Eurographics Symposium on Rendering - Experimental Ideas & Implementations (2017)
P. Sander and M. Zwicker (Editors)

Extended Path Integral Formulation for Volumetric Transport

Toshiya Hachisuka[1] Iliyan Georgiev[2] Wojciech Jarosz[3] Jaroslav Křivánek[4] Derek Nowrouzezahrai[5]

[1]The University of Tokyo [2]Solid Angle [3]Dartmouth College [4]Charles University in Prague [5]McGill University

Abstract

*We propose an extension of the path integral formulation amenable to the expression of volumetric light transport with photon beam estimates. Our main contribution is a generalization of Hachisuka et al.'s extended path space formulation [HPJ12] to light transport in participating media. Our formulation supports various point- and beam-based volumetric density estimators, unifying them with path integration in the spirit of the work by Křivánek et al. [KGH*14]. One unique and useful property of our formulation is that it recasts beam-based density estimation as Monte Carlo path vertex sampling in a higher-dimensional space, rather than beam merging in a lower-dimensional space, which enables a practical algorithm for beam estimators with 3D-blur kernels. We thus establish a complementary theoretical foundation for the development of rendering algorithms using points, beams, and paths in participating media.*

1. Introduction

Light transport simulation in computer graphics is built atop a recursive integral equation that describes the steady state distribution of radiance at every point in an input virtual scene. This recursive equation can be transformed into a conceptually simpler integral over light transport paths, with vertices lying on objects' surfaces [Vea97] or in volumetric participating media [PKK00]. This path integral formulation allows expressing many different Monte Carlo rendering algorithms under a unified theoretical framework and, in turn, combining them using multiple importance sampling [Vea97].

While the path integral formulation precisely describes algorithms based on Monte Carlo path integration—such as path tracing [Kaj86] and bidirectional path tracing [LW93, VG94]—this formulation is known to be incompatible with algorithms based on photon density estimation [Jen96]. Georgiev et al. [GKDS12] and Hachisuka et al. [HPJ12] concurrently proposed unified formulations for these two classes of algorithms for the case of surface light transport.

Georgiev et al. kept the original path integral formulation, and interpreted photon density estimation as a conceptual merging of subpath vertices, called *vertex merging*. This approach corresponds to a contraction of the path space—a vertex on a light subpath is merged with a nearby vertex on an eye subpath, making full light transport paths in photon density estimation have the same number of vertices as regular light transport paths.

In contrast, Hachisuka et al. proposed an *extended path integral formulation* that handles photon density estimation paths as disconnected segments, without merging. This approach corresponds to an expansion of the original path space since disconnected segments have more vertices than regular light transport paths. More recently, Křivánek et al. [KGH*14] applied the concept of vertex merging to render volumetric scattering by combining photon point- and beam-based density estimators with path integration methods [JNSJ11]. An extended path integral formulation for volumetric light transport, however, has not been available.

We address this theoretical gap and propose an extended path integral formulation for light transport in scattering participating media. Our formulation builds atop Hachisuka et al.'s [HPJ12] idea of *disconnected* path segments. Similarly to its surface counterpart [HPJ12], our formulation can precisely express photon density estimation as a path integral in an extended (volumetric) path space. This formulation allows us to express beam-based density estimation as Monte Carlo vertex sampling in the extended path space.

The focus of this paper is to establish a theoretical foundation, rather than introducing a new rendering algorithm. In fact, for the lowest-dimensional blur kernels, our formulation can be realized with an implementation that is equivalent to the existing merging formulation [KGH*14], and similarly allows combination of path integration and photon density estimation. Our formulation, however, explains how beam estimators can be leveraged with higher-dimensional blur kernels in the extended path space. Křivánek et al. [KGH*14] noted that such higher-dimensional blurs can be mapped to their formulation using Monte Carlo integration over auxiliary variables for additional dimensions, but they left the details of this exploration to future work. Our formulation explains how this additional integration appears naturally in an extended volumetric path space, which also leads to a practical algorithm for beam estimators with 3D-blur kernels, based on classic Monte Carlo integration. Higher-dimensional blur kernels can be more convenient than the lowest-dimensional ones, e.g. for removing the singularities in beam estimators and for consistently using the same 3D blurred path throughput across all estimators with points, beams, and paths.

DOI: 10.2312/sre.20171195

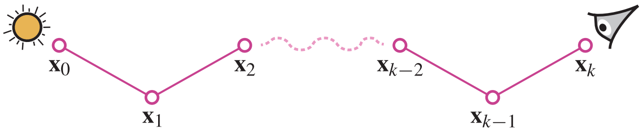

Figure 1: *Regular path for length k. Vertices are fully connected from the sensor to the light sources.*

2. Path Integral Formulation for Volumes

This section briefly reviews the path integral framework. For the sake of brevity, but without loss of generality, we focus on the case where all path vertices are in volumes (i.e. no surface-volume light transport). Readers familiar with this concept can jump to Section 3 where we derive and discuss our extension.

The path integral formulation of light transport expresses a pixel measurement as a multidimensional integral [Vea97, PKK00]:

$$I_j = \int_\Omega f_j(\bar{\mathbf{x}})\, d\mu(\bar{\mathbf{x}}), \tag{1}$$

where j is the pixel index, f_j is a measurement contribution function, $\bar{\mathbf{x}}$ is a light transport path, and μ is the measure of the space of all possible paths of all lengths Ω.

Paths. A path $\bar{\mathbf{x}}$ of length $k \in [1,\infty]$ is a sequence of vertices $\mathbf{x}_i \in \mathcal{V}$:

$$\bar{\mathbf{x}} = \mathbf{x}_0\mathbf{x}_1 \ldots \mathbf{x}_k, \tag{2}$$

where vertex $\mathbf{x}_0$ is on a light source and $\mathbf{x}_k$ is on the eye sensor, and $\mathcal{V}$ is the three-dimensional Euclidean space (see Figure 1).

Path space. Let Ω_k represent the space of length-k paths. The full path space Ω is the union of the sets of paths of all possible lengths:

$$\Omega = \bigcup_{k=1}^{\infty} \Omega_k. \tag{3}$$

Path measure. The differential measure $d\mu_k$ of a length-k path is the product of the standard differential volume measures dV of all its $k+1$ vertices:

$$d\mu_k(\bar{\mathbf{x}}) = dV(\mathbf{x}_0)dV(\mathbf{x}_1)\cdots dV(\mathbf{x}_k), \tag{4}$$

which leads to

$$\mu_k = \underbrace{V\times\cdots\times V}_{k+1 \text{ times}} = V_{k+1}. \tag{5}$$

The measure μ on the path space Ω is the sum of the measures of the paths of each length:

$$\mu(D) = \sum_{k=1}^{\infty} \mu_k(D\cap\Omega_k), \tag{6}$$

where $D \subseteq \Omega$.

Measurement contribution function. The path measurement contribution function f_j for volumes is defined as:

$$f_j(\bar{\mathbf{x}}) = G(\bar{\mathbf{x}})\rho(\bar{\mathbf{x}}), \tag{7}$$

where

$$G(\bar{\mathbf{x}}) = \prod_{i=0}^{k-1} G(\mathbf{x}_i \leftrightarrow \mathbf{x}_{i+1})T(\mathbf{x}_i \leftrightarrow \mathbf{x}_{i+1}), \quad \rho(\bar{\mathbf{x}}) = \prod_{i=0}^{k} \rho(\mathbf{x}_i), \tag{8}$$

and

$$\rho(\mathbf{x}_i) = \begin{cases} L_e(\mathbf{x}_0 \to \mathbf{x}_1) & \text{if } i = 0, \\ W_j(\mathbf{x}_{k-1} \to \mathbf{x}_k) & \text{if } i = k, \\ \sigma_s(\mathbf{x}_i)\,\rho_p(\mathbf{x}_{i-1} \to \mathbf{x}_i, \mathbf{x}_i \to \mathbf{x}_{i+1}) & \text{otherwise.} \end{cases} \tag{9}$$

The geometry term $G(\mathbf{x}_i \leftrightarrow \mathbf{x}_{i+1})$ includes the visibility between $\mathbf{x}_i$ and $\mathbf{x}_{i+1}$, $T(\mathbf{x}_i \leftrightarrow \mathbf{x}_{i+1})$ is the transmittance between $\mathbf{x}_i$ and $\mathbf{x}_{i+1}$, $\sigma_s(\mathbf{x}_i)$ is the scattering coefficient at $\mathbf{x}_i$, and $\rho_p(\mathbf{x}_{i-1} \to \mathbf{x}_i, \mathbf{x}_i \to \mathbf{x}_{i+1})$ is the medium phase function. L_e and W_j are the emission and importance distribution functions at the light source and the eye sensor, respectively.

Path probability density. A Monte Carlo estimator for I_j (Eq. 1) with a random path $\bar{\mathbf{X}}$ sampled with density $p(\bar{\mathbf{X}})$ has the form

$$I_j = \int_\Omega \frac{f_j(\bar{\mathbf{x}})}{p(\bar{\mathbf{x}})} p(\bar{\mathbf{x}})\, d\mu(\bar{\mathbf{x}}) = E\left[\frac{f_j(\bar{\mathbf{X}})}{p(\bar{\mathbf{X}})}\right] \approx \frac{1}{N}\sum_{i=1}^{N} \frac{f_j(\bar{\mathbf{x}}_i)}{p(\bar{\mathbf{x}}_i)}. \tag{10}$$

Obtaining an actual numerical estimate of I_j requires sampling N random paths $\bar{\mathbf{x}}_i$ over Ω and deriving their corresponding probability density $p(\bar{\mathbf{x}}_i)$. Since μ is a product volume measure, the pdf of a path $\bar{\mathbf{x}}$ can be expressed as the joint density of its vertices:

$$p(\bar{\mathbf{x}}) = \frac{dP}{d\mu}(\bar{\mathbf{x}}) = p(\mathbf{x}_0, \ldots, \mathbf{x}_k). \tag{11}$$

Most sampling techniques construct paths bidirectionally and their joint density has the form

$$p(\mathbf{x}_0, \ldots, \mathbf{x}_k) = \left(p(\mathbf{x}_0)\prod_{i=1}^{s-1} p(\mathbf{x}_i|\mathbf{x}_{i-1})\right)\left(p(\mathbf{x}_k)\prod_{i=s}^{s+t-2} p(\mathbf{x}_i|\mathbf{x}_{i+1})\right), \tag{12}$$

where the first grouped terms correspond to the probability density of sequentially sampling s vertices from the light source (light subpaths), and the second grouped terms correspond to the probability density of sampling t vertices from the eye (eye subpaths).

3. Path Space Extension for Volumes

We now extend the path integral formulation to include photon density estimation with points and beams [JNSJ11]. We first present a general formulation that is capable of expressing many possible estimators. The next section shows how this formulation describes each existing algorithm including beam estimators. Similarly to Hachisuka et al.'s approach for the surface case [HPJ12], we define an *extended* path space Ω' and a corresponding extended measurement contribution function f'_j.

While the unified framework of Křivánek et al. [KGH*14] also addresses the same goal, there is a distinct difference between the two formulations: we *increase* the dimensionality of the path space

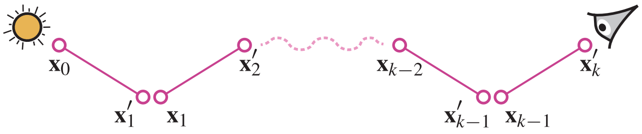

Figure 2: *Extended path of length k. Unlike regular paths, extended paths can be disconnected at every $(\mathbf{x}'_i, \mathbf{x}_i)$. Each pair of vertices $(\mathbf{x}_i, \mathbf{x}'_{i+1})$ forms the i-th disconnected segment.*

in order to include photon density estimation, whereas Křivánek et al. *reduce* the dimensionality of photon density estimation so that it fits within the original path space. Both formulations are amenable to Monte Carlo integral estimation and this difference is analogous to the difference between the two surface-based unification methods [GKDS12, HPJ12].

Our extended path integral formulation has the same general structure as the one in Equation 1:

$$I'_j = \int_{\Omega'} f'_j(\bar{\mathbf{x}}')d\mu'(\bar{\mathbf{x}}'), \tag{13}$$

where I'_j as an integral of an extended measurement contribution function f'_j over an extended path space Ω' with measure μ', all defined below. A Monte Carlo estimator for I' with a random extended path $\bar{\mathbf{X}}' \sim p'(\bar{\mathbf{X}})$ has the form

$$I'_j = \int_{\Omega'} \frac{f'_j(\bar{\mathbf{x}}')}{p'(\bar{\mathbf{x}}')} p'(\bar{\mathbf{x}}')d\mu'(\bar{\mathbf{x}}') = E\left[\frac{f'_j(\bar{\mathbf{X}}')}{p'(\bar{\mathbf{X}}')}\right]. \tag{14}$$

Paths. Using vertices $\mathbf{x}_i \in \mathcal{V}$, an extended path $\bar{\mathbf{x}}'$ of length k is

$$\bar{\mathbf{x}}' = \mathbf{x}_0\mathbf{x}'_1\mathbf{x}_1\mathbf{x}'_2 \ldots \mathbf{x}'_{k-1}\mathbf{x}_{k-1}\mathbf{x}'_k. \tag{15}$$

Compared to the original path definition in Equation 2, a length-k extended path has $2k$ vertices, rather than $k+1$ vertices. Another difference is that an extended path $\bar{\mathbf{x}}'$ is a "disconnected" light transport path, where the locations of disconnection are between every $\mathbf{x}'_i$ and $\mathbf{x}_i$. The vertices $(\mathbf{x}_i, \mathbf{x}'_{i+1})$ essentially represent the start and endpoints of the i-th *disconnected segment*. Figure 2 illustrates this idea. While it is also possible to duplicate vertices at the sensor and light source to have $2(k+1)$ vertices, we intentionally avoid this to make the view of paths as disconnected segments clear, omitting isolated duplicated vertices at the sensor and a light source.

Path space. The set Ω'_k of extended paths of length k becomes a union of the original set Ω_k and $k-1$ additional vertices:

$$\Omega'_k = \Omega_k \cup \mathcal{V}^{k-1}. \tag{16}$$

The full extended space Ω' is then naturally defined as

$$\Omega' = \bigcup_{k=1}^{\infty} \Omega'_k. \tag{17}$$

Path measure. A measure μ'_k on an extended path of length k is a product measure as before, but with $2k$ vertices:

$$\begin{aligned} d\mu'_k(\bar{\mathbf{x}}') &= dV(\mathbf{x}_0)dV(\mathbf{x}'_1)dV(\mathbf{x}_1)\cdots dV(\mathbf{x}'_k) \\ \mu'_k &= \underbrace{V \times \cdots \times V}_{2k \text{ times}} = V_{2k} \\ \mu'(D) &= \sum_{k=1}^{\infty} \mu'_k(D \cap \Omega'_k). \end{aligned} \tag{18}$$

Measurement contribution function. We now define the extended measurement contribution function based on the integral formulation of density estimation [Sil86] at each intermediate vertex:

$$f'_j(\bar{\mathbf{x}}') = G(\bar{\mathbf{x}}')\rho(\bar{\mathbf{x}}')K(\bar{\mathbf{x}}') \tag{19}$$

where

$$G(\bar{\mathbf{x}}') = \prod_{i=0}^{k-1} G(\mathbf{x}_i \leftrightarrow \mathbf{x}'_{i+1})T(\mathbf{x}_i \leftrightarrow \mathbf{x}'_{i+1}) \quad \rho(\bar{\mathbf{x}}') = \prod_{i=0}^{k} \rho(\mathbf{x}_i) \tag{20}$$

and

$$\rho(\mathbf{x}_i) = \begin{cases} L_e(\mathbf{x}_0 \to \mathbf{x}'_1) & \text{if } i = 0, \\ W_e^{(j)}(\mathbf{x}_{k-1} \to \mathbf{x}'_k) & \text{if } i = k, \\ \sigma_s(\mathbf{x}_i)\,\rho_p(\mathbf{x}_{i-1} \to \mathbf{x}'_i, \mathbf{x}_i \to \mathbf{x}'_{i+1}) & \text{otherwise.} \end{cases} \tag{21}$$

The key idea that enables this formulation for beams is that the path blurring kernel K is a product of kernels K_i that each considers *four* neighboring vertices, rather than two vertices as in the surface case [HPJ12]. That is, each K_i takes two disconnected segments $(\mathbf{x}_{i-1}, \mathbf{x}'_i)$ and $(\mathbf{x}_i, \mathbf{x}'_{i+1})$ as input:

$$K(\bar{\mathbf{x}}') = \prod_{i=1}^{k-1} K_i(\mathbf{x}_{i-1}, \mathbf{x}'_i, \mathbf{x}_i, \mathbf{x}'_{i+1}), \tag{22}$$

where we have two segments $(\mathbf{x}_{i-1}, \mathbf{x}'_i)$ and $(\mathbf{x}_i, \mathbf{x}'_{i+1})$ as the input. Similarly to the surface case [HPJ12], using Dirac deltas for all kernels makes the extended formulation identical to the regular formulation. We provide detailed definitions of the blurring kernel K for various existing density estimators in Section 4.

Path probability density. The probability density $p(\bar{\mathbf{x}}')$ of an extended path $\bar{\mathbf{x}}'$ is defined similarly as before:

$$p(\bar{\mathbf{x}}') = \frac{dP}{d\mu'}(\bar{\mathbf{x}}') = p(\mathbf{x}_0, \ldots, \mathbf{x}'_k). \tag{23}$$

Regular path sampling does not specify how to relate two vertices $\mathbf{x}'_i, \mathbf{x}_i \in \bar{\mathbf{x}}'$ at each disconnection of an extended path. We thus need to redefine path sampling to generate extended paths as we elaborate in the next section.

4. Reformulation of Existing Techniques

We show how the extended path integral formulation can express different kinds of volumetric scattering estimators, including Monte Carlo path integration as well as photon density estimation using points and beams. This formulation allows us to use the exact same measurement contribution function for bidirectional path tracing,

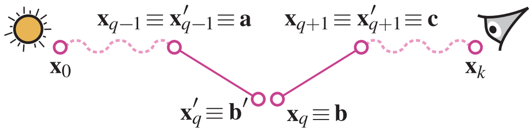

Figure 3: *Simplified configuration for an extended path of length k. The disconnection is only at* $(\mathbf{x}'_q, \mathbf{x}_q)$ *and subpaths are connected until this point.*

volumetric photon mapping, the beam-point estimator with 3D blur, and the beam-beam estimator with 3D blur [JNSJ11]. The only difference among those estimators is then in their corresponding probability densities.We also show that the beam-point estimator with 2D blur and the beam-beam estimator with 1D blur can be formulated via blurring kernels with Dirac delta functions. These delta functions cancel out with corresponding delta functions in the probability densities for these techniques.

Note that our formulation suggests the possibility of estimators with multiple disconnections for a light transport path. While this generalization can lead to the development of such novel estimators, all the existing estimators consider at most one disconnection, which corresponds to a location where density estimation is performed. Therefore, in the following, we consider only this special case with one disconnection at $\mathbf{x}'_q, \mathbf{x}_q$ as in Figure 3:

$$\bar{\mathbf{x}}' = \mathbf{x}_0 \ldots \mathbf{x}_{q-1}\mathbf{x}'_q\mathbf{x}_q\mathbf{x}_{q+1} \ldots \mathbf{x}_k. \tag{24}$$

For the sake of brevity, we define $\mathbf{a} \equiv \mathbf{x}_{q-1}$, $\mathbf{b}' \equiv \mathbf{x}'_q$, $\mathbf{b} \equiv \mathbf{x}_q$, and $\mathbf{c} \equiv \mathbf{x}_{q+1}$, so that we can use

$$\bar{\mathbf{x}}' = \mathbf{x}_0 \ldots \mathbf{a}\mathbf{b}'\mathbf{b}\mathbf{c} \ldots \mathbf{x}_k. \tag{25}$$

in the following derivations. This special case of a single disconnection corresponds to a blurring kernel K whose sub-kernels K_i are all Dirac deltas at $\mathbf{x}_i$, except for K_q. As a consequence, we can simplify $\bar{\mathbf{x}}'$ by omitting the coinciding vertices. Figure 4 summarizes the configurations and notation used in the following. Since the vertices $\Xi \equiv \mathbf{x}_0, \ldots, \mathbf{a}, \mathbf{c}, \ldots, \mathbf{x}_k$ are commonly given in all the techniques, we focus on the (conditional) probability density of $p(\mathbf{b}', \mathbf{b}|\Xi)$ given those vertices in the following.

4.1. Extended Bidirectional Path Tracing

Since the regular paths constructed by the sampling techniques in bidirectional path tracing do not have an extra vertex $\mathbf{b}'$, we consider a *perturbation* of $\mathbf{b}$ using some probability density $p(\mathbf{b}'|\mathbf{b})$:

$$p(\mathbf{b}', \mathbf{b}|\Xi) = p(\mathbf{b}'|\mathbf{b})p(\mathbf{b}|\Xi). \tag{26}$$

This algorithm is essentially bidirectional path tracing in the extended path space (Fig. 4b). The kernel in this case simplifies to

$$K(\mathbf{a}, \mathbf{b}', \mathbf{b}, \mathbf{c}) = K(\mathbf{b}'|\mathbf{b}). \tag{27}$$

It reduces to regular bidirectional path tracing if we use $p(\mathbf{b}'|\mathbf{b}) = K(\mathbf{b}'|\mathbf{b}) = \delta(\|\mathbf{b}' - \mathbf{b}\|)$. Note that $p(\mathbf{b}'|\mathbf{b}) = \delta(\|\mathbf{b}' - \mathbf{b}\|)$ means that we deterministically set $\mathbf{b}' = \mathbf{b}$. This formulation is analogous to the surface case [HPJ12]; the difference is that the kernel is volumetric.

4.2. Volume Photon Mapping (Point-Point 3D)

The probability density of an extended path in this technique is simply a product of two probability densities:

$$p(\mathbf{b}', \mathbf{b}|\Xi) = p(\mathbf{b}'|\Xi)p(\mathbf{b}|\Xi). \tag{28}$$

The kernel remains exactly the same as in extended bidirectional path tracing. Under this formulation, there is no fundamental difference between extended bidirectional path tracing and volumetric photon mapping, except for probability densities. Specifically, bidirectional path tracing samples $\mathbf{b}'$ by perturbing (possibly with a delta distribution) the end vertex of the eye path $\mathbf{b}$. Here, $\mathbf{b}'$ is sampled by extending the subpath from $\mathbf{a}$ (Fig. 4c).

4.3. Beam-Point 3D

We can express beams in the extended path integral formulation by sampling $\mathbf{b}'$ within the support of the kernel (Fig. 4d). For this technique, we have exactly the same measurement contribution function as extended bidirectional path tracing and volumetric photon mapping. The only difference is that we sample the vertex $\mathbf{b}'$ with the following specific pdf:

$$p(\mathbf{b}', \mathbf{b}|\Xi) = p(\omega_\mathbf{a}|\mathbf{a})G(\mathbf{a}, \mathbf{b}')k(t_\mathbf{a}|\mathbf{a}, \omega_\mathbf{a})p(\mathbf{b}|\Xi), \tag{29}$$

where $\omega_\mathbf{a}$ is the direction from $\mathbf{a}$ to $\mathbf{b}'$ and $t_\mathbf{a}$ is the distance from $\mathbf{a}$ to $\mathbf{b}'$ along this direction. The pdf $k(t_\mathbf{a}|\mathbf{a}, \omega_\mathbf{a})$ is non-zero only within the support of $K(\mathbf{b}'|\mathbf{b})$. One simple form is $k(t_\mathbf{a}|\mathbf{a}, \omega_\mathbf{a}) = 1/(t_\mathbf{a}^+ - t_\mathbf{a}^-)$ where $[t_\mathbf{a}^-, t_\mathbf{a}^+]$ is the intersection interval of the ray and the kernel support. We can see that, under our formulation, this technique is no longer fundamentally different from Point-Point 3D—only the distance sampling $t_\mathbf{a}$ is different, which in Point-Point 3D is proportional to the medium transmittance.

4.4. Beam-Beam 3D

Under our formulation, the Beam-Beam 3D estimator is also not fundamentally different from Beam-Point 3D, and measurement contribution function remains the same. The only difference is the distance sampling, this time for both $t_\mathbf{a}$ and $t_\mathbf{c}$ (Fig. 4e):

$$\begin{aligned} p(\mathbf{b}', \mathbf{b}|\Xi) = p(\omega_\mathbf{a}|\mathbf{a})G(\mathbf{a}, \mathbf{b}')k(t_\mathbf{a}|\mathbf{a}, \omega_\mathbf{a}) \\ p(\omega_\mathbf{c}|\mathbf{c})G(\mathbf{b}, \mathbf{c})k(t_\mathbf{c}|\mathbf{c}, \omega_\mathbf{c}), \end{aligned} \tag{30}$$

where $\omega_\mathbf{c}$ is the direction from $\mathbf{c}$ to $\mathbf{b}$ and $t_\mathbf{c}$ is the distance from $\mathbf{c}$ to $\mathbf{b}$ along this direction. Similarly to Beam-Point 3D, the pdfs $k(t_\mathbf{a}|\mathbf{a}, \omega_\mathbf{a})$ and $k(t_\mathbf{c}|\mathbf{c}, \omega_\mathbf{c})$ are non-zero only within the support of $K(\mathbf{b}'|\mathbf{b})$. For example, we can simply use uniform distributions $k(t_\mathbf{a}|\mathbf{a}, \omega_\mathbf{a}) = 1/(t_\mathbf{a}^+ - t_\mathbf{a}^-)$ and $k(t_\mathbf{c}|\mathbf{c}, \omega_\mathbf{c}) = 1/(t_\mathbf{c}^+ - t_\mathbf{c}^-)$ as the intersecting intervals of the support of the kernel.

Point-Point 3D typically uses importance sampling of the transmittance terms for those distances. Our formulation shows that Beam-Beam 3D samples distances such that $\mathbf{b}'$ and $\mathbf{b}$ are inside the kernel. The original formulation [JNSJ11] relies on analytical integration over the kernel support, which is a special case of achieving perfect importance sampling of distances $t_\mathbf{a}$ and $t_\mathbf{c}$ according to the path contributions within the kernel support.

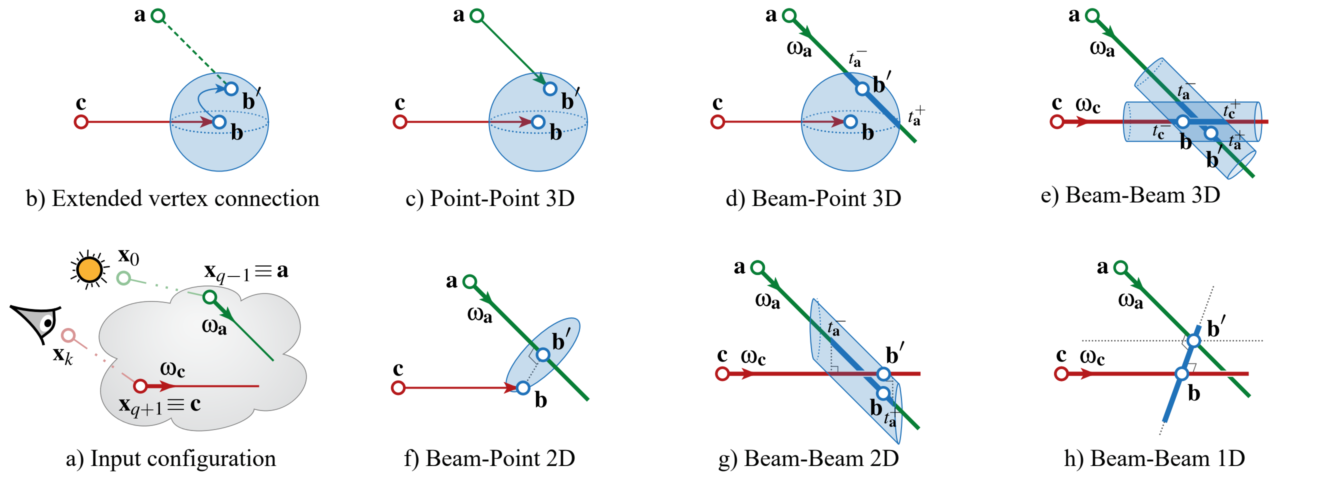

Figure 4: *Configurations and notation used to explain the existing estimators. We skip Point-Beam 3D, Point-Beam 2D, and the alternative Beam-Beam 2D in this paper since they are trivially defined from the other estimators based on the reciprocity in our formulation. We followed the same style and notation used for the merging formulation [KGH*14].*

4.5. Beam-Point 2D

In this estimator, the kernel support becomes a disc rather than a sphere (Fig. 4f). The original formulation [JNSJ11] simply introduced it as a 2D blurring kernel, but in order to express this disc kernel consistently in the extended path space, we flatten the 3D kernel by multiplying with a delta function:

$$K(\mathbf{a},\mathbf{b}',\mathbf{b},\mathbf{c}) = K_{2\mathrm{D}}(\mathbf{b}'|\mathbf{b})\delta(t_\mathbf{a} - d_{\mathbf{b}\perp}). \tag{31}$$

In this technique, the distance $t_\mathbf{a}$ is deterministically given, since it is given as the distance from **a** to the projection of **b** onto the ray $(\mathbf{a},\omega_\mathbf{a})$ which we denote as $d_{\mathbf{b}\perp}$. Deterministic sampling is expressed as a Dirac delta pdf, which results in

$$p(\mathbf{b}',\mathbf{b}|\Xi) = p(\omega_\mathbf{a}|\mathbf{a})G(\mathbf{a},\mathbf{b}')\delta(t_\mathbf{a} - d_{\mathbf{b}\perp})p(\mathbf{b}|\Xi) \tag{32}$$

The delta function in the pdf cancels out the delta function in the measurement contribution function.

4.6. Beam-Beam 2D

Under our formulation, this technique is very similar to Beam-Point 2D. The only difference is that $t_\mathbf{c}$ is sampled from a pdf that is non-zero only if $K_{2\mathrm{D}}(\mathbf{b}'|\mathbf{b})$ is non-zero (Fig. 4g). In Beam-Point 2D, this pdf is often given by importance sampling of the transmittance term along $\omega_\mathbf{c}$. Beam-Beam 2D instead importance samples the support of the kernel $K_{2\mathrm{D}}(\mathbf{b}'|\mathbf{b})$:

$$\begin{aligned} p(\mathbf{b}',\mathbf{b}|\Xi) = p(\omega_\mathbf{a}|\mathbf{a})G(\mathbf{a},\mathbf{b}')\delta(t_\mathbf{a} - d_{\mathbf{b}\perp}) \\ p(\omega_\mathbf{c}|\mathbf{c})G(\mathbf{b},\mathbf{c})k(t_\mathbf{c}|\mathbf{c},\omega_\mathbf{c}). \end{aligned} \tag{33}$$

The pdf $k(t_\mathbf{c}|\mathbf{c},\omega_\mathbf{c})$ can, for example, be a uniform distribution $k(t_\mathbf{c}|\mathbf{c},\omega_\mathbf{c}) = 1/(t_\mathbf{a}^+ - t_\mathbf{a}^-)$ where $[t_\mathbf{a}^-, t_\mathbf{a}^+]$ again is the intersection interval of the ray $(\mathbf{a},\omega_\mathbf{a})$ and the kernel support. The measurement contribution function remains the same as Beam-Point 2D and the delta function still cancels out.

4.7. Beam-Beam 1D

In this technique, distances along both rays are deterministically given by the closest points on the rays (Fig. 4h). We denote them as $d_\mathbf{a}$ and $d_\mathbf{c}$, which results in the kernel

$$K(\mathbf{a},\mathbf{b}',\mathbf{b},\mathbf{c}) = K_{1\mathrm{D}}(\mathbf{b}'|\mathbf{b})\delta(t_\mathbf{a} - d_\mathbf{a})\delta(t_\mathbf{c} - d_\mathbf{c}). \tag{34}$$

The probability density function contains two delta functions for distance sampling along rays:

$$\begin{aligned} p(\mathbf{b}',\mathbf{b}|\Xi) = p(\omega_\mathbf{a}|\mathbf{a})G(\mathbf{a},\mathbf{b}')\delta(t_\mathbf{a} - d_\mathbf{a}) \\ p(\omega_\mathbf{c}|\mathbf{c})G(\mathbf{b},\mathbf{c})\delta(t_\mathbf{c} - d_\mathbf{c}). \end{aligned} \tag{35}$$

Again, those delta functions cancel out as long as we use this specific technique for sampling. Note that the singularity [JNSJ11] appears due to the normalization factor of $K_{1\mathrm{D}}$, not due to the pdf.

5. Discussion

5.1. Merging vs Perturbation

Our extended path space formulation and the merging formulation [KGH*14] have complementary properties regarding how they model actual implementations of the estimators. This difference appears as discrepancy between the theory and the implementation for each formulation.

Since the merging formulation uses the regular path integral, it precisely models path integration based methods such as bidirectional path tracing. In order to model photon density estimation in the same path space, Křivánek et al. [KGH*14] exclude the photon from the path by interpreting it as a Russian roulette variable. This approach thus, in fact, formulates an unbiased form of photon density estimation. Regular photon density estimation does not exactly implement this unbiased formulation, which appears as an approximation in the Russian roulette probability in the implementation of the merging formulation.

In contrast, our extended path formulation exactly models photon density estimation including beams by explicitly modifying the path integral to include a blurring kernel. However, to model path integration methods, we need to use biased analogs for the original path integral formulation. *Virtual perturbation* [HPJ12] "approximates" those biased analogs by existing unbiased path integration methods. For instance, we do not perform vertex perturbation in extended bidirectional path tracing, but approximate it by regular bidirectional path tracing with virtual perturbation.

Those approximations in the two formulations are introduced for the sake of using existing algorithms. If we implement unbiased photon density estimation or extended bidirectional path tracing as-is, the theory and the implementation match exactly under both formulations. Our claim in this paper is that the extended path formulation explains photon density estimation better than the merging formulation, thus making blur kernels more natural to consider.

5.2. Connections to the Photon Beam Theories

Beam×Beam 3D. Jarosz et al. [JNSJ11] mentioned that this particular estimator is intractable, since their formulation uses the analytical solution of the measurement contribution function marginalized over two distance PDFs. Our formulation, however, points out that we can simply use Monte Carlo integration instead of analytically solving the marginalization. This allows us to avoid singularities present in the 1D blur for beams.

Beam×Point 2D. For a constant kernel, our formulation results in the same solution to the one by Jarosz et al. Our formulation also reveals that this technique as well as Beam×Beam 1D cannot be theoretically combined with other estimators with 3D blur under MIS, since the measurement contribution functions are all different. The delta functions in the measurement contribution function of 2D or 1D blur only cancel out with the pdf when using the sampling techniques corresponding to these lower-dimensional kernels.

Short beams. Beams in photon density estimation have two variations; short beams and long beams. The above formulation models long beams without any additional operation. We can also trivially model short beams as rejection sampling according to the transmittance term. In addition to importance sampling the distance t according to the corresponding kernel as we explained above, short beams essentially performs the additional Russian roulette of this sampled distance with the probability $\exp(-\sigma_t t)$. This cancels out the transmittance term since it is also exactly $\exp(-\sigma_t t)$ at the distance t. Jarosz et al. [JNT*11] showed in Appendix E of their paper that this process is equivalent to rejection sampling of the sampled distance t according to the transmittance.

6. Conclusion

We proposed an extended path integral formulation for volumetric transport with beam density estimators, unifying Monte Carlo path integration and photon density estimation, much like current merging formulations. Similar to surface transport counterparts, our extended formulation introduces additional vertices into light transport paths in order to treat sets of disconnected transport *segments* as light transport paths. Photon density estimators (including beam variants) naturally reside in our extended path space without any further modification, whereas Monte Carlo path integral estimates must additionally consider vertex perturbations. This result mirrors existing unified formulations, where Monte Carlo path integral estimates are more naturally expressed, but the inclusion of photon density estimators requires an additional Russian roulette process.

Our formulation is unique in that it recasts beam estimators as simple Monte Carlo vertex sampling processes. This formulation leads to a more practical implementation of photon beam estimators with 3D blur kernels. We believe that our formulation complements the growing foundation for a unified Monte Carlo path integration and photon density estimation framework for volumetric light transport. For instance, we expect that the ability to recast beam estimators as vertex sampling processes is immediately useful when analyzing (and improving upon) generalizations of conventional point-to-point photon density estimators to beam estimators.

Acknowledgments

This work was partly supported by the Charles University grant SVV-2017-260452 and the Czech Science Foundation grant 16-18964S.

References

[GKDS12] GEORGIEV I., KRIVÁNEK J., DAVIDOVIC T., SLUSALLEK P.: Light transport simulation with vertex connection and merging. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2012) 31*, 6 (2012), 192. 1, 3

[HPJ12] HACHISUKA T., PANTALEONI J., JENSEN H. W.: A path space extension for robust light transport simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2012) 31*, 6 (2012), 191. 1, 2, 3, 4, 6

[Jen96] JENSEN H. W.: Global illumination using photon maps. In *Rendering Techniques 96 (Proceedings of Eurographics Workshop on Rendering)*. Springer, 1996, pp. 21–30. 1

[JNSJ11] JAROSZ W., NOWROUZEZAHRAI D., SADEGHI I., JENSEN H. W.: A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Transactions on Graphics (TOG) 30*, 1 (2011), 5. 1, 2, 4, 5, 6

[JNT*11] JAROSZ W., NOWROUZEZAHRAI D., THOMAS R., SLOAN P.-P., ZWICKER M.: Progressive photon beams. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2011) 30*, 6 (2011), 181. 6

[Kaj86] KAJIYA J. T.: The rendering equation. *Computer Graphics (Proceedings of SIGGRAPH '86) 20*, 4 (1986), 143–150. 1

[KGH*14] KŘIVÁNEK J., GEORGIEV I., HACHISUKA T., VÉVODA P., ŠIK M., NOWROUZEZAHRAI D., JAROSZ W.: Unifying points, beams, and paths in volumetric light transport simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014) 33*, 4 (2014), 103. 1, 2, 5

[LW93] LAFORTUNE E. P., WILLEMS Y. D.: Bi-directional path tracing. In *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics '93)* (Alvor, Portugal, December 1993), pp. 145–153. 1

[PKK00] PAULY M., KOLLIG T., KELLER A.: Metropolis light transport for participating media. In *Rendering Techniques 2000 (Proceedings of Eurographics Workshop on Rendering)*. Springer, 2000, pp. 11–22. 1, 2

[Sil86] SILVERMAN B. W.: *Density estimation for statistics and data analysis*, vol. 26. CRC press, 1986. 3

[Vea97] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1997. 1, 2

[VG94] VEACH E., GUIBAS L. J.: Bidirectional estimators for light transport. In *Photorealistic Rendering Techniques (Proceedings of Eurographics Workshop on Rendering)*. Springer, 1994, pp. 147–162. 1