

Virtual Ray Lights for Rendering Scenes with Participating Media

Jan Novák^{1,2}

Derek Nowrouzezahrai^{1,3}

Carsten Dachsbacher²

Wojciech Jarosz¹

¹Disney Research Zürich

²Karlsruhe Institute of Technology

³Université de Montréal

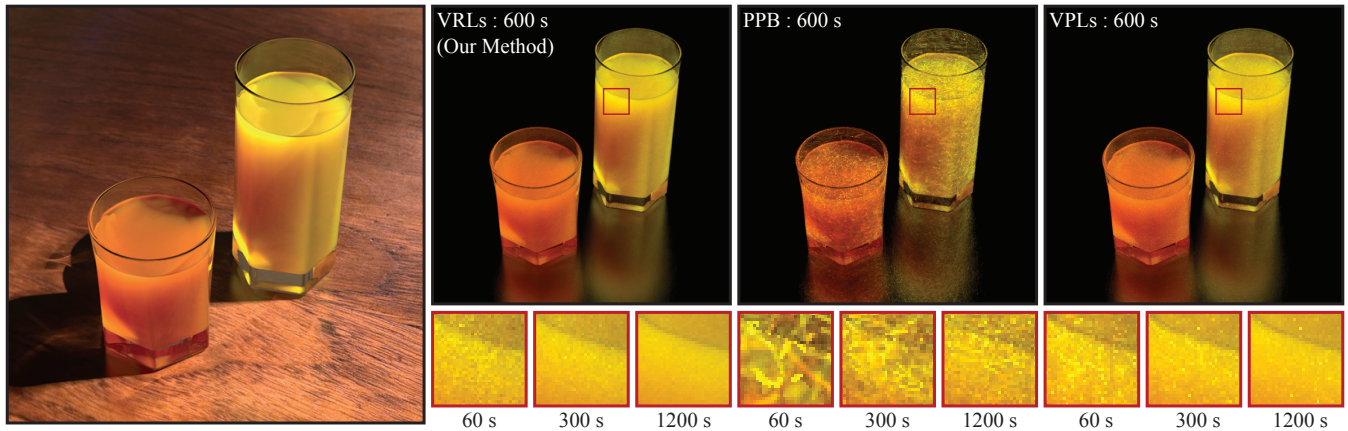


Figure 1: The FRUIT JUICE scene featuring both single and multiple scattering (left). We evaluate the quality of multiple scattering. Our approach can compute the multiple scattering within the media in about a minute (middle-left), while in equal time, both the previous state-of-the-art approaches for general media (progressive photon beams, middle-right; virtual point lights, right) contain significant artifacts.

Abstract

We present an efficient many-light algorithm for simulating indirect illumination in, and from, participating media. Instead of creating discrete virtual point lights (VPLs) at vertices of random-walk paths, we present a continuous generalization that places *virtual ray lights* (VRLs) along each path segment in the medium. Furthermore, instead of evaluating the lighting independently at discrete points in the medium, we calculate the contribution of each VRL to entire camera rays through the medium using an efficient Monte Carlo product sampling technique. We prove that by spreading the energy of virtual lights along both light and camera rays, the singularities that typically plague VPL methods are significantly diminished. This greatly reduces the need to clamp energy contributions in the medium, leading to robust and *unbiased* volumetric lighting not possible with current many-light techniques. Furthermore, by acting as a form of final gather, we obtain higher-quality multiple-scattering than existing density estimation techniques like progressive photon beams.

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

Keywords: participating media, global illumination, virtual lights

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

1 Introduction

Volumetric scattering is responsible for many of the intricate visual effects we observe in natural phenomena like clouds and fog, in the interaction of light with drinks and foods (Figure 1), or even with air. This complex light transport is governed by the radiative transport [Chandrasekar 1960] and rendering equations [Kajiya 1986]. Unfortunately, solving these accurately is notoriously expensive.

Though variants of path tracing [Lafortune and Willems 1993; Veach and Guibas 1994; Lafortune and Willems 1996] provide a general solution for most lighting configurations, they converge slowly. Alternatives using diffusion theory [Jensen et al. 2001; Donner and Jensen 2005] provide efficient approximations but are only applicable with strong constraints on the scattering properties of the medium. For transport in general participating media, including multiple scattering, two-pass approaches like volumetric photon mapping (VPM) or virtual point light (VPL) methods are often the preferred choice in offline rendering.

VPM approaches [Jensen and Christensen 1998; Jarosz et al. 2008] shoot photons from light sources, and then compute the lighting by estimating the local density of stored photons at each spatial location. Unfortunately, a large number of photons may be required to accurately resolve complex lighting. The recent photon beams approach [Jarosz et al. 2011a; Jarosz et al. 2011b] stores entire paths of photons instead of just the vertices, significantly increasing data density and therefore the accuracy of this class of algorithms.

VPL methods [Keller 1997; Walter et al. 2005; Walter et al. 2006; Raab et al. 2008] use the same framework as photon mapping (shooting photons from light sources), but then convert each stored photon into a “virtual” point light source. Multiple scattering is estimated using standard direct lighting techniques from this collection of VPLs. Unfortunately, VPL techniques suffer from distracting artifacts in the form of local singularities of high intensity. These can be avoided by clamping or blurring [Hašan et al. 2009], at the expense of introducing bias. Though this bias can be corrected, available techniques are either approximate [Novák et al. 2011] or incur substantial computational cost [Kollig and Keller 2006].




Figure 2: *Virtual point light methods (left) convert the vertices of a random-walk pre-process into a collection of virtual point lights (VPLs). Our method (right) converts entire segments of the random-walk into virtual ray lights instead. This results in denser sampling, provably weaker singularities, and higher quality when estimating illumination from the collection of virtual lights (bottom).*

Our approach combines the strengths of these two classes of algorithms. We shoot and store random-walk paths from the light, then convert them into “virtual ray lights” (VRLs), which we use to compute unbiased multiply scattered light from the media onto itself and onto surfaces (Figure 2). Compared to both VPL methods and photon beams, we obtain higher quality multiple scattering at a fraction of the cost by exploiting these important benefits:

- using line segments instead of points allows us to more densely sample media radiance with fewer virtual lights;
- distributing energy over line segments (instead of concentrating it at discrete points) provably reduces the amount *and* order of singularities, diminishing the need to blur or clamp energy;
- using a product importance sampling scheme to efficiently integrate over the 2D space of camera rays \times VRLs allows us to more robustly handle the traditionally challenging cases (for VPL methods) of scattering from anisotropic phase functions;
- by acting like a final gather pass over photons beams, we obtain higher quality multiple scattering; and,
- unlike most virtual light approaches, the resulting algorithm is unbiased and therefore trivially supports progressive updates.

We efficiently compute the contribution of media radiance by solving an integral over a 2D domain formed by each camera/light ray pair, instead of discretely decomposing this into two nested 1D problems. We present a product sampling method to efficiently compute these integrals in a Monte Carlo framework. Our unbiased approach is effective in the presence of highly-anisotropic scattering/lighting and outperforms other variance reduction techniques, such as multiple importance sampling [Veach 1997], by more accurately modeling the underlying distribution of the 2D domain.

2 Related Work

Participating media rendering has been studied extensively. We discuss only works most relevant to ours, referring the reader to a comprehensive survey [Cerezo et al. 2005] for a broader introduction.

Photon Mapping. Photon mapping [Jensen 1996] is general, relatively efficient, and one of the few practical algorithms which can robustly handle caustic light paths. Several extensions seek to decrease errors caused by density estimation [Lastra et al. 2002;

Havran et al. 2005; Herzog et al. 2007; Moon and Marschner 2006], but ultimately some amount of bias remains. Progressive photon mapping [Hachisuka et al. 2008b; Hachisuka and Jensen 2009; Knaus and Zwicker 2011] eliminates this error by shooting and discarding photons while progressively refining radiance statistics so as to converge in the limit. Since our approach is unbiased, it is naturally convergent and trivially supports progressive updates.

Jensen and Christensen [1998] first generalized photon mapping to participating media and Jarosz et al. [2008] later improved this approach by finding all photons along the length of each camera ray in one query. More recently, Jarosz et al. [2011a] introduced the photon beams method which performs density estimation over photon path segments instead of just the path vertices. A similar idea was concurrently suggested by Sun et al. [2010] to simulate single scattering and caustics by finding nearly intersecting camera and light paths. Due to the increased sample density, beams significantly improve quality in media compared to photons, and can also be formulated progressively [Jarosz et al. 2011b] to obtain convergent results. We also leverage the increased sampling density provided by photon paths but apply it within a virtual point light framework.

Virtual Point Lights. Since the introduction of the original VPL rendering algorithm, instant radiosity [Keller 1997], VPLs have formed the basis of a large number of surface rendering algorithms. In the interactive setting [Wald et al. 2002], acceleration is typically achieved by selecting a small but representative set of hundreds to thousands of VPLs [Hašan et al. 2007], by making approximations to visibility [Ritschel et al. 2008; Dong et al. 2009], or by providing progressive updates [Laine et al. 2007; Dammertz et al. 2010]. Offline algorithms [Walter et al. 2005; Walter et al. 2006; Ou and Pelacini 2011] typically use millions of VPLs for higher quality and then hierarchically prune this set during rendering to reduce visibility computation. In our algorithm, each camera ray/VRL pair actually parametrizes a continuous family of light/gather point pairs, and hence our 2D product sampling approach can be thought of as an implicit clustering and pruning of an infinite collection of VPLs.

Unfortunately, all VPL methods suffer from “weak singularities”, resulting in distracting artifacts in the form of spikes of high intensity. These are typically eliminated by clamping or blurring [Hašan et al. 2009], at the risk of introducing significant bias and modifying material appearance [Křivánek et al. 2010]. In some situations this bias can be compensated approximately [Novák et al. 2011], but full compensation [Kollig and Keller 2006] reverts to brute-force path tracing, especially in non-Lambertian environments.

Given the extensive VPL research for surface rendering, surprisingly little work exists on using VPLs in participating media. Arbrece et al. [2008] applied lightcuts to approximate subsurface scattering using the diffusion dipole [Jensen et al. 2001]. Though the hierarchical pruning of multidimensional lightcuts [Walter et al. 2006] can be used for participating media, handling general multiple scattering is incredibly difficult because VPL singularities become visible virtually everywhere. As with surfaces, bias compensation can recover lost energy in either exact [Raab et al. 2008] or approximate form [Engelhardt et al. 2010], but is necessary for virtually all pixels and becomes impractical for anisotropic scattering. We robustly handle anisotropic scattering and show that by spreading energy across both camera and light rays, singularities become provably less severe, diminishing the need for these corrections.

3 Overview and Problem Statement

Radiative Transport. The light arriving at a location \mathbf{x} in a direction $\vec{\omega}$ can be expressed as a sum of two terms:

$$L(\mathbf{x}, \vec{\omega}) = T_r(s)L_s(\mathbf{x}_s, \vec{\omega}) + L_m(\mathbf{x}, \vec{\omega}). \quad (1)$$

The first term is the radiance, L_s , from the nearest surface along the ray at $\mathbf{x}_s = \mathbf{x} - s\vec{\omega}$, attenuated due to transmittance T_r . Evaluating transmittance involves integrating the extinction coefficient σ_t along the line connecting two points. For brevity, we use a simplified notation (assuming homogeneous medium) here.

We focus on L_m , which accounts for radiance from the medium,

$$L_m(\mathbf{x}, \vec{\omega}) = \int_0^s \sigma_s(u) T_r(u) L_i(\mathbf{x}_u, \vec{\omega}) du, \quad (2)$$

where $\mathbf{x}_u = \mathbf{x} - u\vec{\omega}$, and σ_s is the scattering coefficient. The in-scattered radiance L_i recursively integrates all the incident light at \mathbf{x}_u multiplied by the normalized phase function f_s :

$$L_i(\mathbf{x}_u, \vec{\omega}) = \int_{\Omega_{4\pi}} f_s(\theta) L(\mathbf{x}_u, \vec{\omega}') d\vec{\omega}', \quad (3)$$

where $\theta = \cos^{-1}(\vec{\omega} \cdot \vec{\omega}')$ and $\Omega_{4\pi}$ is the set of all unit directions.

Photon Mapping. Photon mapping approaches start by shooting “photons” from light sources, scattering the photons within the medium or at surfaces, and recording the resulting scattering locations in a photon map, or the resulting path segments in a photon beams map. During rendering, Equations (2) or (3) are approximated by computing the local density of points or beams.

Virtual Point Lights. VPL methods consider vertices (photons) of the random walk paths and approximate in-scattered indirect radiance as a sum over VPLs. The contribution of a single VPL is:

$$L_i \approx \frac{\Phi f_s(\theta_p) f_s(\theta_u) T_r(w) V}{w^2}, \quad (4)$$

where Φ is the power of the photon (i.e. $\Phi f_s(\theta_p)$ is the intensity of the VPL), w and V are the distance and binary visibility from \mathbf{x}_u to the VPL respectively, $f_s(\theta_p)$ is the phase function evaluated at the VPL, and $f_s(\theta_u)$ accounts for scattering at the evaluation location.

Virtual Ray Lights. We propose to estimate the in-scattered indirect radiance using entire path segments created during photon tracing. The contribution of one VRL is:

$$L_i \approx \Phi \int_0^t \frac{\sigma_s(v) f_s(\theta_v) f_s(\theta_u) T_r(w_u(v)) T_r(v) V_u(v)}{w_u(v)^2} dv, \quad (5)$$

where t is the length of the VRL parametrized by v , and V_u and w_u compute the binary visibility and distance from \mathbf{x}_u to the point v on the VRL. The two phase functions evaluate scattering at the VRL connection θ_v and at the camera connection θ_u . An analogous equation arises for the surface radiance L_s due to a VRL:

$$L_s \approx \Phi \int_0^t \frac{\sigma_s(v) f_s(\theta_v) f_r T_r(w_u(v)) T_r(v) V_u(v)}{w_u(v)^2} dv, \quad (6)$$

where the phase function $f_s(\theta_u)$ has been replaced by the cosine-weighted BRDF f_r . At a high level, Equations (5) and (6) act like a final gather pass [Reichert 1992] over infinitesimally thin photon beams, whereas Equation (4) is a final gather over photon points.

Inserting Equation (5) into Equation (2) we see that we ultimately need to solve a double-integral along each camera ray and VRL:

$$L_m \approx \Phi \int_0^s \int_0^t \frac{\sigma_s(u) \sigma_s(v) f_s(\theta_u) f_s(\theta_v) T_r(u) T_r(v) T_r(w) V}{w(u, v)^2} dv du, \quad (7)$$

where $V(u, v)$ and $w(u, v)$ are the binary visibility and distance between a point u along the camera and a point v along the VRL (we omit these function parameters in the numerator for brevity). We illustrate the geometry of the double integral in Figure 3.




Figure 3: Illustration of the terms involved in Equation (7) for computing the transport between an entire camera ray (green) and a VRL (orange) parametrized by u and v respectively.

Algorithm Overview. Equations (6) and (7) form the basis for our more efficient method for simulating multiple scattering in participating media. We first emit random-walk photon paths from light sources which scatter at surfaces and within the medium. We use these random walk paths with progressive photon beams (PPB) and progressive photon mapping (PPM) to simulate volume and surface caustics. For volume-to-volume transport (multiple scattering) we evaluate Equation (7) for each camera- and photon-path segment using an efficient importance sampling strategy that we describe in Section 4. For volume-to-surface transport (i.e. indirect illumination from the volume) we solve the simpler 1D problem expressed in Equation (6) at the surface hit points. These steps can be trivially repeated in a progressive fashion to provide interactive preview with convergent results.

In order to resolve the visibility present in Equations (6) and (7) in a robust manner, we use ray tracing, as opposed to shadow mapping which introduces bias, and whose performance benefits diminish with increased scene complexity or in the presence of heterogeneous media. Ray tracing also allows us to integrate more readily into existing physically-based rendering frameworks. For example, we support progressive radiance estimation and integrate PPB and PPM to form a fast and general algorithm for computing light transport in and between participating media and surfaces.

As we will analyze mathematically in Section 5 and illustrate with results in Section 6, VRLs provide faster and more robust convergence for indirect light paths when compared to VPL methods and PPB. This is because VRLs act as a form of final gather over photon beams and, by integrating along light and camera rays, the singularities that plague VPL methods become provably weaker.

4 Importance Sampling VRL Transport

Equation (7) defines a 2D integration domain, where one axis is the length u along a camera ray and the other axis is the length v along a VRL. The integrand within this domain incorporates the scattering coefficients, phase functions, and transmittance along both the camera and light rays, and the squared distance, visibility, and transmittance between the two points. We visualize this domain for a single, unoccluded VRL/ray pair in Figures 4 and 7. Unfortunately, a closed-form solution to this double integral is not currently known. As such, we compute this integral using Monte Carlo integration.

If we denote the integrand of Equation 7 as $g(u, v)$, we are interested in evaluating the following unbiased Monte Carlo estimator:

$$L_m(\mathbf{x}, \vec{\omega}) \approx \frac{1}{N} \sum_{i=1}^N \frac{g(u_i, v_i)}{\text{pdf}(u_i, v_i)}, \quad (8)$$

where $\text{pdf}(u_i, v_i)$ is the probability of choosing a point (u_i, v_i) in the domain. To evaluate this estimator efficiently, the PDF should include as many properties of g as possible. We begin with the case of isotropic phase functions, generalizing later to arbitrary phase functions.




Figure 4: Importance sampling multiple scattering within the 2D uv domain for the isotropic (left) and anisotropic (right) Cornell box scenes shown in (c) with all light transport for reference. Current approaches (a) sample according to inverse-squared distance only along one axis, u , and uniformly along the other, v , resulting in suboptimal distributions and noticeable noise. Our analytic method (b) for isotropic scattering accounts for variation along v , providing a better distribution and less noise in the same amount of time as (a). For anisotropic scattering however, our isotropic method (d) fails since significant variation is due to the phase functions. Our generalized product sampling technique (e) faithfully accounts for this more complex density with only a marginal computational overhead (34% for this simple scene).

4.1 Isotropic Scattering

In the isotropic case, inverse-squared distance term causes the most variation in the domain and so we target a pdf $(u_i, v_i) \propto w(u, v)^{-2}$. Unfortunately, sampling according to this PDF using the inversion method is not possible since computing, and then *inverting*, the corresponding CDF eludes analytical computation.

Kulla et al. [2011] recently proposed an equiangular approach for importance sampling a camera ray according to the inverse-squared distance to a point light, and also applied this idea to a randomly sampled square area light. We could trivially apply this approach to our context of a linear light (the VRL) by first uniformly choosing a random location v_i along the VRL and then importance sampling the location u_i along the camera ray according to inverse-squared distance. Unfortunately, by sampling the location v_i on the VRL without regard for the camera ray, the sampling density does not account for variation in $w(u, v)^{-2}$ along the v axis, resulting in a suboptimal distribution. We visualize this in the 2D uv domain in Figure 4 (a) along with a corresponding Cornell box rendering, illustrating the effects of the suboptimal sampling distribution.

To incorporate variation in $w(u, v)^{-2}$ along both v and u , we propose to construct a joint distribution for the entire 2D uv domain and sample it using a two-stage procedure. We first distribute a sample v_i along the VRL using a marginal PDF and then sample a position u_i along the camera ray according to a conditional PDF based on the inverse-squared distance to v_i . In order to make the sampling routine efficient, we strive for an analytic marginal PDF.

We start by applying a change of variables $\hat{u} = u - u_h$ and $\hat{v} = v - v_h$, and similarly for the VRL start and end points \hat{v}_0 and \hat{v}_1 . Here, u_h and v_h are the parameters of the two closest points along the camera ray and VRL, which are separated by a Euclidean distance h (see Figure 3). Using the law of cosines we can now define the squared distance as $w(\hat{u}, \hat{v}, h, \theta)^2 = h^2 + \hat{u}^2 + \hat{v}^2 - 2\hat{u}\hat{v}\cos\theta$, where $\cos\theta$ is the dot-product of the camera and VRL ray directions. The marginal PDF we seek can now be expressed as:

$$\text{pdf}(\hat{v}, \hat{v}_0, \hat{v}_1) = \frac{\int_{\hat{u}_0}^{\hat{u}_1} w(\hat{u}, \hat{v}, h, \theta)^{-2} d\hat{u}}{\int_{\hat{v}_0}^{\hat{v}_1} \int_{\hat{u}_0}^{\hat{u}_1} w(\hat{u}, \hat{v}, h, \theta)^{-2} d\hat{u} d\hat{v}}. \quad (9)$$

To best of our knowledge, there is no analytic solution to such integrals in the current math literature [Gradshteyn and Ryzhik 1994]. Therefore, we opt to simplify the inner integral by assuming the

camera ray is infinite. With this change the numerator evaluates to:

$$\int_{-\infty}^{\infty} w(\hat{u}, \hat{v}, h, \theta)^{-2} d\hat{u} = \frac{\pi}{\sqrt{h^2 + \hat{v}^2 \sin^2 \theta}}, \quad (10)$$

and the normalization term in the denominator evaluates to:

$$\int_{\hat{v}_0}^{\hat{v}_1} \int_{-\infty}^{\infty} w(\hat{u}, \hat{v}, h, \theta)^{-2} d\hat{u} d\hat{v} = \pi \frac{A(\hat{v}_1) - A(\hat{v}_0)}{\sin \theta}, \quad (11)$$

where $A(x) = \sinh^{-1}(\frac{x}{h} \sin \theta)$.

We obtain the marginal CDF by integrating Equation (9):

$$\text{cdf}(\hat{v}, \hat{v}_0, \hat{v}_1) = \frac{A(\hat{v}_0) - A(\hat{v})}{A(\hat{v}_0) - A(\hat{v}_1)}, \quad (12)$$

which can be readily solved for the inverse CDF:

$$\text{cdf}^{-1}(\xi, \hat{v}_0, \hat{v}_1) = \frac{h \sinh(\text{lerp}(A(\hat{v}_0), A(\hat{v}_1), \xi))}{\sin \theta}. \quad (13)$$

We sample a position v_i on the VRL by generating a random number $\xi_{i,1} \in [0, 1)$ and passing it into the inverse CDF. The result can be interpreted as a random VPL located at v_i along the VRL.

In the second step we apply Kulla et al.'s method to obtain a sample location u_i along the camera ray that is distributed according to the inverse-squared distance to v_i . This involves generating another random number $\xi_{i,2} \in [0, 1)$ and inserting it into their inverse CDF:

$$\text{cdf}^{-1}(\xi, \hat{u}_0, \hat{u}_1) = h \tan(\text{lerp}(B(\hat{u}_0), B(\hat{u}_1), \xi)), \quad (14)$$

where $B(x) = \tan^{-1}(x/h)$. Note that this is actually equivalent to the well-known truncated Cauchy distribution. The final PDF is simply the product of the PDFs from these two sampling steps.

We illustrate the effect of incorporating variation along the v axis in Figure 4 (b). Notice that the sample distribution more closely matches the target $w(u, v)^{-2}$ density, which results in less noise in the rendered Cornell box image.

4.2 Anisotropic Scattering

The anisotropic case is unfortunately significantly more complex and, in this case, we aim to sample according to the *product* of the two phase functions as well as the inverse-squared distance:

$$\text{pdf}(u_i, v_i) \propto \frac{f_s(\theta_u) f_s(\theta_v)}{w(u, v)^2}. \quad (15)$$




Figure 5: Importance sampling the phase function (orange) at a point (left) requires generating samples only along the spherical arc (blue) formed by the projection of a ray onto the point. We importance sample the product of two phase functions (green and orange) within this angular domain, viewed (right) within the plane containing the ray and v_i .

In order to illustrate the impact of the anisotropic phase functions, we visualize this density function in Figure 4 (d,e) in the same 2D uv domain as for the isotropic case.

Though importance sampling is possible for many commonly used phase functions, these routines consider the entire spherical domain. Even if we only considered the phase function at a single point light, the samples would need to be constrained to lie along an arbitrary spherical arc (e.g. the projection of the camera ray onto the point as shown in Figure 5 (left)). This type of domain restriction is uncommon hence, to our knowledge, such sampling routines do not currently exist for phase functions. Unfortunately, our context is even more complex since we need to consider not only the 1D camera ray domain, but also the 1D VRL domain and the product of the phase functions aligned along each of these two domains.

A common alternative to sampling the integrand exactly is to use *multiple importance sampling* (MIS) [Veach 1997]. After analyzing and experimenting with this option (as we will discuss further in Section 5), we found that with low sample counts MIS results in significant variance for our problem. We therefore developed a specialized method to directly importance sample the product of the two phase functions and the inverse-squared distance which produces better results than MIS. Our solution is a simple generalization of the two-step isotropic approach described in Section 4.1.

We first proceed as in the isotropic case, choosing a location v_i along the VRL according to the inverse-squared distance to an infinite camera ray (Equation (13)). Note that we ignore the length of the camera ray and phase functions in this first step. The resulting point v_i can be interpreted as an anisotropic VPL along the VRL.

In the second step we wish to sample a location u_i along the camera ray according to the product of the inverse-squared distance and both phase functions (one at v_i aligned to the VRL, and another aligned to the camera ray, see Figure 5 (right)). Working in the angular domain about v_i (along the spherical arc formed by the camera ray’s projection onto v_i , shown in blue), the inverse-squared distance term drops out. We denote the remaining product of the two phase functions as $f_{uv}(\theta(u_i))$ and show a few examples in Figure 6 as solid blue curves. Note that in isotropic cases this product is constant and sampling simplifies to the method in Section 4.1.

For anisotropic scattering, our solution is to construct a compact piecewise-linear PDF which closely approximates f_{uv} . We first evaluate f_{uv} at a fixed number of variably-spaced points, $\theta_1 \dots \theta_M$, along the spherical arc. Then we fit a piecewise linear PDF to these evaluations, and distribute samples by integrating and inverting the corresponding piecewise-quadratic CDF. For this approach to be practical, we must fit f_{uv} accurately (to obtain noise-free results) and efficiently (since this operation is performed for every camera ray/VRL pair during rendering).

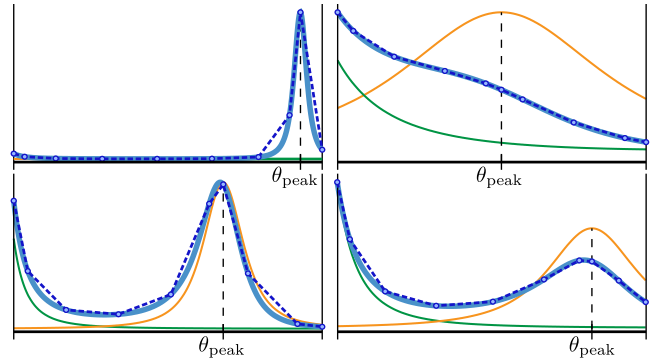


Figure 6: Four example configurations of anisotropic phase functions (Henyey-Greenstein with $g = 0.95$) plotted along the angular domain about a point v_i on the VRL. The product of the phase function along the camera ray (green) and along the VRL (orange) results in the product (solid blue) to which we fit a piecewise-linear PDF (dashed blue). With only 10 vertices, we can reconstruct the product robustly for arbitrary configurations.

We experimented with many settings for M and found that even for very anisotropic phase functions (Henyey-Greenstein with $g = \pm 0.95$), a 10-point piecewise linear fit was sufficient to guarantee high accuracy (see the dashed poly-lines in Figure 6). To efficiently obtain accurate fits we carefully place the vertices to avoid missing important features of f_{uv} , as described in Appendix A.

Figure 4 (e) shows the resulting sample distribution in the uv domain and a corresponding rendering of the Cornell box. Note that while allowing us to handle the product of two arbitrary phase functions, our fitting approach is not much slower than our analytic method for isotropic scattering (Figure 4 (d)). In the Cornell box scene, we incur only a 34% overhead in computing samples, and this overhead becomes negligible as the scene complexity increases.

4.3 VRL to Surface Contribution

Though our primary focus is on multiple scattering, we can apply a largely identical procedure to compute Equation (6), the volume-to-surface transport due to a VRL. In this case, we only have a 1D domain since one of the points is fixed on a surface, and hence we only need the second step of our two-step procedure. Note that this is actually a dual configuration, as we need to sample the location along the VRL for a fixed point on the surface. Aside from this swapping, the only required change is that we replace the phase function of the VPL with the foreshortened BRDF of the surface.

5 Analysis of Sampling and Singularities

Importance Sampling Alternatives. In Section 4 we proposed to importance sample the transport in arbitrary media by first fixing a point on the VRL using our marginal analytic CDF, and then sampling the camera ray according to the product of phase functions. We also considered an alternative: importance sampling according to the transmittance and scattering coefficients along the VRL and the camera ray. Since their product is separable, we can split the 2D PDF into two 1D PDFs (one along each of the rays), avoiding an expensive construction of a numeric PDF for the entire 2D domain.

In Figure 7 we visualize all the individual terms of the integrand from Equation (7) and compare the outcomes of sampling from two PDFs for: (a) the product of phase functions divided by the squared distance, and (b) the product of scattering and transmittance along the VRL and the camera ray. The complete integrand, i.e. the product of (a), (b), and the visibility and transmittance along the connecting segment, is shown in (c). All plots to the left of the dashed line are individually normalized.

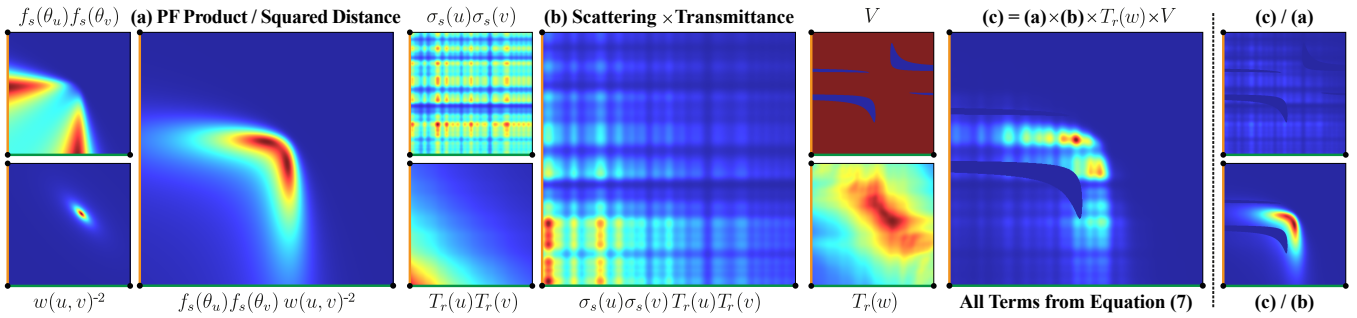


Figure 7: Visualizing terms from Equation (7) and their impact on the integrand for a scene with occluders, heterogeneous media, and anisotropic scattering. We compare importance sampling of (a) the product of phase functions divided by squared distance (our method), to (b) the scattering and transmittance terms along the VRL and camera ray. The complete integrand is shown in (d), and the relative efficiency of sampling using the two PDFs is visualized in the last column, where the same scaling is used for both ratios to allow comparison.




Figure 8: Comparison of sampling in the second stage of our method, i.e. according to the phase functions along the camera ray and the VRL. In the first step, we sampled points along the VRL (very few, for illustration purposes). Multiple importance sampling using the balance heuristic (left) performs much worse than a piecewise-linear approximation (10 points) of the product of both phase functions (center). The right image shows the “reference” solution where the PDF has been evaluated at 1000 points.

The last column depicts the ratios of the integrand to each of the PDFs (i.e. the summand in Equation (8)). With ideal importance sampling this ratio would be constant. For our proposed method (a) the ratio varies only due to the scattering and transmittance, which have relatively low impact on the integrand. In contrast, the ratio of (c) to (b) has much higher variation. This clearly demonstrates that considering the phase function product divided by the squared distance is crucial for efficient numeric evaluation, and supports our choice for importance sampling described in Section 4. We also compare different sampling schemes in the result section.

Multiple Importance Sampling. MIS is a common strategy to importance sample complex integrands by generating samples according to different subsets of the integrand [Veach 1997]. We evaluated the feasibility of an MIS approach for the second stage of our VRL method, i.e. after choosing a point on the VRL. We are not aware of appropriate sampling routines for generating samples along the camera ray according to variation of $f_s(\theta_u)$ towards a fixed point on the VRL, and analogously along a VRL. However, we can use tabulated representations of the two PDFs, which we can both integrate and sample using the inversion method. The results of this experiment, where we randomly chose a strategy and combined the results using MIS, are shown in Figure 8. We tried both the power and balance heuristic, where the latter resulted in less noise. For comparison, Figure 8 (right) shows the result of directly sampling the product (also computed here in a brute force tabulated fashion). This constitutes the best we could possibly hope to achieve. We see that combining strategies using MIS, which may initially seem like a good choice, results in significantly more noise than sampling the product directly. This is because the effective PDF of MIS is actually a linear combination of the two PDFs, and




Figure 9: We analyze four different cases of energy transfer; for illustration we use an equal number (10k) of VPLs and VRLs to render a sphere in participating media. The left column shows surface illumination only. Top row: light transport computed using VPLs; bottom row: VRLs illuminating surfaces and media.

not the product (Figure 6). Furthermore, analytic sampling would require deriving an efficient arc sampling technique for every desired phase function. Due to these drawbacks, we instead resort to the more general but fast approach for sampling the product directly, without the need for MIS.

Singularities. Similar to VPLs, we obtain VRLs directly from photon tracing. However, the distribution of the energy along VRLs suggests that singularities should be less pronounced compared to VPLs where the energy is concentrated at a single point. In the spirit of Jarosz et al. [2011a] we can distinguish four cases of energy transfer that we analyze: point-to-point (P2P, VPLs illuminating surfaces, Figure 9a), point-to-line (P2L, VPLs illuminating participating media, Figure 9b) and its dual line-to-point (L2P, VRLs illuminating surfaces, Figure 9c), and line-to-line (L2L, VRLs illuminating participating media, Figure 9d).

In Appendix B we derive the order of the respective singularities of energy transfer. We observe that P2P transfer can expectedly cause the strongest singularities, while L2P and P2L result in significantly less pronounced singularities. We also show that singularities in L2L transfer are lower still, except for the special case where both lines are parallel. Figure 9 shows the singularities when rendering

a sphere embedded in a homogeneous participating media using VPLs and VRLs. It can be seen that VPLs in media (top-right) and VRLs for surface lighting (bottom-left) cause less distracting artifacts than VPLs with surfaces (top-left); when using VRLs for participating media (bottom-right) the singularities are not visible at all. In the same spirit as most VPL methods, we could simply clamp the VRLs’ contributions to avoid any singularities. As we have seen, these are generally lower than for VPLs, and this means that clamping would remove less energy from the solution, and consequently also has less impact on material appearance [Křivánek et al. 2010]. However, we decided not to clamp in all our renderings, as the singularities vanish quickly when the number of VRLs increases. We leave the development of bias compensation techniques for future work.

6 Algorithm and Results

Light Path Splitting. The best choice of methods for computing light transport varies with the respective characteristics of the individual phenomena. For example, progressive photon beams (PPB) are well suited for rendering volumetric caustics, but have no significant advantage over traditional photon mapping for volume-to-surface transport (the endpoints of the beams are simply surface photons). VRLs, on the other hand, yield significantly better results for multiple scattering and require less virtual light sources (than VPLs) for surface lighting. Fortunately, both PPB and VRLs can be created from the same photon tracing step and naturally fit complementarily into one rendering framework.

Implementation. We implemented our algorithm in a hybrid CPU-GPU framework. We first trace random-walk paths from light sources, scattering photons at surfaces and in the media using a CPU ray-tracer. This step corresponds to photon shooting in standard photon mapping [Jensen and Christensen 1998; Jensen 2001]. Photons form VPLs or VRLs where, in the latter case, the photon path segments are used instead of the photon locations. In addition to VRL radiance estimation, we apply PPB to the photon segments and PPM to the segment endpoints to simulate volume and surface caustics. As proposed by Jarosz et al. [2011b], we render directly visible single-scattered volume caustics with rasterization and use CPU ray-tracing to handle beams only visible after specular reflection and/or refraction (e.g. all lighting in the glasses in Figure 1).

For volume-to-volume transport (multiple scattering), we evaluate Equation (7) for each VRL/camera ray pair (all of which are uploaded to the GPU) by sampling according to our novel product importance sampling routine (Section 4) and resolving visibility using Aila and Laine’s [2009] efficient GPU ray-tracer. We use the same framework when evaluating volume-to-volume transport with VPLs, and both techniques benefit from the same acceleration structures and similar code paths. We furthermore use Russian roulette, based on the minimum distance between each camera/light ray pair (or surface point/light ray for VPLs), to probabilistically prune virtual lights with low contribution.

Evaluating the transmittances in Equations (6) and (7) can easily become a bottleneck in the case of heterogeneous medium. We improve the performance by precomputing the transmittance along camera rays and VRLs: using Woodcock tracking we sample and cache a number of distances (typically 16) that are later used to approximate the transmittance along the ray in an unbiased manner [Jarosz et al. 2011b]. Notice that transmittance between sample points on the VRL and the sample points on the ray (or the surface point) cannot be cached. In this case we use either fewer Woodcock samples or employ fast ray-marching, which improves the performance at the cost of introducing a small and imperceptible amount of bias into the transmittance computation.

We handle volume-to-surface transport by solving the simpler 1D problem (Equation (6)) at the surface hit points. We wrap all of the methods in a progressive framework, providing interactive previews that converge to ground truth all in the same renderer. Our results show all the energy in the scene, i.e. no clamping is used.

Results. We compare the quality and performance of our algorithm against VPLs and PPB on an Intel Core i7 CPU @ 2.8GHz with 8GB RAM and an NVIDIA GTX 470. The CPU ray-tracer and density estimation for PPM and PPB are parallelized over all CPU cores. All our results use a Henyey-Greenstein phase function.

Figure 1 shows the FRUIT JUICE scene with glasses of orange and grapefruit juice rendered at 512×512 resolution. We simulate anisotropic scattering for both materials. The media parameters are $\sigma_a = (0.41, 0.72, 5.18)\text{cm}^{-1}$, $\sigma_s = (0.36, 0.50, 0.18)\text{cm}^{-1}$, and $g = 0.5$ for the orange juice; and $\sigma_a = (0.41, 0.95, 4.73)\text{cm}^{-1}$, $\sigma_s = (0.45, 0.32, 0.23)\text{cm}^{-1}$, and $g = 0.6$ for the grapefruit juice.

We show an equal time comparison for volume-to-volume transport computed using VRLs, PPB, and VPLs. Note that PPB is a biased (but consistent) method and yields darker results prior to convergence; VPLs (without clamping) suffer from bright pixels, even after long render times. In contrast, our method produces unbiased results and does not suffer from concentrated bright pixels. For the final image (left), we combine our multiple scattering result with a single scattering (volume caustic) component using PPB, and surface illumination and surface caustics, computed using PPM.

The SMOKY ROOM scene (Figures 10, 11, and 12) is filled with a heterogeneous isotropic medium (computed using Perlin noise) whose density decreases with height. Figure 10 shows volume-to-volume transport only, i.e. paths whose last two interactions before reaching the camera happened in the medium. At equal time, VRLs yield higher quality results than both PPB and VPLs. Despite the fact that we accelerate PPB using hardware rasterization, it requires millions of beams to sufficiently cover the entire image, and the convergence rate is thus low. PPB works extremely well for volumetric caustic illumination since the beams are concentrated, reducing fill rate, while ensuring high density. For multiple scattering, photon paths quickly become incoherent, which introduces low-frequency noise into the local density estimate. VRLs do not rely solely on local density information and can therefore obtain higher quality multiple scattering using significantly less beams. In contrast to VPLs, our method does not suffer from visible artifacts due to the singularities, trading this for a slight uniform noise distributed evenly across the image. This is because we sample each VRL/ray pair with only one random sample. The final solution consists of volume-to-volume (100s) and volume-to-surface (600s) illumination computed using VRLs, and single scattering (100s) and surface-to-volume (300s) light transport computed using PPB.

Figure 11 shows volume-to-surface illumination, i.e. light paths with the last two scattering events being in volume and then on a surface. When using VRLs, the volume-to-surface transport converges faster than with VPLs, and very little, if any, clamping is required, while it does not seem to be necessary at all for volume-to-volume transport. This is an important improvement over VPLs, since clamping removes energy and changes the appearance of materials and media. Bias compensation techniques can recover the lost energy, but at significant additional expense.

In Figure 12 we compare different sampling strategies for evaluating the VRL-ray transport using a single sample. For each VRL and camera ray we construct a 1D piecewise linear PDF (with 100 vertices) for the transmittance (\times the scattering coefficient, cf. Figure 7(b)) along the ray, and sample each ray independently. This leads to much higher variance than sampling according to the product of phase functions divided by the squared distance (cf. Fig-




Figure 10: For the SMOKY ROOM scene (left) we compare the volume-to-volume transport computed using VRLs, PPB, and VPLs at equal time. While PPB would require many more photon beams to sufficiently fill the scene, and VPLs suffer from singularities, our method provides artifact-free results faster, benefiting from explicit gathering from the light path segments and denser sampling of the space, respectively.




Figure 11: These images show volume-to-surface transport in the SMOKY ROOM scene computed in only 600s; some regions are dark also due to absorption and out-scattering in the participating media. The close-ups compare the singularities of VRLs and VPLs (both without clamping).

Table 1: Performance breakdown of volume-to-volume transport.

	FRUIT JUICE ¹	SMOKY ROOM ²
Photon (VRL) shooting	6%	2%
Ray tracing ¹ / Rasterization ²	55%	1%
Importance sampling	9%	4%
Visibility test	23%	52%
Light transport	7%	42%

ure 7(a)), which has significant impact on the integrand from Equation 7. Admittedly, there can be situations when sampling according to the scattering \times transmittance can also improve the efficiency. For cases where this additional expense is warranted, the individual sampling strategies can be combined using MIS.

In the supplemental video, we compare the convergence behavior of our progressive algorithm to both PPB and VPLs. Another common problem with VPL techniques is temporal flickering due to stochastic under-sampling the indirect lighting. The increased sample density provided by VRLs significantly diminishes these artifacts even when using fewer virtual lights.




Figure 12: Importance sampling strategies for the VRL/ray transport: sampling each of the rays independently according to the transmittance (\times scattering) suffers from high noise. Sampling according to the phase function product divided by the squared distance (our method) leads to significantly less noise, verifying the theoretical analysis from Figure 7.

In Table 1 we show the performance breakdown for computing volume-to-volume transport. Our importance sampling approach is efficient, occupying between 4–9% of the total render time whereas 23–52% of the time is spent evaluating visibility on the GPU, which would be even higher if performed on the CPU.

7 Discussion and Future Work

VRLs are a light transport representation well-suited for computing unbiased multiple scattering and volume-to-surface lighting. They can be easily combined into a single framework with PPB for rendering single scattering and in particular volume caustics.

Analytic Integration, Duality. Equations (5) and (6) can also be seen as the dual of the “airlight integral,” for which closed-form solutions exist when constrained to homogeneous media and no occluders [Sun et al. 2005; Pegoraro and Parker 2009]. We have considered leveraging these analytic methods, however, we found that each analytic integration is quite expensive (especially when incorporating anisotropic scattering [Pegoraro et al. 2009; Pegoraro et al. 2010]), and ultimately the homogeneous/visibility assumptions are too restrictive for the general setting we consider. Furthermore, by solving Equation (5) in isolation we would actually be considering only a small portion of a larger 2D integration problem.

General Bidirectional Algorithm. We believe that our idea of sampling transport between rays is an important first step towards developing new bidirectional MC approaches for rendering in the presence of participating media. At the moment, we consider only

camera rays and specular camera paths. Formulating a general bidirectional algorithm evaluating transport between arbitrary segments of light and camera paths is an interesting area for future work.

Visibility. In our implementation we did not use shadow mapping, which is often used to accelerate VPL rendering, as this introduces bias and prevents a general rendering framework. Moreover, standard shadow maps cannot be used for heterogeneous media. Nevertheless, dedicated shadow techniques for linear lights exist [Heidrich et al. 2000] and seem to be worth investigating. Furthermore, incorporating mutual visibility into sampling could possibly speed-up the convergence for highly occluded scenes. We believe that our approach could possibly be combined with ideas from Georgiev et al. [2012], though this would warrant its own in-depth investigation.

Singularities. We have seen that singularities are virtually unnoticeable when using VRLs for volume rendering, yet they can remain visible on surfaces, however, even in this case they are less noticeable than singularities on surfaces due to VPLs. This empirically supports our proof of progressively lower order of singularities. Though we are able to obtain unbiased images using relatively few VRLs, convergence is slower for volume-to-surface transport. We believe that if a little amount of bias is acceptable, we could obtain artifact-free images faster by spatially blurring VRL energy.

Clustering of VRLs. Lightcuts clusters millions of VPLs and hierarchically chooses the best representatives for each gather point. Though we do not explicitly cluster virtual lights, each VRL/camera ray pair can be seen as a continuous family of light and gather point pairs. By importance sampling the resulting 2D domain, we are effectively choosing the best representative from this family for each location along a camera ray. Nonetheless, it would be interesting to see if explicit clustering of VRLs, in the spirit of Lightcuts [Walter et al. 2005; Walter et al. 2006], is possible.

Reconstruction with Sample Reuse. Our 2D uv domain shares similarities with epipolar slices. This geometric interpretation has previously been used to accelerate volumetric shadows [Engelhardt and Dachsbacher 2010; Baran et al. 2010], and sampling and reconstruction of motion blur and depth-of-field [Hachisuka et al. 2008a]. It would be interesting to consider more sophisticated, but perhaps biased reconstruction of the 2D uv slice, or whether occluders could be directly rasterized into this warped domain.

8 Conclusion

We presented a new lighting primitive — virtual ray lights — for unbiased many-light rendering of indirect illumination in, and from, participating media. VRLs are created from path segments of photon random walks in the medium, and we showed how to compute their contribution to entire camera rays through the medium and to surfaces. VRLs are less prone to singularities and with our method we obtain higher-quality multiple-scattering than previously existing techniques.

Acknowledgements. We would like to thank Marios Papas, Ilya Baran, Steve Marschner, and Ladislav Kavan for feedback on an early draft of the paper; Benjamin Müller for initial investigations and experiments; and Krzysztof Jarosz for discussions relating to the singularity analysis in Appendix B.

References

AILA, T., AND LAINE, S. 2009. Understanding the efficiency of ray traversal on gpus. In *Proc. of High Performance Graphics*.

ARBREE, A., WALTER, B., AND BALA, K. 2008. Single-pass scalable subsurface rendering with lightcuts. *Computer Graphics Forum* 27, 2.

BARAN, I., CHEN, J., RAGAN-KELLEY, J., DURAND, F., AND LEHTINEN, J. 2010. A hierarchical volumetric shadow algorithm for single scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 29, 5.

CEREZO, E., PEREZ-CAZORLA, F., PUEYO, X., SERON, F., AND SILLION, F. 2005. A survey on participating media rendering techniques. *The Visual Computer* 21.

CHANDRASEKAR, S. 1960. *Radiative Transfer*. Dover Publications.

DAMMERTZ, H., KELLER, A., AND LENSCH, H. P. A. 2010. Progressive point-light-based global illumination. *Computer Graphics Forum* 29, 8.

DONG, Z., GROSCH, T., RITSCHEL, T., KAUTZ, J., AND SEIDEL, H.-P. 2009. Real-time indirect illumination with clustered visibility. In *Proc. Vision, Modeling, and Visualization Workshop*.

DONNER, C., AND JENSEN, H. W. 2005. Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 24, 3.

ENGELHARDT, T., AND DACHSBACHER, C. 2010. Epipolar sampling for shadows and crepuscular rays in participating media with single scattering. In *Symposium on Interactive 3D Graphics and Games*.

ENGELHARDT, T., NOVÁK, J., AND DACHSBACHER, C. 2010. Instant multiple scattering for interactive rendering of heterogeneous participating media. Tech. rep., Karlsruhe Institute of Technology, Dec.

GEORGIEV, I., KŘIVÁNEK, J., POPOV, S., AND SLUSALLEK, P. 2012. Importance caching for complex illumination. *Computer Graphics Forum (Proc. of Eurographics)* 31, 2.

GRADSHTEYN, I. S., AND RYZHIK, I. M. 1994. *Table of Integrals, Series, and Products, Fifth Edition*. Academic Press.

HACHISUKA, T., AND JENSEN, H. W. 2009. Stochastic progressive photon mapping. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 28, 5.

HACHISUKA, T., JAROSZ, W., WEISTROFFER, R. P., DALE, K., HUMPHREYS, G., ZWICKER, M., AND JENSEN, H. W. 2008. Multidimensional adaptive sampling and reconstruction for ray tracing. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 27, 3.

HACHISUKA, T., OGAKI, S., AND JENSEN, H. W. 2008. Progressive photon mapping. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 27, 5.

HAVRAN, V., BITTNER, J., HERZOG, R., AND SEIDEL, H.-P. 2005. Ray maps for global illumination. In *Proc. of Eurographics Symposium on Rendering*.

HAŠAN, M., PELLACINI, F., AND BALA, K. 2007. Matrix row-column sampling for the many-light problem. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 26, 3.

HAŠAN, M., KŘIVÁNEK, J., WALTER, B., AND BALA, K. 2009. Virtual spherical lights for many-light rendering of glossy scenes. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 28, 5.

HEIDRICH, W., BRABEC, S., AND SEIDEL, H.-P. 2000. Soft shadow maps for linear lights. In *Proc. of Eurographics Workshop on Rendering*.

HERZOG, R., HAVRAN, V., KINUWAKI, S., MYŠKOWSKI, K., AND SEIDEL, H.-P. 2007. Global illumination using photon ray splatting. *Computer Graphics Forum (Proc. of Eurographics)* 26, 3.

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., NOWROUZEZHRAI, D., SADEGHI, I., AND JENSEN, H. W. 2011. A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 30, 1.
- JAROSZ, W., NOWROUZEZHRAI, D., THOMAS, R., SLOAN, P.-P., AND ZWICKER, M. 2011. Progressive photon beams. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 30, 6.
- 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*.
- JENSEN, H. W. 2001. *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters, Ltd., Natick, MA, USA.
- KAJIYA, J. T. 1986. The rendering equation. In *Computer Graphics (Proc. of SIGGRAPH)*.
- KELLER, A. 1997. Instant radiosity. In *Proc. of SIGGRAPH '97*.
- KNAUS, C., AND ZWICKER, M. 2011. Progressive photon mapping: A probabilistic approach. *ACM Transactions on Graphics* 30, 3.
- KOLLIG, T., AND KELLER, A. 2006. Illumination in the presence of weak singularities. In *Monte Carlo and Quasi-Monte Carlo Methods 2004*.
- KULLA, C., AND FAJARDO, M. 2011. Importance sampling of area lights in participating media. In *ACM SIGGRAPH 2011 Talks*.
- KŘIVÁNEK, J., FERWERDA, J. A., AND BALA, K. 2010. Effects of global illumination approximations on material appearance. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 29, 4.
- LAFORTUNE, E. P., AND WILLEMS, Y. D. 1993. Bi-directional path tracing. In *Computographics '93*.
- LAFORTUNE, E. P., AND WILLEMS, Y. D. 1996. Rendering participating media with bidirectional path tracing. In *Proceedings of the eurographics workshop on Rendering techniques*.
- LAINE, S., SARANSAARI, H., KONTKANEN, J., LEHTINEN, J., AND AILA, T. 2007. Incremental instant radiosity for real-time indirect illumination. In *Proc. of Eurographics Symposium on Rendering*.
- LASTRA, M., UREÑA, C., REVELLES, J., AND MONTES, R. 2002. A particle-path based method for monte carlo density estimation. In *Proc. of Eurographics Workshop on Rendering*.
- MOON, J. T., AND MARSCHNER, S. R. 2006. Simulating multiple scattering in hair using a photon mapping approach. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, ACM, NY.
- NOVÁK, J., ENGELHARDT, T., AND DACHSBACHER, C. 2011. Screen-space bias compensation for interactive high-quality global illumination with virtual point lights. In *Proc. of ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*.
- OU, J., AND PELLACINI, F. 2011. Lightslice: Matrix slice sampling for the many-lights problem. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 30, 6.
- PEGORARO, V., AND PARKER, S. G. 2009. An analytical solution to single scattering in homogeneous participating media. *Computer Graphics Forum (Proc. of Eurographics)* 28, 2.
- 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*.
- PEGORARO, V., SCHOTT, M., AND PARKER, S. G. 2010. A closed-form solution to single scattering for general phase functions and light distributions. *Computer Graphics Forum (Proc. of Eurographics Symposium on Rendering)* 29, 4.
- RAAB, M., SEIBERT, D., AND KELLER, A. 2008. Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer.
- REICHERT, M. C. 1992. *A Two-pass Radiosity Method Driven by Lights and Viewer Position*. Master's thesis, Cornell University.
- RITSCHEL, T., GROSCHE, T., KIM, M. H., SEIDEL, H.-P., DACHSBACHER, C., AND KAUTZ, J. 2008. Imperfect shadow maps for efficient computation of indirect illumination. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 27, 5.
- 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. 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. SIGGRAPH)* 29, 4.
- VEACH, E., AND GUIBAS, L. 1994. Bidirectional estimators for light transport. In *Proc. of Eurographics Rendering Workshop*.
- VEACH, E. 1997. *Robust Monte Carlo methods for light transport simulation*. PhD thesis, Stanford, CA, USA.
- WALD, I., KOLLIG, T., BENTHIN, C., KELLER, A., AND SLUSALLEK, P. 2002. Interactive global illumination using fast ray tracing. In *Proc. of Eurographics Workshop on Rendering*.
- WALTER, B., FERNANDEZ, S., ARBREE, A., BALA, K., DONIKIAN, M., AND GREENBERG, D. P. 2005. Lightcuts: a scalable approach to illumination. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 24, 3.
- WALTER, B., ARBREE, A., BALA, K., AND GREENBERG, D. P. 2006. Multidimensional lightcuts. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3.

A Piecewise Linear PDF Construction

We first evaluate f_{uv} at the two ends of the spherical arc to obtain vertices θ_1 and θ_M . We distribute the remaining $\theta_2 \dots \theta_{M-1}$ interior vertices non-uniformly. Most importantly, we must avoid missing the peak of the underlying function. Though we do not know the exact location of the peak, we do know that the function is a product of two phase functions and that, of these, $f_s(\theta_u)$ must be monotonic between θ_1 and θ_M . Hence, the peak is primarily determined by the remaining phase function, $f_s(\theta_v)$. If the direction θ_{peak} maximizing $f_s(\theta_v)$ lies on the arc (i.e. between θ_1 and θ_M), we can easily find it as:

$$\theta_{\text{peak}} = \cos^{-1}(\vec{a} \cdot \vec{e}) \quad (16)$$

$$\vec{e} = \text{normalize}((\vec{c} \times \vec{d}) \times \vec{c}), \quad (17)$$

$$\vec{c} = \vec{a} \times \vec{b} \quad (18)$$

where \vec{a} and \vec{b} are the directions towards the origin and the end of the ray, and \vec{d} is the direction of the VRL. If the peak does not fall within $[\theta_1, \theta_M]$, we invert \vec{d} and repeat the above procedure to find the possible minimum of $f_s(\theta_v)$ (negative peak). If neither the peak, nor the minimum fall within $[\theta_1, \theta_M]$, we distribute the




Figure 13: The canonical geometric configuration for the three different transport types.

interior vertices with a cosine-warped uniform spacing¹:

$$\theta_j = \frac{\theta_M - \theta_1}{2} \left(1 - \cos \left(\frac{\pi(j-1)}{M} \right) \right). \quad (19)$$

If $\theta_{\text{peak}} \in [\theta_1, \theta_M]$ then we place the vertex with index

$$j_{\text{peak}} = \left\lfloor \frac{\theta_{\text{peak}} - \theta_1}{\theta_M - \theta_1} M - 0.5 \right\rfloor \quad (20)$$

at θ_{peak} , and we distribute the remaining vertices to either side using a similar cosine-warped distribution as in Equation 19, but within each subinterval.

This procedure ensures that we always sample the expected location of the peak, and distribute the remaining vertices where the function is expected to be varying the most rapidly.

B Singularity Analysis

We analyze the singularities present in VPL and VRL light transport. There are three types of transfer we are interested in comparing: point-to-point, point-to-line, and line-to-line as highlighted in the renderings in Figure 9 and illustrated in Figure 13. Without loss of generality, we assume a canonical geometric configuration and eliminate terms which do not affect limit behavior as the distance between the elements approaches 0. This corresponds to the following three equations:

$$f_{\text{P2P}}(w) = \frac{1}{w^2} \quad (21)$$

$$f_{\text{P2L}}(w) = \int_0^1 \frac{1}{u^2 + w^2} du = \frac{1}{w} \tan^{-1} \left(\frac{1}{w} \right) \quad (22)$$

$$f_{\text{L2L}}(w, \theta) = \int_0^1 \int_0^1 \frac{1}{(v \cos \theta - u)^2 + (v \sin \theta)^2 + w^2} dudv \quad (23)$$

In order, these functions correspond to: the contribution of a VPL to a *point* on a surface/volume; the contribution of a VRL to a *point* on a surface/volume, or the contribution of a VPL integrated along a camera *ray*; and, the contribution of a VRL integrated along a camera *ray*. Note that each successive function integrates over an additional linear domain. We plot these functions in Figure 14 (left).

We will prove that $f_{\text{L2L}}(w) \in o(f_{\text{P2L}}(w))$, and $f_{\text{P2L}}(w) \in o(f_{\text{P2P}}(w))$ as $w \rightarrow 0$; where for non-zero functions f and g :

$$f(x) \in o(g(x)) \text{ as } x \rightarrow 0 \iff \lim_{x \rightarrow 0} \frac{f(x)}{g(x)} = 0. \quad (24)$$

More informally, we will show that the singularity at $w = 0$ for VPLs on surfaces dominates that of VRLs at surfaces, which dominates that of VRLs integrated along a camera ray in a volume.

¹Applying a cosine to equally-spaced points pushes more points towards the boundary of the interval, where we expect the most variation.

B.1 Point-to-Line vs Point-to-Point Transfer

Plugging in Equations (21) and (22) into Equation (24) we have:

$$\lim_{w \rightarrow 0} \frac{f_{\text{P2L}}(w)}{f_{\text{P2P}}(w)} = \lim_{w \rightarrow 0} \left[w \tan^{-1} \left(\frac{1}{w} \right) \right] = 0. \quad (25)$$

Hence, $f_{\text{P2L}}(w) \in o(f_{\text{P2P}}(w))$: the point-to-line contribution is dominated by the point-to-point contribution as the distance w goes to 0. This limit is illustrated as the blue curve in Figure 14 (right).

Furthermore, since

$$\lim_{w \rightarrow 0} \frac{f_{\text{P2L}}(w)}{w} = \frac{\pi}{2}, \quad (26)$$

we know that $f_{\text{P2L}}(w) \in \Theta(1/w)$ as $w \rightarrow 0$; or, in other words, the point-to-line singularity is bounded from above and below by $1/w$ asymptotically. Note that this is in contrast to the point-to-point singularity, which is of order $\Theta(1/w^2)$.

B.2 Line-to-Line vs Point-to-Line Transfer

The inner integral of Equation (23) has the closed form solution:

$$\frac{\sqrt{2} \left(\tan^{-1} \left(\frac{\sqrt{2}v \cos(\theta)}{d(v,w,\theta)} \right) - \tan^{-1} \left(\frac{\sqrt{2}(v \cos(\theta) - 1)}{d(v,w,\theta)} \right) \right)}{d(v,w,\theta)}, \quad (27)$$

where $d(v,w,\theta) = \sqrt{v^2(1 - \cos 2\theta) + 2w^2}$. Since the numerator is bounded from above by $\sqrt{2}\pi$, we have

$$f_{\text{L2L}}(w, \theta) \leq \sqrt{2}\pi \int_0^1 \frac{1}{d(v,w,\theta)} dv \quad (28)$$

$$= \pi \csc \theta \operatorname{csch}^{-1}(w \csc \theta), \quad (29)$$

and, for any fixed $\theta \in (0, \frac{\pi}{2}]$, we can obtain

$$\lim_{w \rightarrow 0} \frac{f_{\text{L2L}}(w, \theta)}{f_{\text{P2L}}(w)} \leq \lim_{w \rightarrow 0} \frac{\pi \csc \theta \operatorname{csch}^{-1}(w \csc \theta)}{\frac{1}{w} \tan^{-1} \left(\frac{1}{w} \right)} = 0. \quad (30)$$

Hence, $f_{\text{L2L}}(w, \theta) \in o(f_{\text{P2L}}(w))$: the line-to-line contribution is dominated by the point-to-line contribution as the distance w goes to 0. This limit is illustrated as the red curve in Figure 14 (right).

Following the same simplification, we can see that

$$\lim_{w \rightarrow 0} \frac{f_{\text{L2L}}(w, \theta)}{-\ln(w)} \leq \frac{\pi}{\sin \theta}, \quad (31)$$

which is non-negative and finite for any fixed $\theta \in (0, \frac{\pi}{2}]$ and hence, $f_{\text{L2L}}(w, \theta) \in O(-\ln(w))$ as $w \rightarrow 0$; or, in other words, the line-to-line singularity is bounded from above by $-\ln(w)$.

For the special case where the two lines approach being parallel, the term in the numerator of Equation (30) is bounded by a constant: $\lim_{\theta \rightarrow 0} f_{\text{L2L}}(w, \theta) \leq \frac{\pi}{w}$. Hence, in this special case it can be shown that $f_{\text{L2L}}(w, 0) \in \Theta(f_{\text{P2L}}(w))$ as $w \rightarrow 0$: the singularity of f_{L2L} becomes asymptotically equivalent to that of f_{P2L} .




Figure 14: Asymptotic behavior of the point-to-point, point-to-line and line-to-line contributions (left), and (right) limits of the ratios show that the singularities become successively weaker.