering events.

3.2 Path Probability Density for Existing Techniques

The variance of a Monte Carlo path integral estimator (the amount of image noise) depends on the probability density for sampling paths. The path probability density $p(\bar{\mathbf{x}})$ is given by the joint density of the individual vertices, expressed with respect to the path space measure μ , i.e. $p(\bar{\mathbf{x}}) = p(\mathbf{x}_0, \dots, \mathbf{x}_k)$. This joint PDF is fully determined by the *path sampling technique* used to generate the path. For traditional techniques based on incremental, vertex-by-vertex path construction, the joint PDF is the product of the conditional PDFs $p(\mathbf{x}_i | \text{vertices sampled before } \mathbf{x}_i)$ for sampling each path vertex. Ideally the joint PDF should be chosen to be exactly proportional to the full path contribution function $f(\bar{\mathbf{x}})$ in Equation (1), however this is not feasible in practice and so only certain terms of $f(\bar{\mathbf{x}})$ are importance sampled. The resulting estimator’s variance depends directly on the variability of the ignored (and/or approximated) terms.

We will analyze the different joint path PDFs induced by existing unbiased participating media rendering algorithms. We will motivate our new sampling techniques by identifying terms in the contribution function that are *not* importance sampled by each existing technique’s joint PDF, and interpret all techniques as general path sampling and connection strategies. Since we focus on media, and in the interest of simplifying our exposition, we assume from this point forward that all scattering vertices correspond to volume scattering events (i.e. that $G(\mathbf{x}, \mathbf{y}) = 1/\|\mathbf{x} - \mathbf{y}\|^2$). We use an example length-3 path \mathbf{abcd} depicted in Figure 2 throughout our discussions.

Measures and measure conversion. We express all directional PDFs $p(\omega)$ w.r.t. the solid angle measure, all distance PDFs $p(t)$ w.r.t. the Euclidean length on \mathbb{R}^1 and all medium vertex PDFs $p(\mathbf{x})$ w.r.t. the Euclidean volume on \mathbb{R}^3 . Converting from the solid angle \times length product measure to the volume measure requires multiplication by a corresponding geometry term G .

Unidirectional transmittance-based path sampling. A simple method to construct paths is to first sample a vertex \mathbf{d} on the eye sensor from a PDF $p(\mathbf{d})$ proportional to importance, and to incrementally add vertices until we randomly hit a light source, yielding a simple version of *path tracing* [Kajiya 1986] (Figure 2a). Given vertex \mathbf{c} in the medium with incident direction $\omega_{\mathbf{dc}}$, the next path vertex \mathbf{b} is sampled as follows: first, a direction $\omega_{\mathbf{cb}}$ is sampled proportionally to the phase function at \mathbf{c} , i.e. $p(\omega_{\mathbf{cb}} | \omega_{\mathbf{dc}}, \mathbf{c}) \propto \rho(\mathbf{c})$. The propagation distance $t_{\mathbf{cb}}$ along $\omega_{\mathbf{cb}}$ is then sampled proportionally to transmittance¹ with PDF $p(t_{\mathbf{cb}} | \mathbf{c}, \omega_{\mathbf{cb}})$. The direction $\omega_{\mathbf{cb}}$

¹Surface scattering occurs if $t_{\mathbf{cb}}$ is beyond the nearest surface intersection along the ray $(\mathbf{c}, \omega_{\mathbf{cb}})$. We disregard this technicality here for clarity.




Figure 2: A comparison of various existing techniques for sampling light transport paths in a medium. The pink connection vertices are sampled from line distributions that are conditioned on the endpoints and/or end-lines of the light and eye subpaths.

and distance t_{cb} determine the next vertex \mathbf{b} , with conditional PDF:

$$p(\mathbf{b} | \omega_{dc}, \mathbf{c}) = p(\omega_{cb} | \omega_{dc}, \mathbf{c}) p(t_{cb} | \mathbf{c}, \omega_{cb}) G(\mathbf{b}, \mathbf{c}), \quad (3)$$

where $G(\mathbf{b}, \mathbf{c})$ is needed to convert the PDF to the volume measure.

The joint path PDF of this purely unidirectional sampling procedure is the product of $p(\mathbf{d})$ (proportional to emitted importance) with the PDF in Equation (3) for every other vertex (Figure 2a). This joint PDF is proportional to all the terms in the path contribution of Equation (1) except the emitted radiance $\rho(\mathbf{a}) = L_e(\mathbf{a} \rightarrow \mathbf{b})$. Unfortunately, in many practical cases, this approach yields estimators with extremely high variance due to the negligible probability of hitting small light sources. Unidirectional sampling can alternatively be used to build paths starting from light sources: the resulting *light tracing* algorithm [Dutré et al. 2006] usually has even higher variance because of the small size of the eye sensor.

In practice, both path and light tracing use next event estimation with explicit shadow connections to lights or sensors. We consider these as special cases of bidirectional path sampling, discussed below.

Bidirectional transmittance-based path sampling. To address the high variance of unidirectional sampling, both emitted radiance and sensor importance must be included in the path PDF. This can be achieved with bidirectional sampling, illustrated in Figure 2b. Unidirectional sampling is first used to construct one subpath with s vertices from a light source and another one with t vertices from the eye. The ends of these independent subpaths are then deterministically connected with a segment, completing a full path. Next event estimation for path and light tracing corresponds to such a *vertex-vertex* connection technique, with $s = 1$ or $t = 1$ respectively.

Vertex-vertex connections can lead to high variance since none of the terms for the connecting segment are importance sampled, as it is a *0-random-decision* subpath connecting the light and eye subpaths (see Figure 2b). In fact, the variance is infinite due to the geometry term $G(\mathbf{b}, \mathbf{c})$ diverging as \mathbf{b} and \mathbf{c} approach each other [Kalos 1963]. To resolve this, *bidirectional path tracing* exploits the fact that any full path can be generated using a number of techniques, each corresponding to a unique combination of different lengths for the light and eye subpaths. The MIS-weighted combination of these estimators yields finite and often relatively low variance.

Equiangular sampling. Several researchers [Kalli and Cashwell 1977; Rief et al. 1984; Kulla and Fajardo 2012] have proposed importance sampling the geometry term between a vertex and positions on a ray. Kulla and Fajardo used this technique for rendering single scattering, and we generalize it to a *vertex-segment* connection technique in bidirectional sampling (Figure 2c).

Given the end vertices \mathbf{b} and \mathbf{d} on the light and eye subpaths, along with a direction ω_{dc} from \mathbf{d} , the equiangular technique constructs a *1-random-decision connection subpath* by sampling the propagation distance t_{dc} along ω_{dc} to create vertex \mathbf{c} with a PDF proportional to the inverse squared length of the connecting edge \mathbf{bc} ,

$$p(t_{dc} | \mathbf{b}, \mathbf{d}, \omega_{dc}) \propto G(\mathbf{b}, \mathbf{c}) = 1/\|\mathbf{b} - \mathbf{c}\|^2. \quad (4)$$

Equiangular distance sampling often results in lower variance than transmittance sampling. However, this technique does not importance sample any other terms in the connection, namely: $T_r(\mathbf{b}, \mathbf{c})$, $V(\mathbf{b}, \mathbf{c})$, $T_r(\mathbf{d}, \mathbf{c})$, $V(\mathbf{d}, \mathbf{c})$, $\rho(\mathbf{b})$, and $\rho(\mathbf{c})$. Therefore, in some cases, transmittance sampling may be preferred, so Kulla and Fajardo combine both techniques using MIS.

Kulla and Fajardo [2012] originally developed this PDF to “cancel out” the weak singularity in $G(\mathbf{b}, \mathbf{c})$; however, as we will show in Section 5.1, the singularity remains “hidden” in the orientation of ω_{dc} , which we eliminate when importance sampling ω_{dc} .

Virtual ray lights. VRLs [Novák et al. 2012] are a many-light method that uses a one-sample Monte Carlo estimator to calculate the total energy transfer between a “virtual ray light” and an eye ray. We reinterpret this method as a general bidirectional *segment-segment* connection technique (Figure 2d): given the end vertex \mathbf{a} of a light subpath along with direction ω_{ab} (i.e. the virtual ray light), and an end vertex \mathbf{d} on an eye subpath along with direction ω_{dc} , this *2-random-decision* technique samples the propagation distances t_{ab} and t_{dc} from a joint conditional distribution,

$$p(t_{ab}, t_{dc} | \mathbf{a}, \omega_{ab}, \mathbf{d}, \omega_{dc}) \propto G(\mathbf{b}, \mathbf{c}) \rho(\mathbf{b}) \rho(\mathbf{c}) = \frac{\rho(\mathbf{b}) \rho(\mathbf{c})}{\|\mathbf{b} - \mathbf{c}\|^2}. \quad (5)$$

The approximate proportionality stems from various simplifications in the VRL approach, which we improve upon in Section 6.3.

In Figure 2 we observe that the only difference between vertex-vertex connection, equiangular sampling, and VRLs is the PDF for sampling the distances t_{ab} and t_{dc} . This PDF is either proportional to transmittance, or to the scattering and/or geometry term(s) involved in the connection. Depending on the geometry configuration and media properties, either technique may lead to a lower-variance estimator, so Novák et al. [2012] combine them using MIS.

Figure 2 also shows that each of these techniques can be viewed as creating connections between light and eye subpaths with a varying number of random decisions (0, 1, or 2). In the next sections, we go a step further and develop a practical method for creating connection subpaths involving *3 random decisions*, where the extra decision is direction sampling. We then show how to extend unidirectional and bidirectional path tracing with such connections in Section 7, leading to significant variance reduction.

4 Joint Path Vertex Sampling

We consider a configuration where a light subpath $\bar{x}_l \mathbf{a}$ and an eye subpath $\mathbf{d} \bar{x}_e$ are given. Our aim is to construct a full path $\bar{x}_l \mathbf{a} \mathbf{b} \mathbf{c} \mathbf{d} \bar{x}_e$ by connecting the endpoints \mathbf{a} and \mathbf{d} via a *3-random-decision connection subpath* with two new vertices \mathbf{b} and \mathbf{c} , as shown in Figure 3a. The subpaths \bar{x}_l and \bar{x}_e can have arbitrary lengths; zero lengths correspond to the case where \mathbf{a} is on a light source and \mathbf{d} is on the eye lens. We also assume that a direction ω_{dc} at vertex \mathbf{d} is given (sampled as part of the eye subpath). The input to our sampling techniques is thus a vertex \mathbf{a} with incident direction ω_{la} and a vertex \mathbf{d} with outgoing direction ω_{dc} .

Since the new vertex \mathbf{c} must reside on the ray $(\mathbf{d}, \omega_{dc})$, our task reduces to sampling a distance t_{dc} from \mathbf{d} along ω_{dc} and another vertex \mathbf{b} . The PDF of the resulting full path is then

$$p(\bar{x}) = p(\bar{x}_l, \mathbf{a}, \omega_{dc}, \mathbf{d}, \bar{x}_e) p(\mathbf{b}, t_{dc} | \omega_{la}, \mathbf{a}, \omega_{dc}, \mathbf{d}) G(\mathbf{c}, \mathbf{d}), \quad (6)$$

where the geometry term $G(\mathbf{c}, \mathbf{d})$ again arises due to measure conversion. Note that $p(\bar{x}_l, \mathbf{a}, \omega_{dc}, \mathbf{d}, \bar{x}_e)$ is given by the chosen rendering algorithm, not by our sampling techniques. How the light and camera subpaths are created is orthogonal to our problem; our sampling techniques only require that the input configuration $(\omega_{la}, \mathbf{a}, \omega_{dc}, \mathbf{d})$ be provided. Since all our PDFs are conditioned on this input configuration, we will use the following shorthand notation for it:

$$\Xi \equiv \omega_{la}, \mathbf{a}, \omega_{dc}, \mathbf{d}. \quad (7)$$

Ideally, we would like to define $p(\mathbf{b}, t_{dc} | \Xi) G(\mathbf{c}, \mathbf{d})$ (i.e., the terms that our techniques introduce) to be proportional to the path throughput $f(\mathbf{abcd})$. Unfortunately, defining such a PDF is not feasible since its normalization constant is the solution of the path integral. Inspired by previous work [Kulla and Fajardo 2012; Novák et al. 2012], we propose to importance sample only the product of the geometry terms and the phase functions, such that

$$p(\mathbf{b}, t_{dc} | \Xi) G(\mathbf{c}, \mathbf{d}) = C_{\Xi} G(\mathbf{abcd}) \rho(\mathbf{abc}), \quad (8)$$

where the normalization factor C_{Ξ} (which we derive in Section 5) depends on the input configuration Ξ , and where $G(\mathbf{abcd}) = G(\mathbf{a}, \mathbf{b}) G(\mathbf{b}, \mathbf{c}) G(\mathbf{c}, \mathbf{d})$ and $\rho(\mathbf{abc}) = \rho(\mathbf{a}) \rho(\mathbf{b}) \rho(\mathbf{c})$. Our contribution is the definition of the joint PDF $p(\mathbf{b}, t_{dc} | \Xi)$ along with its corresponding sampling techniques. Note that $\rho(\mathbf{d})$ does not appear in the above equation, since we assume that both the outgoing direction, ω_{dc} , and the incident direction at \mathbf{d} are given.

4.1 Factorizations of the Joint PDF

Given the input configuration Ξ , there are a number of possible ways to obtain the connection subpath vertices. Each approach corresponds to factorizing the joint PDF in Equation (8) in a different way, which influences the definition of the conditional PDFs for the individual vertices and their corresponding sampling routines. We consider two such possible factorizations, depicted in Figure 3. We focus only on factorization here, and we will later explain how to obtain and sample from each factor of the joint PDF.

Unidirectional factorization. One way to sample \mathbf{b} and \mathbf{c} is to first sample a distance t_{dc} , which effectively samples \mathbf{c} in combination with the given ω_{dc} , then sample a direction ω_{cb} , and finally sample a distance t_{cb} to obtain \mathbf{b} . This process is illustrated in Figure 3b. We call this *unidirectional factorization* since \mathbf{c} and \mathbf{b} are sampled from the same direction along the path starting from \mathbf{d} .

Since we sample \mathbf{b} as a distance and direction from \mathbf{c} , we need to change the measure of the joint PDF in Equation (8) to

$$p(\mathbf{b}, t_{dc} | \Xi) = p(t_{cb}, \omega_{cb}, t_{dc} | \Xi) G(\mathbf{b}, \mathbf{c}). \quad (9)$$

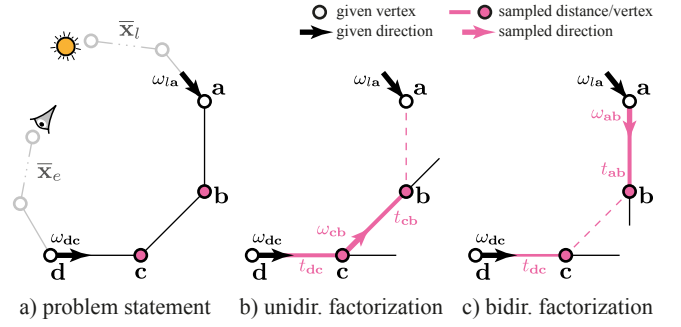


Figure 3: Given the vertices \mathbf{a} and \mathbf{d} , with an incident and an outgoing direction respectively, we aim to sample vertices \mathbf{b} and \mathbf{c} from a joint distribution proportional to the geometry and scattering terms of the resulting connection subpath. We derive the sampling techniques corresponding to two possible factorizations of this PDF.

In order to sample the connection subpath, we factorize the joint PDF into the product of three conditional PDFs, corresponding to first sampling t_{dc} , followed by ω_{cb} , and finally t_{cb} :

$$p(t_{dc}, \omega_{cb}, t_{cb} | \Xi) = p(t_{dc} | \Xi) \quad (\mathcal{U}_3) \quad (10)$$

$$p(\omega_{cb} | t_{dc}, \Xi) \quad (\mathcal{U}_2) \quad (11)$$

$$p(t_{cb} | \omega_{cb}, t_{dc}, \Xi). \quad (\mathcal{U}_1) \quad (12)$$

We introduce a shorthand notation \mathcal{U} above, where the subscripts denote the order in which we will derive the conditional PDFs.

Bidirectional factorization. Another way to sample \mathbf{b} and \mathbf{c} is to first sample a direction ω_{ab} from \mathbf{a} , then a distance t_{ab} , which yields \mathbf{b} , and finally a distance t_{dc} to obtain \mathbf{c} . We call this *bidirectional factorization* since \mathbf{b} and \mathbf{c} are sampled from opposite directions.

In this process, illustrated in Figure 3c, \mathbf{b} is sampled as a distance and direction from \mathbf{a} , and the change of measure of the target PDF is

$$p(\mathbf{b}, t_{dc} | \Xi) = p(t_{ab}, \omega_{ab}, t_{dc} | \Xi) G(\mathbf{a}, \mathbf{b}). \quad (13)$$

The bidirectional factorization samples the connection subpath by factorizing this PDF into the product of three conditional PDFs, corresponding to first sampling ω_{ab} , followed by t_{ab} and t_{dc} :

$$p(\omega_{ab}, t_{ab}, t_{dc} | \Xi) = p(\omega_{ab} | \Xi) \quad (\mathcal{B}_3) \quad (14)$$

$$p(t_{ab} | \omega_{ab}, \Xi) \quad (\mathcal{B}_2) \quad (15)$$

$$p(t_{dc} | \omega_{ab}, t_{ab}, \Xi). \quad (\mathcal{B}_1) \quad (16)$$

We likewise introduce a shorthand \mathcal{B} for these conditional PDFs.

4.2 Conditional Distributions and Marginalization

Given the factorizations defined in Equations (10–12) and (14–16), we now need to derive the individual conditional PDFs. We do this through the process of *marginalization*, which we briefly review in general terms here.

Suppose we have a set of random variables distributed according to a joint distribution $p(x, y, z)$. By definition, any conditional distribution is given by the ratio of the joint distribution and the marginalized joint. For example, we can have

$$p(x | y, z) = \frac{p(x, y, z)}{p(y, z)}. \quad (17)$$

The denominator $p(y, z)$ is given by *marginalizing out* the variable x from the original joint distribution as follows:

$$p(y, z) = \int_{\mathbb{R}} p(x, y, z) dx. \quad (18)$$

In the following two sections, we will derive the conditional PDFs for both factorizations above using these marginalization relations.

5 Analytic Importance Sampling

In this section, we show that for the case of isotropic scattering, we can derive *fully analytic* formulas for the PDFs and the sampling routines for the unidirectional factorization above. In Section 6, we will show how to take advantage of the symmetries in the geometry of the sampling problem to construct low-dimensional PDF tabulations for both factorizations for the general case of anisotropic scattering.

5.1 Unidirectional Factorization PDFs

We first derive analytic expressions for the conditional PDFs in Equations (10–12) for the case of isotropic scattering. Since the phase function is constant, Equation (8) simplifies to²

$$p(\mathbf{b}, t_{dc} | \Xi) G(\mathbf{c}, \mathbf{d}) = C_{\Xi} G(\mathbf{a}, \mathbf{b}) G(\mathbf{b}, \mathbf{c}) G(\mathbf{c}, \mathbf{d}). \quad (19)$$

Using Equation (9), and canceling out geometry terms, the above further simplifies to

$$p(t_{cb}, \omega_{cb}, t_{dc} | \Xi) = C_{\Xi} G(\mathbf{a}, \mathbf{b}). \quad (20)$$

We start with \mathcal{U}_1 and derive \mathcal{U}_2 and \mathcal{U}_3 in this order via repeated marginalization.

\mathcal{U}_1 : Derivation of $p(t_{cb} | \omega_{cb}, t_{dc}, \Xi)$. Following the standard definition of conditional PDFs in Equation (17), we have

$$p(t_{cb} | \omega_{cb}, t_{dc}, \Xi) = \frac{p(t_{cb}, \omega_{cb}, t_{dc} | \Xi)}{p(\omega_{cb}, t_{dc} | \Xi)}. \quad (21)$$

From Equations (18) and (20) it follows that the denominator above is obtained by integrating out t_{cb} :

$$\begin{aligned} p(\omega_{cb}, t_{dc} | \Xi) &= \int_0^{\infty} p(t_{cb}, \omega_{cb}, t_{dc} | \Xi) dt_{cb} \\ &= C_{\Xi} \int_0^{\infty} \underbrace{\frac{1}{h_{ca^{\perp}}^2 + (t_{ca^{\perp}} - t_{cb})^2}}_{G(\mathbf{a}, \mathbf{b})} dt_{cb} = C_{\Xi} \frac{\pi - \theta_{cb}}{t_{ca} \sin \theta_{cb}}, \end{aligned} \quad (22)$$

where $h_{ca^{\perp}}$ is the distance between \mathbf{a} and the ray $(\mathbf{c}, \omega_{cb})$, $t_{ca^{\perp}}$ is the distance between \mathbf{c} and this projection, and θ_{cb} is the angle between ω_{cb} and the line between \mathbf{c} and \mathbf{a} of length t_{ca} (see Figure 4c). By inserting Equations (20) and (22) into (21), we obtain a simple expression for the PDF of t_{cb} :

$$p(t_{cb} | \omega_{cb}, t_{dc}, \Xi) = \frac{t_{ca} \sin \theta_{cb}}{\pi - \theta_{cb}} \frac{1}{h_{ca^{\perp}}^2 + (t_{ca^{\perp}} - t_{cb})^2}. \quad (23)$$

The denominator of the second term above is simply the squared distance $t_{da^{\perp}}^2$. Note that this PDF is exactly equal to that of equiangular sampling [Kalli and Cashwell 1977; Rief et al. 1984; Kulla and Fajardo 2012] provided \mathbf{a} , ω_{cb} , and \mathbf{c} are given.

\mathcal{U}_2 : Derivation of $p(\omega_{cb} | t_{dc}, \Xi)$. The geometric configuration for sampling the direction ω_{cb} is illustrated in Figure 4b. Again, following the definition of conditional PDFs (17), we have

$$p(\omega_{cb} | t_{dc}, \Xi) = \frac{p(\omega_{cb}, t_{dc} | \Xi)}{p(t_{dc} | \Xi)}. \quad (24)$$

²Note that the constant C_{Ξ} now additionally incorporates the product of the three isotropic phase functions $(4\pi)^{-3}$.




Figure 4: Unidirectional factorization of the joint PDF in Equation (8). (a) First, a distance t_{dc} is sampled along the ray $(\mathbf{d}, \omega_{dc})$ which yields vertex \mathbf{c} . (b) Then, a direction ω_{cb} is sampled from \mathbf{c} , and finally (c) a distance t_{cb} is sampled along it, yielding vertex \mathbf{b} .

Following the marginalization chain, from Equation (22) we get

$$p(t_{dc} | \Xi) = \int_{S^2} p(\omega_{cb}, t_{dc} | \Xi) d\omega_{cb} \quad (25)$$

$$= \int_0^{2\pi} \int_0^{\pi} C_{\Xi} \frac{\pi - \theta_{cb}}{t_{ca} \sin \theta_{cb}} \sin \theta_{cb} d\theta d\phi = \frac{C_{\Xi} \pi^3}{t_{ca}}. \quad (26)$$

The resulting conditional PDF for sampling ω_{cb} , obtained by dividing Equation (22) by Equation (26), thus depends only on \mathbf{a} and \mathbf{c} since the normalization factors again cancel out:

$$p(\omega_{cb} | t_{dc}, \Xi) = \frac{\pi - \theta_{cb}}{\pi^3 \sin \theta_{cb}}. \quad (27)$$

Note that the above PDF diverges when θ_{cb} approaches zero. This same singularity also appears in the path contribution f (when f is expressed as a function of the angle θ_{cb}) [Rief et al. 1984]. By sampling ω_{cb} from this PDF (Section 5.2) we can effectively cancel out this singularity, reducing variance. Note also that the equiangular PDF (23) has no singularity since ω_{cb} is fixed.

\mathcal{U}_3 : Derivation of $p(t_{dc} | \Xi)$. The final marginalization step gives us the PDF for sampling the distance t_{dc} along the ray $(\mathbf{d}, \omega_{dc})$. To complete the definition for this PDF, which is already given by Equation (26), we only need to derive the normalization factor C_{Ξ} , which we do by enforcing the PDF (26) to integrate to 1:

$$\int_0^{\infty} p(t_{dc} | \Xi) dt_{dc} = \int_0^{\infty} \frac{C_{\Xi} \pi^3}{\sqrt{h_{da^{\perp}}^2 + (t_{da^{\perp}} - t_{dc})^2}} dt_{dc} = 1, \quad (28)$$

where $h_{da^{\perp}}$ is the distance between \mathbf{a} and its projection onto the ray $(\mathbf{d}, \omega_{dc})$, and $t_{da^{\perp}}$ is the distance between \mathbf{d} and this projection (see Figure 4a). Unfortunately, the above integral diverges at infinity; however, by setting a maximum sampling distance t_{dc}^{\max} along $(\mathbf{d}, \omega_{dc})$, we can obtain an analytical expression for it:

$$C_{\Xi} \pi^3 = \frac{1}{C_{t_{dc}^{\max}}} = \frac{1}{\operatorname{asinh}\left(\frac{t_{dc}^{\max} - t_{da^{\perp}}}{h_{da^{\perp}}}\right) - \operatorname{asinh}\left(\frac{-t_{da^{\perp}}}{h_{da^{\perp}}}\right)}. \quad (29)$$

Substituting Equation (29) into (26) yields the PDF for t_{dc} as simply

$$p(t_{dc} | \Xi) = \frac{1}{C_{t_{dc}^{\max}}} \frac{1}{\sqrt{h_{da^{\perp}}^2 + (t_{da^{\perp}} - t_{dc})^2}}. \quad (30)$$

The denominator of the second term above is just t_{ca} (see Figure 4a). Interestingly, this means that for double-scattering the propagation distance should be chosen proportionally to the *inverse distance* to \mathbf{a} , whereas for single-scattering the equiangular technique samples proportionally to the *inverse squared distance* to \mathbf{a} .

Joint unidirectional PDF. Now that we have completed the marginalization process, we can obtain a closed-form expression for the joint PDF in Equation (9) by multiplying the conditional PDFs \mathcal{U}_1 , \mathcal{U}_2 and, \mathcal{U}_3 from Equations (23), (27), and (30) respectively:

$$\begin{aligned} p(\mathbf{b}, t_{dc} | \Xi) &= p(t_{cb}, \omega_{cb}, t_{dc} | \Xi) G(\mathbf{b}, \mathbf{c}) \\ &= \frac{G(\mathbf{a}, \mathbf{b}) G(\mathbf{b}, \mathbf{c})}{\pi^3 C_{t_{dc}}^{\max}}. \end{aligned} \quad (31)$$

Substituting back into Equation (6), we can confirm that the final path PDF indeed includes the geometry term $G(\mathbf{abcd})$.

Note that we subsume Kalos' [1963] OCF approach by sampling two distances and a direction to construct two intermediate vertices between \mathbf{a} and \mathbf{d} . Kalos' OCF corresponds to sampling the vertex \mathbf{b} proportionally to the product of geometry terms to \mathbf{a} and a fixed \mathbf{c} . This has a PDF: $p(\mathbf{b} | t_{dc}, \Xi) = p(\mathbf{b}, t_{dc} | \Xi) / p(t_{dc} | \Xi)$, which arises by dividing Equation (31) by Equation (30).

5.2 Unidirectional Sampling Techniques

In order to use our importance sampling techniques, we not only need the above PDFs, but also the corresponding sampling routines derived from their inverse cumulative distribution functions (CDFs).

\mathcal{U}_3 : Sampling a distance t_{dc} . A random distance t_{dc} along the ray $(\mathbf{d}, \omega_{dc})$ can be sampled using the inverse CDF of Equation (30), derived by integrating the PDF over the ray and solving for t_{dc} :

$$t_{dc} = t_{da} \sinh(\xi C_{t_{dc}}^{\max}) - t_{da\perp} (1 + \cosh(\xi C_{t_{dc}}^{\max})), \quad (32)$$

where $t_{da} = \sqrt{h_{da\perp}^2 + t_{da\perp}^2}$ is the distance between \mathbf{d} and \mathbf{a} (see Figure 4a), and $\xi \in [0, 1]$ is a uniform random number.

\mathcal{U}_2 : Sampling a direction ω_{cb} . The PDF (27) is circularly symmetric (i.e. it depends only on θ_{cb}). Its corresponding CDF is:

$$P(\omega_{cb}) = \int_0^{2\pi} \int_0^{\theta_{cb}} \frac{\pi - \theta}{\pi^3 \sin \theta} \sin \theta \, d\theta \, d\phi = \frac{(2\pi - \theta_{cb})\theta_{cb}}{\pi^2}, \quad (33)$$

which we invert to solve for θ_{cb} , yielding the simple transformation

$$\theta_{cb} = \pi(1 - \sqrt{\xi_1}), \quad \text{and} \quad \phi_{cb} = 2\pi\xi_2, \quad (34)$$

where ξ_1 and ξ_2 are uniform random numbers, and θ_{cb} and ϕ_{cb} represent standard spherical coordinates with respect to a local frame where the direction from \mathbf{c} to \mathbf{a} is the z -axis (Figure 4b).

\mathcal{U}_1 : Sampling a distance t_{cb} . Given a uniform random number ξ , a distance t_{cb} along ω_{cb} is sampled using the inverse CDF of the equiangular PDF (23), yielding vertex \mathbf{b} :

$$t_{cb} = t_{ca\perp} + h_{ca\perp} \tan\left(\xi(\pi - \theta_{cb}) + \theta_{cb} - \frac{\pi}{2}\right). \quad (35)$$

The end result is a fully analytic method to importance sample the geometry terms along 3-segment paths (double-scattering) in isotropic participating media which generalizes and complements Kulla and Fajardo's [2012] equiangular sampling for 2-segment paths (single-scattering), as well as Kalos' [1963] OCF estimator.

5.3 Bidirectional Factorization PDFs

Similarly to the unidirectional case, the three conditional PDFs in Equations (14–16) are defined by the chain of marginalization of the joint PDF in Equation (13). The derivation of the first conditional PDF, \mathcal{B}_1 , is the same as for \mathcal{U}_1 . The resulting PDF is again equal to that of equiangular sampling given \mathbf{b} , ω_{dc} , and \mathbf{d} . However, in contrast to the unidirectional case, we were unable to obtain analytic expressions for the remaining two conditional PDFs, \mathcal{B}_2 and \mathcal{B}_3 . In the next section, we introduce a novel tabulation method that can handle both factorizations, while additionally supporting anisotropic phase functions, without requiring analytic formulations.

6 Tabulated Importance Sampling

To handle anisotropic scattering, we introduce a method to efficiently tabulate our target PDF in Equation (8) using both the unidirectional and the bidirectional factorizations. In constructing our tables, we assume a circularly-symmetric 1D phase function (i.e. depending only on the deflection angle between the incident and outgoing directions). Isotropic scattering is a special case where the phase function is simply $\rho = 1/4\pi$. As we will demonstrate in our results, our tabulations achieve nearly ideal importance sampling of the product of the geometry and scattering terms.

6.1 General Approach

As illustrated in Figure 3, both factorizations of our target PDF (8) consider two types of sampling events: (1) a distance along a given ray and (2) a direction from a given vertex. For a medium with a given phase function, an entire *family* of PDFs exists for the geometric configuration of each sampling event. For instance, given a ray $(\mathbf{c}, \omega_{cb})$, there is in general a different 1D PDF along that line for every possible relative location of vertex \mathbf{a} and its incident direction ω_{la} . Our task is then to construct a table that holds tabulated line PDFs for a discrete set of positions \mathbf{a} and directions ω_{la} .

Table parameterization. Even though each conditional PDF is one- or two-dimensional (corresponding to distance or direction sampling), the additional dimensionality of conditional variables makes the naive tabulation intractable. Our key idea to address this problem is to exploit the various symmetries in each geometric configuration of conditional variables by designing a suitable canonical coordinate system. This coordinate system allows us to dramatically reduce the dimensionality of our tables and their construction time.

PDF parameterization. Depending on the geometric configuration, the PDFs we tabulate may have large variations. This is because we consider the product of the geometric and scattering terms, both of which can have sharp peaks. Such peaks can significantly reduce the accuracy of the tabulation; more importantly, when the geometric term has a singularity, accurate tabulation is not even possible. We address this problem by warping the PDF domains via suitable reparameterizations that analytically eliminate geometric variations.

6.2 Unidirectional Factorization

For this factorization, illustrated in Figure 4, we seek to tabulate the factors given in Equations (10–12) for the joint PDF:

$$p(t_{cb}, \omega_{cb}, t_{dc} | \Xi) = C_{\Xi} G(\mathbf{a}, \mathbf{b}) \rho(\mathbf{a}) \rho(\mathbf{b}) \rho(\mathbf{c}). \quad (36)$$

The difference from Equation (20) is that now we also take the phase functions into account. For notation simplicity, in this section we will denote the PDFs and the normalization constant C_{Ξ} with the same symbols as in Section 5, in spite of them being different.

6.2.1 \mathcal{U}_1 : Tabulation for $p(t_{cb} | \omega_{cb}, t_{dc}, \Xi)$

This PDF is the anisotropic generalization of the PDF defined in Equation (21), and illustrated in Figure 4c. For the tabulation we can safely ignore the denominator, since—as a function of t_{cb} —it is merely a normalization constant, and we have to normalize our PDFs numerically anyway. We also ignore the terms C_Ξ and $\rho(\mathbf{c})$ in the numerator (36), as both are constant w.r.t. t_{cb} . We seek a PDF

$$p(t_{cb} | \omega_{cb}, t_{dc}, \Xi) \propto G(\mathbf{a}, \mathbf{b})\rho(\mathbf{a})\rho(\mathbf{b}). \quad (37)$$

Novák et al. [2012] proposed an on-the-fly tabulation of the same PDF. However, their approach suffers from limited accuracy and the overhead of on-the-fly construction. We avoid these issues by precomputing accurate approximations of a number of such PDFs once and storing them in a compact table which we can then use to sample a distance along any given ray. Our key observation is that this configuration, in general parameterized by $(\mathbf{a}, \mathbf{c}, \omega_{la}, \omega_{cb})$, can in fact be expressed using just two angles representing ω_{la} in a suitable canonical coordinate system. This allows us to precompute the entire family of PDFs into a compact 2D table of 1D PDFs.

Table construction. Figure 5c shows the canonical coordinate system used for the tabulation. Given a ray $(\mathbf{c}, \omega_{cb})$ and vertex \mathbf{a} , we first rescale the problem such that the distance between the line and the vertex is one. We place the origin at \mathbf{a} , the z -axis is aligned with the ray, and the x -axis lies in the same plane. The table is indexed by the direction vector ω_{la} . Each entry in the table is a PDF $p(u)$ over the (signed) distance u along the line, measured from the projection of \mathbf{a} onto the line. Similarly to Novák et al. [2012], we map the infinite line to a finite angular domain from the x -axis with $\theta = \text{arccot } u$. Each entry in our 2D table is a piecewise-linear approximation of the following PDF:

$$p(\theta(u)) = p(u) \left| \frac{du}{d\theta} \right| = p(u) \frac{1}{G(\mathbf{a}, \mathbf{b}(u))} \propto \rho(\mathbf{a}) \rho(\mathbf{b}(u)). \quad (38)$$

We tabulate the entire domain of $\theta \in [0; \pi]$, employing adaptive refinement to improve accuracy. Note that with the above change of variable, our tabulation is constant if the scattering is isotropic. In the special case where scattering at \mathbf{a} is isotropic (e.g. it is a point light source), each entry in the 2D table would correspond to the same 1D distribution.

Sampling. For any given ray $(\mathbf{c}, \omega_{cb})$, vertex \mathbf{a} , and direction ω_{la} , we first construct the coordinate system as described above. We then index the table with the canonical spherical coordinates of ω_{la} . We draw an angle θ_b from the retrieved tabulated PDF, constrained to the interval $\theta_b \in [0; \text{arccot } u_c]$ (see Figure 5c), and transform it to a canonical distance $u_b = \cot \theta_b$ to obtain the desired scattering distance along the ray $t_{cb} = -u_c + u_b$. We finally need to transform the PDF from the angular measure to the Euclidean length measure:

$$p(u_b) = p(\theta_b) \left| \frac{d\theta_b}{du_b} \right| = p(\theta_b) [G(\mathbf{a}, \mathbf{b}) h_{ca^\perp}], \quad (39)$$

where h_{ca^\perp} is the distance between \mathbf{a} and the ray (see Figure 4c).

6.2.2 \mathcal{U}_2 : Tabulation of $p(\omega_{cb} | t_{dc}, \Xi)$

This spherical distribution, used to sample the scattering direction ω_{cb} at vertex \mathbf{c} , is the anisotropic variant of the PDF defined in Equation (24), illustrated in Figure 4b. Once again, since we will normalize the PDF after tabulation, we can ignore the division by




Figure 5: Canonical coordinate systems for tabulating the conditional PDFs in the unidirectional factorization (10–12) of the joint PDF in Equation (8), for the case of anisotropic scattering. For each configuration we construct a table indexed by the entities colored in blue. The table entries are tabulated PDFs that are used for making the corresponding sampling decisions illustrated in Figure 4.

the normalization factor $p(t_{dc} | \Xi)$. We therefore seek to tabulate:

$$p(\omega_{cb} | t_{dc}, \Xi) \propto \int_0^\infty p(t_{dc}, \omega_{cb}, t_{cb} | \Xi) dt_{cb} \quad (40)$$

$$\propto \rho(\mathbf{c}) \underbrace{\int_0^\infty G(\mathbf{a}, \mathbf{b})\rho(\mathbf{a})\rho(\mathbf{b}) dt_{cb}}_{P(\omega_{cb})} \quad (41)$$

where $\rho(\mathbf{c})$ is in front of the integral, as it does not depend on t_{cb} . Note that the integral on the second line, $P(\omega_{cb})$, is the normalization of a PDF in our previous tabulation, \mathcal{U}_1 , and can be readily looked up without recomputation. The directional distribution is therefore proportional to the product $\rho(\mathbf{c})P(\omega_{cb})$.

Table construction. Figure 5b shows the canonical coordinate system for the tabulation. The origin is at \mathbf{c} , the z -axis is the direction from \mathbf{c} to \mathbf{a} , the x -axis is coplanar with the z -axis and ω_{la} . The PDF family (41) then has three degrees of freedom—the direction ω_{dc} and the angle θ_{la} of ω_{la} from the z -axis. We would thus need a 3D table of tabulated 2D PDFs, which is unfortunately impractical in terms of both storage and computation. We address this problem by tabulating $\rho(\mathbf{c})$ and $P(\omega_{cb})$ separately, and sample from their product. We build a 1D table of spherical PDFs for $P(\omega_{cb})$, indexed by θ_{la} , where for every tabulated direction ω_{cb} we evaluate $P(\omega_{cb})$ by looking up the corresponding PDF normalization from the \mathcal{U}_1 table. An intuitive interpretation of this process is that we compute the single-scattered incident radiance field from \mathbf{a} (omitting transmittance) at \mathbf{c} , for a number of lobe orientations at \mathbf{a} . We store the PDFs in the table as fitted mixtures of von Mises Fischer (vMF) distributions. This allows for both compact storage and efficient sampling from the product with the vMF-fitted phase function, which can be easily rotated on-the-fly [Tsai et al. 2008].

As we discussed in Section 5.1, this spherical PDF has a singularity around the pole. A naive discretization would therefore result in a highly inaccurate approximation, as it cannot capture the infinite variation of the function. To make tabulation possible, we warp the domain using Equation (34): $\omega_{u,v} = \omega(\pi(1 - \sqrt{u}), 2\pi v)$. With this change of variables, the PDFs we store are

$$p(\omega_{u,v}) = p(\omega) \left| \frac{d\omega}{d\omega_{u,v}} \right| = p(\omega) \sin \theta_{cb}. \quad (42)$$

We found that for scattering anisotropy of $g = 0.9$, every PDF in the 1D table for $P(\omega_{cb})$ can be accurately represented with 20

vMF lobes, and the phase function with 3 lobes. We fit a (randomly initialized) vMF mixture to every tabulated spherical PDF using a standard iterative expectation-maximization algorithm.

Sampling. Before sampling, we first transform the given configuration into the canonical coordinate system, and index the two tables to retrieve a product vMF distribution (41). We then sample (u, v) with PDF $p(u, v)$, which yields $\omega_{cb} = \omega_{u, v}$ using the above transformation. The solid angle PDF of the sampled direction is

$$p(\omega_{cb}) = p(u, v) \left| \frac{du dv}{d\omega_{cb}} \right| = p(u, v) \left[\frac{\pi - \theta_{cb}}{\pi^3 \sin \theta_{cb}} \right]. \quad (43)$$

Special case for $\rho(\mathbf{a}) = 1/4\pi$. When the scattering at \mathbf{a} is isotropic (e.g. it is a point light), the PDF family (41) has only one degree of freedom, the angle θ_{dc} between ω_{dc} and the z -axis. The whole family can thus be stored in a compact 1D table of 2D PDFs, for a set of angles θ_{dc} , without the need for product sampling.

6.2.3 \mathcal{U}_3 : Tabulation of $p(t_{dc} | \Xi)$

This last PDF in the marginalization chain is used to sample a distance along the ray $(\mathbf{d}, \omega_{dc})$. It is the anisotropic variant of the PDF defined in Equations (25) and (30) and illustrated in Figure 4a:

$$p(t_{dc} | \Xi) \propto \int_{\mathcal{S}^2} \int_0^\infty p(t_{dc}, \omega_{cb}, t_{cb} | \Xi) dt_{cb} d\omega_{cb} \quad (44)$$

$$\propto \int_{\mathcal{S}^2} p(\omega_{cb} | t_{dc}, \Xi) d\omega_{cb}. \quad (45)$$

The values of this PDF can be readily obtained by looking up in the corresponding normalizations from the \mathcal{U}_2 table described above.

Table construction. The geometric configuration for this family of PDFs, illustrated in Figure 5a, is identical to the one for t_{cb} , and we use the same coordinate system as in Figure 5c. The only difference is that this time we eliminate variations due to the geometric term by reparameterizing the tabulation domain using a different transformation: $v = \operatorname{asinh} u$. The stored PDFs are then

$$p(v(u)) = p(u) \left| \frac{du}{dv} \right| = p(u) \sqrt{1 + u^2} = p(u) t_a, \quad (46)$$

where t_a is the distance between the corresponding point and \mathbf{a} .

Sampling. Sampling a distance $t_{dc} = -u_d + u_c$ along a ray $(\mathbf{d}, \omega_{dc})$ proceeds analogously to the case for t_{cb} . We first construct the canonical coordinate system and retrieve the distribution corresponding to ω_{ba} . We then sample v_c from the corresponding tabulated PDF, constrained to the interval $v_c \in [0; \operatorname{asinh} u_d]$, and compute $u_c = \sinh v_c$. The line PDF of the sampled distance is:

$$p(t_{dc}) = p(u_c) = p(v_c) \left| \frac{dv_c}{du_c} \right| = p(v_c) \left[\frac{1}{t_{ca}} \right]. \quad (47)$$

6.3 Bidirectional Factorization

In the bidirectional factorization, we tabulate the factorized conditional PDFs in Equations (14–16) for the joint PDF:

$$p(\omega_{ab}, t_{ab}, t_{dc} | \Xi) = C_\Xi G(\mathbf{b}, \mathbf{c}) \rho(\mathbf{a}) \rho(\mathbf{b}) \rho(\mathbf{c}). \quad (48)$$

As in the isotropic case, the conditional PDF \mathcal{B}_1 is the same as \mathcal{U}_1 but with different variable names. We illustrate the coordinate system in Figure 6c for reference. We can thus readily reuse the PDF and CDF tables, and continue with the next PDF in the marginalization chain.




Figure 6: Canonical coordinate systems for tabulating the conditional PDFs in the bidirectional factorization (14–16) of the joint PDF in Equation (8). For each configuration we construct a table indexed by the entities colored in blue. Each table entry is a tabulated PDFs that is used for sampling distances or directions.

6.3.1 \mathcal{B}_2 : Tabulation of $p(t_{ab} | \omega_{ab}, \Xi)$

This PDF is used to sample a distance along a ray $(\mathbf{a}, \omega_{ab})$, given another ray $(\mathbf{d}, \omega_{dc})$ (see Figure 3c). Following the definition of conditional PDFs in Equation (17), we have:

$$\begin{aligned} p(t_{ab} | \omega_{ab}, \Xi) &= \frac{p(\omega_{ab}, t_{ab}, \Xi)}{p(\omega_{ab}, \Xi)} \propto p(\omega_{ab}, t_{ab}, \Xi) \\ &= \int G(\mathbf{b}, \mathbf{c}) \rho(\mathbf{a}) \rho(\mathbf{b}) \rho(\mathbf{c}) dt_{dc} \propto \int G(\mathbf{b}, \mathbf{c}) \rho(\mathbf{b}) \rho(\mathbf{c}) dt_{dc}. \end{aligned} \quad (49)$$

The marginal VRL PDF of Novák et al. [2012] roughly corresponds to this PDF, with $(\mathbf{d}, \omega_{dc})$ being the eye ray and $(\mathbf{a}, \omega_{ab})$ the VRL. Unlike Novák et al., however, we take into account anisotropic scattering at vertex \mathbf{b} and consider the semi-infinite extent of line $(\mathbf{d}, \omega_{dc})$ (which Novák et al. approximate as an infinite line). Our approach dramatically reduces variance, as we show in our results.

Table construction. Figure 6b illustrates the geometric configuration. In the canonical coordinate system the x -axis is aligned with the shortest connecting line between $(\mathbf{a}, \omega_{ab})$ and $(\mathbf{d}, \omega_{dc})$. The origin \mathbf{o} is the end of the connecting line on $(\mathbf{a}, \omega_{ab})$ and the z -axis is aligned with ω_{dc} . Similarly to the previous cases, an entire family of these PDFs exists, parameterized by the rays $(\mathbf{a}, \omega_{ab})$ and $(\mathbf{d}, \omega_{dc})$. Once again, the symmetries in the geometry configuration and scale invariance allow us to parameterize this family of PDFs by only a few parameters. In fact, with Novák et al.'s assumption that the ray $(\mathbf{d}, \omega_{dc})$ is infinite, we would only need a single table parameter, the θ_{ab} angle between ω_{ab} and ω_{dc} . However, the shape of the PDFs critically depends on the actual location of vertex \mathbf{d} along the $(\mathbf{d}, \omega_{dc})$ line. For that reason, we use the distance u_d from the projection of the origin onto $(\mathbf{d}, \omega_{dc})$ as an additional parameter, and build a 2D table of 1D PDFs. For every combination of angle θ_{ab} and distance u_d , the associated PDF assigns probability density to distance u from the origin on the (infinite) line $(\mathbf{a}, \omega_{ab})$. We build a piecewise-linear approximation of this PDF, by looking-up the integral (49), to which the constructed PDF is proportional, from the previously constructed \mathcal{B}_1 table.

The PDF $p(u)$ in the canonical coordinate system is defined over an entire infinite $(\mathbf{a}, \omega_{ab})$ line. To allow tabulation, we map this to a finite domain by parameterizing the position u using $v = \operatorname{asinh} u$. Using this change of variables, the PDFs stored in the table are

$$p(v(u)) = p(u) \left| \frac{du}{dv} \right| = p(u) \sqrt{1 + u^2} = p(u) t_{o^\perp}, \quad (50)$$




Figure 7: Comparison of our analytic and tabulated sampling techniques against state of the art techniques in path tracing and virtual ray light (VRL) [Novák et al. 2012] renderers on a simple scene with an infinite medium and an isotropic point light source. The number in each image indicates its relative render time. Note the poor performance of the analytic techniques with anisotropic phase functions; our tabulations perform particularly well, since transmittance (not importance sampled) has a smaller relative impact on the variability of the light transport.

This is the same transformation as in Equation (46), with t_{\perp} denoting the distance between the corresponding point and \mathbf{o}^{\perp} (Figure 6b). Any possible variation in the tabulation is solely due to the anisotropy of the phase function and the position of \mathbf{d} along the line $\omega_{\mathbf{d}\mathbf{c}}$.

Sampling. Given two rays $(\mathbf{d}, \omega_{\mathbf{d}\mathbf{c}})$ and $(\mathbf{a}, \omega_{\mathbf{a}\mathbf{b}})$, we construct the canonical coordinate system and retrieve the distribution corresponding to $\theta_{\mathbf{a}\mathbf{b}}$ and $u_{\mathbf{d}}$. We then sample $v_{\mathbf{b}}$ from the tabulated PDF, constraining it to the interval $v_{\mathbf{b}} \in [0; \text{asinh}u_{\mathbf{a}}]$, and compute $u_{\mathbf{b}} = \sinh v_{\mathbf{b}}$. The line PDF of the sampled distance $t_{\mathbf{a}\mathbf{b}} = -u_{\mathbf{a}} + u_{\mathbf{b}}$ is

$$p(t_{\mathbf{a}\mathbf{b}}) = p(u_{\mathbf{b}}) = p(v_{\mathbf{b}}) \left| \frac{dv_{\mathbf{b}}}{du_{\mathbf{b}}} \right| = p(v_{\mathbf{b}}) \left[\frac{1}{t_{\mathbf{b}\mathbf{o}^{\perp}}} \right]. \quad (51)$$

6.3.2 \mathcal{B}_3 : Tabulation of $p(\omega_{\mathbf{a}\mathbf{b}} | \Xi)$

The final PDF in the marginalization chain is used to sample a direction $\omega_{\mathbf{a}\mathbf{b}}$ at vertex \mathbf{a} (see Figure 3c). We seek to tabulate

$$p(\omega_{\mathbf{a}\mathbf{b}} | \Xi) = \frac{p(\omega_{\mathbf{a}\mathbf{b}}, \Xi)}{p(\Xi)} \propto \int p(\omega_{\mathbf{a}\mathbf{b}}, t_{\mathbf{a}\mathbf{b}} | \Xi) dt_{\mathbf{a}\mathbf{b}} \propto \rho(\mathbf{a}) \underbrace{\int G(\mathbf{b}, \mathbf{c}) \rho(\mathbf{b}) \rho(\mathbf{c}) dt_{\mathbf{a}\mathbf{b}}}_{P(\omega_{\mathbf{a}\mathbf{b}})} \quad (52)$$

where $\rho(\mathbf{a})$ is in front of the integral, as it does not depend on $t_{\mathbf{a}\mathbf{b}}$. Similarly to Section 6.2.2, the integral on the second line, denoted $P(\omega_{\mathbf{a}\mathbf{b}})$, is the normalization of a corresponding PDF in the \mathcal{B}_2 table above, and can be readily looked up. The desired distribution is proportional to the product $\rho(\mathbf{a})P(\omega_{\mathbf{a}\mathbf{b}})$.

Table construction. Figure 6a shows the canonical coordinate system used for the tabulation. The origin is at \mathbf{a} , the z -axis is aligned with $\omega_{\mathbf{d}\mathbf{c}}$, and the x -axis lies in the same plane. The PDF family (52) is then spanned by the direction $\omega_{l_{\mathbf{a}}}$ and the position of \mathbf{d} along $\omega_{\mathbf{d}\mathbf{c}}$, given by the distance $u_{\mathbf{d}}$ in the figure. Once again, this means that we generally need a 3D table of tabulated spherical PDFs. However, in this case the tabulated data cannot be efficiently represented by vMF mixtures, necessitating a different (e.g. wavelet-based) product sampling approach. We leave this for future work, and build our PDF tables assuming isotropic scattering at \mathbf{a} .

Similarly to the special case in 6.2.2, the problem reduces to constructing a 1D table of tabulated 2D PDFs, this time indexed by the distance $u_{\mathbf{d}}$. The spherical distributions in this family have a singularity around the z -axis, which is due to the $\sin^{-1} \theta$ factor in Novak et al.'s [2012] analytic inverse CDF. We eliminate this singularity by warping the spherical domain using the transformation $\omega_{u,v} = \omega(\pi\sqrt{1-u}, 2\pi v)$. With this change of variable, we store

$$p(\omega_{u,v}) = p(\omega) \left| \frac{d\omega}{d\omega_{u,v}} \right| = p(\omega) \sin \theta_{\mathbf{a}\mathbf{b}}. \quad (53)$$

Sampling. For this sampling decision, we are given a ray $(\mathbf{d}, \omega_{\mathbf{d}\mathbf{c}})$ and a point \mathbf{a} . After transforming this input configuration into the canonical coordinate system, we retrieve a PDF by indexing our table with $u_{\mathbf{d}}$. We then sample (u, v) with PDF $p(u, v)$, yielding $\omega_{\mathbf{a}\mathbf{b}} = \omega_{u,v}$ using the above transformation. The solid angle PDF of the sampled direction is

$$p(\omega_{\mathbf{a}\mathbf{b}}) = p(u, v) \left| \frac{du dv}{d\omega_{\mathbf{a}\mathbf{b}}} \right| = p(u, v) \left[\frac{1}{2\pi^2 \sin \theta_{\mathbf{a}\mathbf{b}}} \right]. \quad (54)$$




Figure 8: Unidirectional path tracing with isotropic (left, 15 min.) and anisotropic (right, 30 min.) media using traditional 0-random-decision (shadow) connections and our extensions incorporating 1- and 3-random-decision connection subpaths, resulting in $4\times$ to $6\times$ and $15\times$ to $38\times$ reduction in RMSE error (bottom) respectively.

7 Applications and Results

Our unidirectional and bidirectional sampling techniques described in the previous two sections provide an effective way to construct 3-random-decision connection subpaths. The most direct application of this method is to compute single- and double-scattering in participating media by connecting eye rays to random locations on light sources. This corresponds to the special case where the prefix and postfix subpaths \bar{x}_l and \bar{x}_e introduced in Section 4 are both of zero length, i.e. \mathbf{a} is on a light source and \mathbf{d} is on the eye lens.

Though our routines do not currently consider chains of random decisions longer than 3, we can still leverage them to improve higher-order multiple scattering. We do this by generalizing deterministic (0-random-decision) shadow connections used in unidirectional and bidirectional path tracing, into 1-, 2-, and 3-random-decision subpaths connecting light and eye subpaths of arbitrary lengths.

In this section we describe how to incorporate our sampling routines for practical improvements in rendering. We compare our enhanced algorithms to classical unidirectional and bidirectional path tracing as well as VRL rendering. Figure 7 summarizes the connection techniques at our disposal, demonstrating their ideal-case performance. Due to space limitations, we omit the images from the bidirectional 3-random-decision subpath technique, which produces similar results to the unidirectional 3-random-decision subpath technique. Table 1 summarizes the resolution, memory footprint, and computation time for the PDF tabulations described in Section 6. All measurements have been obtained on a mobile quad-core Intel Core i7-2820QM 2.3GHz processor using a CPU ray tracer.

7.1 Unidirectional Path Tracing

0 random decisions. The most common forward path tracer uses unidirectional path sampling of transmittance and phase function starting at the eye, combined with deterministic (0-random-decision) shadow connections (a special case of Figure 2b). We use this algorithm as a baseline for our comparisons.

Figures 1 and 8 show a scene with isotropic (left) and highly anisotropic (right) phase functions and scattering albedo 0.57. Even




Figure 9: A scene containing an anisotropic medium ($g = 0.9$) and an anisotropic point light, rendered with unidirectional path tracing using Novák et al.’s [2012] on-the-fly tabulated 1-connections (left) and our improved tabulated 3-connections (right).

with explicit shadow connections, traditional path tracing (far left) suffers from significant spike noise as it cannot importance sample the geometry term or phase function when making the connection.

1 random decision. Our first extension to unidirectional path tracing incorporates 1-random-decision connection subpaths. We generate a random-walk path starting at the eye, just as with standard unidirectional path tracing. However, instead of connecting each eye path vertex to the light with a deterministic connection, we independently connect each eye path segment to the light using a 1-random-decision subpath, sampled with the \mathcal{U}_1 technique. This corresponds to using (isotropic) equiangular sampling or Novák et al.’s [2012] anisotropic generalization as a connection strategy between an eye path and a light. As in previous work, we combine with transmittance importance sampling using MIS.

This extension corresponds to interpreting previous work as a general connection strategy. On top of this, our precomputed tabulation improves render time and affords a more accurate approximation of the anisotropic PDF since we avoid coarse on-the-fly tabulation. The center columns in Figure 8 shows unidirectional path tracing with 1-random-decision connections. This reduces RMS error by a factor of 4 to 6 (a $16\times$ to $36\times$ improvement in rendering convergence time) compared to using only deterministic shadow connections.

2 and 3 random decisions. Our main contribution for unidirectional path tracing additionally incorporates our \mathcal{U}_2 and \mathcal{U}_3 sampling techniques for up to 3-random-decision connections (Figure 3b) between each eye path segment and a light source. We also sample a 1-random-decision connection for the first eye path segment to account for single scattering. We use Equations (32–35) for isotropic media and our tabulated equivalents for anisotropic media. For each decision, we randomly choose between our phase function and geometry term sampling routines and traditional transmittance and phase function sampling, and use the one-sample MIS balance heuristic to combine the estimates. The resulting algorithm fully subsumes traditional path tracing (i.e. with 0-random-decision connections), which corresponds to always selecting the traditional techniques.

Figures 1 and 8 show the result of rendering the Dragon scene using 3-random-decision connection subpaths. This provides a substantial variance reduction compared to using just 1-random-decision connections, resulting in a $15\times$ (isotropic) to $38\times$ (anisotropic) reduction in total RMS error compared to deterministic shadow connections. This corresponds to a $225\text{--}1444\times$ speedup in rendering convergence time. The Lighthouse scene in Figure 9 has an anisotropic light source and albedo 0.8, and Figure 10 shows the

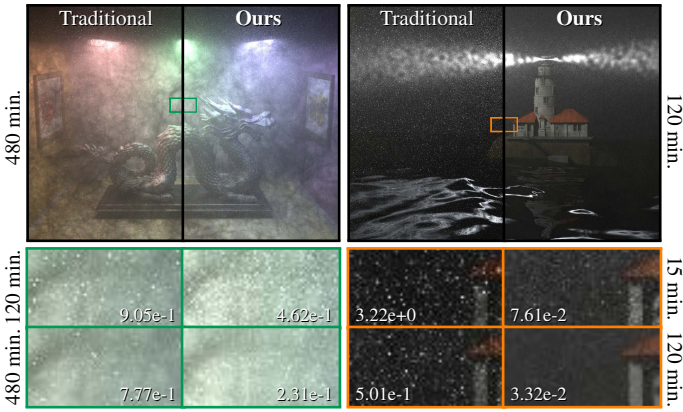


Figure 10: Two scenes with anisotropic ($g = 0.9$) heterogeneous media, rendered with unidirectional path tracing. In contrast to our improved (3-connection) techniques, the traditional (0-connection) techniques miss most of the double-scattered illumination.

two scenes with heterogeneous media; both figures show significant variance reduction using our 3-connection techniques. For heterogeneous media, we do not use Kulla and Fajardo’s [2012] transmittance tabulation, though a combination with our techniques would be interesting to investigate. In Figure 11 we demonstrate longer bounces (path lengths 1–8) with and without our improvements.

7.2 Virtual Ray Lights

We incorporated our tabulated bidirectional sampling strategy into a many-light renderer. This renderer generates random-walk paths from the lights with transmittance and phase function importance sampling, and stores them as a collection of VRLs. During rendering, eye rays are connected to the VRLs using a one-sample Monte Carlo estimator. See Novák et al. [2012] for complete details.

Our \mathcal{B}_1 , \mathcal{B}_2 , and \mathcal{B}_3 tabulations provide three concrete improvements over Novák et al.’s [2012] approach in Equation (5): 1) our tabulations are precomputed, affording higher precision and avoiding expensive on-the-fly tabulation; 2) we eliminate approximations in the marginal distribution (see Section 6.3.1) which significantly reduces noise; and 3) we can importance sample one additional random decision (the emitted direction of the VRL) which allows us to sample a set of VRLs relevant to the viewpoint. We evaluate these improvements compared to a baseline VRL renderer.

In Figure 12, we compare Novák et al.’s approach to our improved 2-random-decision (left) and 3-random-decision (right) connections (albedo 0.89). The right image in particular considers an uncorrected variation of VRLs, where a VRL is generated independently for each eye ray, instead of using a single set of VRLs for the entire image. Our 3-random-decision technique not only importance samples the connection to the VRLs, but also importance samples the emitted direction of the VRLs themselves. In both cases, we see substantial quality improvement in the same rendering time.

Table 1: Resolution, memory footprint, and construction time statistics for the PDF tabulations illustrated in Figures 5 and 6.

	Table resolution	PDF resolution	Memory	Time
$\mathcal{U}_1, \mathcal{B}_1$	$\theta_{1a}: 200, \phi_{1a}: 400$	$u_b: 35$ (avg)	22 MB	320 ms
\mathcal{U}_2	$\theta_{1a}: 90$	$\omega_{cb}: 20$ lobes	36 KB	6 s
\mathcal{U}_3	$\theta_{1a}: 90, \phi_{1a}: 180$	$u_c: 100$	13 MB	580 ms
\mathcal{B}_2	$u_d: 300, \theta_{ab}: 100$	$u_b: 200$	24 MB	530 ms
\mathcal{B}_3	$u_d: 100$	$\theta: 200, \phi: 400$	32 MB	680 ms




Figure 11: The Dragon scene with an isotropic and anisotropic medium, rendered with unidirectional path tracing. We visualize path lengths 1–8 and compare the traditional (0-connection) sampling techniques to our extended (3-connection) method.

7.3 Bidirectional Path Tracing

Bidirectional path tracing (BPT) also traditionally relies on deterministic connections (see Figure 2b), but typically converges much faster than unidirectional path tracing, as it can combine many techniques for constructing the same path. We will show that by incorporating 2-random-decision connections into BPT we can obtain significant variance reduction even for this generally more robust algorithm.

0 random decisions. Our baseline for comparison is a standard bidirectional path tracer. Both the light \bar{x}_l and eye \bar{x}_e subpaths are created using transmittance and phase function importance sampling, and then each pair of vertices are coupled with a deterministic 0-random-decision connection to construct a full path from a light to the eye. We combine the various ways of constructing paths using MIS with the balance heuristic.

1 and 2 random decisions. Instead of connecting each pair of vertices, we incorporate our \mathcal{B}_1 and \mathcal{B}_2 techniques for 2-random-decision connections between each pair of path segments on the light and eye subpaths, as illustrated in Figure 2d. A light subpath segment defines \mathbf{a} and its outgoing direction ω_{ab} , and an eye subpath segment provides \mathbf{d} and its incoming direction ω_{dc} . Given this configuration, we use our tabulated bidirectional factorization (Section 6) to randomly sample t_{ab} and t_{dc} proportionally to the product of inverse squared distance, $G(\mathbf{b}, \mathbf{c})$, and phase functions, $\rho(\mathbf{b})\rho(\mathbf{c})$. We perform this connection for every eye-light segment pair. This effectively converts VRLs from a many-light algorithm into a bidirectional connection strategy, leveraging our more efficient and accurate tabulated joint PDFs. As with all our other implementations, we combine transmittance and phase function importance sampling via MIS, thereby subsuming 0- and 1-random-decision connections.

Figure 1 shows a scene with medium albedo 0.57 rendered with bidirectional path tracing. Even though BPT is generally much more robust than unidirectional path tracing, incorporating our 2-random-decision connections results in visible noise reduction; RMS error is reduced by roughly a factor of 3.5—a $12\times$ improvement in convergence speed for path lengths 1–3. Since our techniques importance sample single and double scattering, we expect the benefits to be most significant for low-order scattering. Applying our approach to longer paths also results in some (albeit less) variance reduction.




Figure 12: Left: Our tabulated PDFs are more accurate than the ones of Novák et al.’s [2012] and dramatically reduce variance when making connections between correlated VRLs and eye rays. Right: Our ability to importance sample the emitted directions of VRLs further reduces variance for uncorrelated VRL rendering.

8 Discussion and Limitations

Variance due to other terms. Although the geometry and scattering terms are often responsible for high variance, other terms can introduce high variance in certain scene configurations. For example, transmittance sampling may outperform our methods in scenes with highly heterogeneous and dense media. In scenes with uniform light emission (e.g. a constant environment map), traditional unidirectional sampling (Figure 2a) will likely be the best strategy. We try to handle such cases robustly using MIS; however, by splitting samples across multiple strategies we allocate fewer samples to a potentially better strategy. Though our method outperforms existing techniques in all our tests, we obtained less benefit in scenes with a combination of caustic paths and surface-to-media transport, e.g., the noise visible near the floor and stools in the Stage scene in Figure 1.

Longer connections. While our method is tailored to single and double scattering, in Section 7 we demonstrate practical improvement for path lengths up to 8 by using our techniques as generalized “shadow connections”. This improvement will likely diminish for even longer paths; however, since longer paths also have smaller contribution, the final image will still benefit from our improvement in lower-order scattering. We believe it is theoretically possible to extend our techniques to jointly importance sample more than three random decisions by continuing the marginalization chain. The theoretical challenge is to derive either analytic expressions or proper tabulation for the resulting conditional PDFs. Our theoretical framework could be used as a stepping stone for such developments.

Anisotropic light sources. Our current implementation of the bidirectional factorization does not support 3-random-decision connections when the scattering at vertex \mathbf{a} is anisotropic. For this particular case we would need a product importance sampling technique that is better suited to the shape of the distributions. This is mainly a limitation of our implementation, not our theory.

Convergence rates. As discussed earlier, Kalos [1963] showed that the UCF estimator (which corresponds to standard unidirectional path tracing with explicit shadow connections) results in *infinite* variance in the presence of light sources inside participating media. Moreover, the central limit theorem does not hold, leading to a substantially lower asymptotic convergence rate for volumetric path tracing than its surface equivalent ($1/\sqrt[3]{N}$ instead of $1/\sqrt{N}$). By reducing the order of the singularity in the geometry term, Kalos’ OCF estimator provides bounded variance and recovers a $1/\sqrt{N}$ convergence rate. Our techniques inherit this property by subsuming the OCF estimator. Nonetheless, a formal analysis of the exact

convergence rates of our methods would be interesting, especially in combination with multiple importance sampling where singularities are further reduced by low weights. Dubei et al. [1982] and Reif [1984] also showed that the combination of path tracing and equiangular sampling retains the $1/\sqrt{N}$ convergence rate, despite the fact that this combination still leaves an angular singularity and has infinite variance. Our techniques also employ equiangular sampling but additionally remove this angular singularity, leading to reduced variance.

9 Conclusion

We presented joint path importance sampling techniques of the product of geometry and scattering terms for rendering participating media. We revealed that traditional incremental path construction can lead to significant variance, and provided a unifying view of existing methods as subpath connection techniques. We developed a subpath construction strategy to importance sample three sequential random decisions. We derived simple analytic expressions for unidirectional sampling in isotropic media, and introduced a compact tabulation scheme to handle anisotropic scattering and bidirectional sampling by exploiting various geometric symmetries in participating media. We demonstrated that these techniques can be easily incorporated into various rendering algorithms and reduce variance by up to three orders of magnitude compared to currently used methods. We believe that our theoretical and practical contributions create a promising new direction of future work in joint path sampling.

Acknowledgements. We thank the anonymous reviewers for their suggestions on improving the paper, Jan Novák for pointing us to the related work in neutron transport, and Marios Papas and Tomáš Davidovič for running the heterogeneous tests. Jiří Vorba provided the vMF code. We have assembled our scenes using free models from Archive3D.net. This work has been partially supported by the Czech Science Foundation grant P202-13-26189S.

References

- BURKE, D., GHOSH, A., AND HEIDRICH, W. 2005. Bidirectional importance sampling for direct illumination. In *Proc. of Eurographics Symposium on Rendering*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, EGSR’05, 147–156.
- CHANDRASEKAR, S. 1960. *Radiative Transfer*. Dover Publications.
- CLARBERG, P., AND AKENINE-MÖLLER, T. 2008. Practical product importance sampling for direct illumination. *Computer Graphics Forum (Proc. of Eurographics)* 27, 2.
- CLARBERG, P., JAROSZ, W., AKENINE-MÖLLER, T., AND JENSEN, H. W. 2005. Wavelet importance sampling: efficiently evaluating products of complex functions. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 24, 3 (July).
- CLINE, D., TALBOT, J., AND EGBERT, P. 2005. Energy redistribution path tracing. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 24, 3 (July).
- D’EON, E., AND IRVING, G. 2011. A quantized-diffusion model for rendering translucent materials. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 30, 4.
- DUBI, A., ELPERIN, T., AND RIEF, H. 1982. On confidence limits and statistical convergence of Monte Carlo point-flux estimators with unbounded variance. *Annals of Nuclear Energy* 9, 675–682.
- DUTRÉ, P., BALA, K., AND BEKAERT, P. 2006. *Advanced Global Illumination*, 2nd ed. A. K. Peters.

- JAKOB, W., AND MARSCHNER, S. 2012. Manifold exploration: A Markov chain Monte Carlo technique for rendering scenes with difficult specular transport, expanded technical report. Tech. rep., Cornell University, May.
- JAKOB, W., AND MARSCHNER, S. 2012. Manifold exploration: A Markov chain Monte Carlo technique for rendering scenes with difficult specular transport. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 31, 4.
- JAROSZ, W., ZWICKER, M., AND JENSEN, H. W. 2008. The beam radiance estimate for volumetric photon mapping. *Computer Graphics Forum (Proc. of Eurographics)* 27, 2.
- JAROSZ, W., CARR, N. A., AND JENSEN, H. W. 2009. Importance sampling spherical harmonics. *Computer Graphics Forum (Proc. of Eurographics)* 28, 2 (Apr.), 577–586.
- JAROSZ, W., NOWROUZEZAHRAI, D., SADEGHI, I., AND JENSEN, H. W. 2011. A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Transactions on Graphics* 30, 1, 5:1–5:19.
- JENSEN, H. W., AND CHRISTENSEN, P. H. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proc. of SIGGRAPH '98*.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proc. of SIGGRAPH '01*.
- JENSEN, H. W. 1996. Global illumination using photon maps. In *Proc. of Eurographics Rendering Workshop*.
- KAJIYA, J. T. 1986. The rendering equation. In *Computer Graphics (Proc. of SIGGRAPH)*.
- KALLI, H., AND CASHWELL, E. 1977. Evaluation of three Monte Carlo estimation schemes for flux at a point. Tech. Rep. LA-6865-MS, Los Alamos Scientific Lab, New Mexico, USA.
- KALOS, M. H. 1963. On the estimation of flux at a point by Monte Carlo. *Nuclear Science and Engineering* 16, 111–117.
- KŘIVÁNEK, J., HAŠAN, M., ARBREE, A., DACHSBACHER, C., KELLER, A., AND WALTER, B. 2012. Optimizing realistic rendering with many-light methods. In *ACM SIGGRAPH 2012 Courses, SIGGRAPH '12*.
- KULLA, C., AND FAJARDO, M. 2012. Importance sampling techniques for path tracing in participating media. *Computer Graphics Forum (Proc. of Eurographics Symposium on Rendering)* 31, 4.
- NOVÁK, J., NOWROUZEZAHRAI, D., DACHSBACHER, C., AND JAROSZ, W. 2012. Virtual ray lights for rendering scenes with participating media. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 31, 4 (July).
- PAULY, M., KOLLIG, T., AND KELLER, A. 2000. Metropolis light transport for participating media. In *Rendering Techniques (Proc. of Eurographics Workshop on Rendering)*.
- PEGORARO, V., SCHOTT, M., AND PARKER, S. G. 2009. An analytical approach to single scattering for anisotropic media and light distributions. In *Proc. of Graphics Interface*, 71–77.
- PHARR, M., AND HUMPHREYS, G. 2010. *Physically Based Rendering: From Theory to Implementation*, 2nd ed. Morgan Kaufmann, San Francisco, USA.
- PREMOŽE, S., ASHIKHMIN, M., AND SHIRLEY, P. S. 2003. Path integration for light transport in volumes. In *Proc. of Eurographics Symposium on Rendering*, 52–61.
- PREMOŽE, S., ASHIKHMIN, M., RAMAMOORTHY, R., AND NAYAR, S. 2004. Practical rendering of multiple scattering effects in participating media. In *Proc. of Eurographics Symposium on Rendering*, 363–374.
- RAAB, A., AND BEIKERT, G. 1999. Two new Monte Carlo methods for point flux estimation. *Computer Physics Communications* 123, 27–45.
- RAAB, M., SEIBERT, D., AND KELLER, A. 2008. Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer, 591–606.
- RIEF, H., DUBI, A., AND ELPERIN, T. 1984. Track length estimation applied to point detector. *Nuclear Science and Engineering* 87, 59–71.
- SPANIER, J., AND GELBARD, E. M. 1969. *Monte Carlo principles and neutron transport problems*. Addison-Wesley.
- SPANIER, J. 1966. Two pairs of families of estimators for transport problems. *SIAM Journal on Applied Mathematics* 14, 4, 702–713.
- STAM, J. 1995. Multiple scattering as a diffusion process. *Rendering Techniques (Proc. of Eurographics Workshop on Rendering)*, 41–50.
- STEINBERG, H. A., AND KALOS, M. H. 1971. Bounded estimators for flux at a point in Monte Carlo. *Nuclear Science and Engineering* 44, 406–412.
- SUN, B., RAMAMOORTHY, R., NARASIMHAN, S. G., AND NAYAR, S. K. 2005. A practical analytic single scattering model for real time rendering. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 24, 3.
- SUN, X., ZHOU, K., LIN, S., AND GUO, B. 2010. Line space gathering for single scattering in large scenes. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 29, 4.
- TESSENDORF, J. 1987. Radiative transfer as a sum over paths. *Physical Review A* 35 (Jan.), 872–878.
- TSAI, Y.-T., CHANG, C.-C., JIANG, Q.-Z., AND WENG, S.-C. 2008. Importance sampling of products from illumination and brdf using spherical radial basis functions. *The Visual Computer* 24, 7 (July).
- VEACH, E., AND GUIBAS, L. J. 1995. Optimally combining sampling techniques for Monte Carlo rendering. In *SIGGRAPH '95*.
- VEACH, E., AND GUIBAS, L. J. 1997. Metropolis light transport. In *SIGGRAPH '97*.
- VEACH, E. 1997. *Robust Monte Carlo methods for light transport simulation*. PhD thesis, Stanford, CA, USA.
- WALTER, B., ZHAO, S., HOLZSCHUCH, N., AND BALA, K. 2009. Single scattering in refractive media with triangle mesh boundaries. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 28, 3.
- WOODCOCK, E., MURPHY, T., HEMMINGS, P., AND T.C., L. 1965. Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry. In *Applications of Computing Methods to Reactor Problems*, Argonne National Laboratory.