EUROGRAPHICS 2008 / G. Drettakis and R. Scopigno
(Guest Editors)

Volume 27 (2008), Number 2

The Beam Radiance Estimate for Volumetric Photon Mapping

Wojciech Jarosz   Matthias Zwicker   Henrik Wann Jensen

University of California, San Diego

Abstract

We present a new method for efficiently simulating the scattering of light within participating media. Using a theoretical reformulation of volumetric photon mapping, we develop a novel photon gathering technique for participating media. Traditional volumetric photon mapping samples the in-scattered radiance at numerous points along the length of a single ray by performing costly range queries within the photon map. Our technique replaces these multiple point-queries with a single beam-query, which explicitly gathers all photons along the length of an entire ray. These photons are used to estimate the accumulated in-scattered radiance arriving from a particular direction and need to be gathered only once per ray. Our method handles both fixed and adaptive kernels, is faster than regular volumetric photon mapping, and produces images with less noise.

Keywords: participating media, light transport, global illumination, rendering, photon tracing, photon map, ray marching, nearest neighbor, variable kernel method.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: raytracing; color, shading, shadowing, and texture; I.6.8 [Simulation and Modeling]: Monte Carlo; G.1.9 [Numerical Analysis]: Fredholm equations; integro-differential equations.

1. Introduction

The appearance of many natural phenomena, such as human skin, clouds, fire, water, or the atmosphere, are strongly influenced by the interaction of light with volumetric media. Therefore, efficiently rendering scenes with participating media has been an area of interest within computer graphics. This problem is challenging, however, because accurately simulating light transport in participating media is computationally very expensive. Cerezo et al. [CPP*05] provide a recent and comprehensive overview of the wealth of research that has been devoted to address this problem.

Light transport in arbitrary participating media is modeled by the radiative transfer equation [Cha60]. The pioneering contributions by Kajiya and von Herzen [KH84] and Rushmeier and Torrance [RT87] are at the beginning of a long list of work on rendering participating media in the computer graphics community. Some of the most popular techniques to date are based on stochastic path-tracing and Monte Carlo integration [PM93, LW96, PKK00]. These approaches are attractive because of their sound underlying theoretical framework and their generality. They are unbiased and guaranteed to converge to the exact solution. In addition, it is straightforward to include heterogeneous media, anisotropic phase functions, and scattering from surfaces. The downside of these approaches is that they suffer from noise that can only be overcome with a huge computational effort.

One strategy to solve this issue is to make simplifying assumptions about the participating media. For example, homogeneous media with a high scattering albedo can be modeled accurately using a diffusion approximation [Sta95, JMLH01], which leads to very efficient rendering algorithms. Premoze et al. [PARN04], under the assumption that the medium is tenuous and strongly forward scattering, use a path integral formulation to derive efficient rendering algorithms. Sun et al. [SRNN05] render single scattering in real time, but without shadowing effects.

In contrast, photon mapping [JC98] improves the efficiency of path-tracing without making additional assumptions about the properties of the medium being rendered. Similar to Monte Carlo methods, photon mapping handles isotropic, anisotropic, homogeneous, and heterogeneous media of arbitrary albedo. A disadvantage of photon mapping is that it introduces bias to the solution of the radiative transfer equation. In practice, however, this bias is preferable to the noisy solutions of pure Monte Carlo methods.

Intuitively, photon mapping works by splitting the energy emitted by each light source into discrete packets, so called *photons*. In a first pass, the propagation of light is simulated

Symbol	Units	Description
$\sigma_s(\mathbf{x})$	m^{-1}	Scattering coefficient at $\mathbf{x}$
$\sigma_a(\mathbf{x})$	m^{-1}	Absorption coefficient at $\mathbf{x}$
$\sigma_t(\mathbf{x})$	m^{-1}	Extinction coefficient at $\mathbf{x}$
$p(\mathbf{x},\vec{\omega},\vec{\omega}')$	sr^{-1}	Normalized phase function
$\tau(\mathbf{x}\leftrightarrow\mathbf{x}')$	unitless	Optical thickness: $\int_{\mathbf{x}'}^{\mathbf{x}} \sigma_t(x)\,dx$
$T_r(\mathbf{x}\leftrightarrow\mathbf{x}')$	unitless	Transmittance: $e^{-\tau(\mathbf{x}\leftrightarrow\mathbf{x}')}$
$\Omega_{4\pi}$	sr	Sphere of directions
$L(\mathbf{x}\rightarrow\mathbf{x}')$	$Wm^{-2}sr^{-1}$	Outgoing radiance at $\mathbf{x}$ towards $\mathbf{x}'$
$L(\mathbf{x},\vec{\omega})$	$Wm^{-2}sr^{-1}$	Incident radiance at $\mathbf{x}$ from $\vec{\omega}$
$W_e(\mathbf{x},\vec{\omega})$	$m^{-3}sr^{-1}$	Importance at $\mathbf{x}$ towards $\vec{\omega}$
$W_e(\mathbf{x}\rightarrow\mathbf{x}')$	$m^{-3}sr^{-1}$	Importance at $\mathbf{x}$ towards $\mathbf{x}'$

Table 1: Definitions of quantities used throughout the paper.

by scattering photons in the scene. At each scattering event (at surfaces or within the medium) the incident energy carried by a photon is stored in a *photon map*. In a second pass, the photon map is used to evaluate the radiance at discrete points in the scene by locally computing the photon density. To compute the radiance carried along each ray towards the eye, the radiance is estimated within the medium at several sample points along the ray [JC98]. At each point the attenuation through the medium to the eye is computed, and the attenuated radiance is added to the ray. The main disadvantage of this procedure, however, is that it is difficult to find a good distribution of sample points along the ray. On one hand, if not enough sample points are used, the result is likely to be noisy. On the other hand, increasing the number of sample points is very costly and can slow down rendering significantly.

In this paper, we propose a novel approach for computing the contribution of in-scattered radiance. We gather photons along viewing rays and analytically compute their contributions, without point sampling. We present the following contributions:

- We derive a reformulation of volumetric photon mapping as a solution to the measurement equation. This theory allows for arbitrary *measurements* of radiance to be computed within participating media, where a measurement is simply an integral of the radiance multiplied with a weighting function.
- Using this new theory, we present an improved radiance estimate for volumetric photon mapping based on "beam gathering." This technique eliminates the need for stepping through the medium to find photons. Instead, it gathers all photons along a ray. We show how to efficiently implement this new gathering technique for both fixed and adaptive smoothing kernels and demonstrate that our method produces images with less noise than conventional photon mapping.

The rest of this paper is organized as follows. In Section 2, we review the theory of radiance transport within participating media and the volumetric photon mapping method. In Section 3, we reformulate volumetric photon mapping in terms of the measurement equation and show how the photon map can be used to estimate *any* measurement of radiance within the scene. In Section 4, we present our new beam radiance estimate using this theory and describe the data structures needed to evaluate it efficiently. Finally, we show comparisons of our approach to conventional photon mapping in Section 5 and discuss avenues of future work in Section 6.

2. Photon Mapping in Participating Media

Light transport within participating media is described by the *radiative transfer equation* (RTE) [Cha60], which defines the radiance that reaches a point $\mathbf{x}$ from direction $\vec{\omega}$ as a sum of the exitant radiance from the nearest surface from this direction, and the accumulated in-scattered radiance from the medium between the surface and $\mathbf{x}$ (see Figure 1). This can be expressed as:

$$L(\mathbf{x},\vec{\omega}) = T_r(\mathbf{x}\leftrightarrow\mathbf{x}_s)L(\mathbf{x}_s,\vec{\omega})+ \int_0^s T_r(\mathbf{x}\leftrightarrow\mathbf{x}_t)\sigma_s(\mathbf{x}_t)L_i(\mathbf{x}_t,\vec{\omega})\,dt, \tag{1}$$

where T_r is the transmittance, s is the distance through the medium to the nearest surface at $\mathbf{x}_s = \mathbf{x} - s\vec{\omega}$, and $\mathbf{x}_t = \mathbf{x} - t\vec{\omega}$ with $t \in (0,s)$. We define the remaining quantities in Table 1.

The surface radiance, $L(\mathbf{x}_s,\vec{\omega})$, at the boundary of the medium is governed by the *rendering equation* [Kaj86]. The in-scattered radiance, $L_i(\mathbf{x}_t,\vec{\omega})$, depends on the radiance arriving at $\mathbf{x}_t$ from all directions $\vec{\omega}_t$ over the sphere of directions $\Omega_{4\pi}$ and is defined as:

$$L_i(\mathbf{x}_t,\vec{\omega}) = \int_{\Omega_{4\pi}} p(\mathbf{x}_t,\vec{\omega},\vec{\omega}_t)L(\mathbf{x}_t,\vec{\omega}_t)\,d\vec{\omega}_t, \tag{2}$$

where p is the normalized phase function.

Volumetric photon mapping [JC98] solves the RTE using a combination of photon tracing, ray-marching, and density estimation. In a preprocess, packets of energy are shot from light sources, scattered at surfaces and within the medium, and their interactions are stored in a global data structure. During rendering, ray marching is used to numerically integrate Equation 1 for radiance seen directly by the observer,

$$L(\mathbf{x},\vec{\omega}) \approx T_r(\mathbf{x}\leftrightarrow\mathbf{x}_s)L(\mathbf{x}_s,\vec{\omega})+ \left(\sum_{t=0}^{S-1} T_r(\mathbf{x}\leftrightarrow\mathbf{x}_t)\sigma_s(\mathbf{x}_t)L_i(\mathbf{x}_t,\vec{\omega})\Delta_t\right), \tag{3}$$

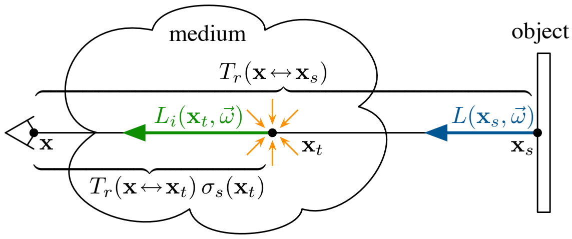

Figure 1: The radiance reaching the eye $L(\mathbf{x},\vec{\omega})$ is the sum of the radiance from the surface $L(\mathbf{x}_s,\vec{\omega})$ and the accumulated in-scattered radiance $L_i(\mathbf{x}_t,\vec{\omega})$ along a ray.

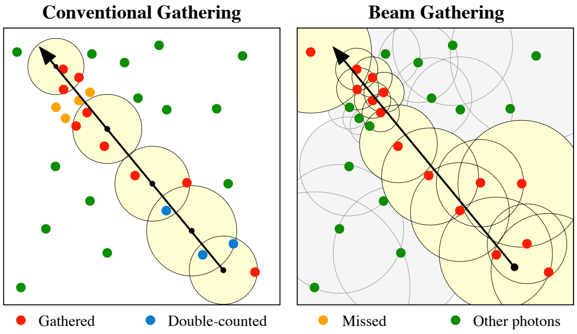

Figure 2: Conventional gathering (left) searches for photons in a sphere around numerous samples along the ray. Our method (right) assigns a radius to each photon and performs a single range search to find all photons along the length of the entire ray.

where Δ_s is the length of each segment along the ray and $\mathbf{x}_0, \ldots, \mathbf{x}_s$ are the sample points for each segment ($\mathbf{x}_0$ is the point where the ray enters the medium and $\mathbf{x}_s$ is a point on a surface past the medium).

The most expensive part to compute in Equation 3 is the in-scattered radiance L_i, because it involves accounting for all light arriving at each point $\mathbf{x}_t$ along the ray from any other point in the scene. Instead of computing these values independently for each location, photon mapping gains efficiency by reusing the computation performed during the photon tracing stage. The in-scattered radiance is approximated using density estimation by gathering photons within a small spherical neighborhood of radius r around each sample location $\mathbf{x}_t$,

$$L_i(\mathbf{x}_t, \vec{\omega}) \approx \sum_{i=1}^{n} \frac{p(\mathbf{x}_t, \vec{\omega}, \vec{\omega}_i)\Delta\Phi_i}{\frac{4}{3}\pi r^3}, \tag{4}$$

where $\Delta\Phi_i$ is the power of photon i, and $\vec{\omega}_i$ is its incident direction [JC98].

Though photon mapping is much more efficient than brute-force techniques like path tracing, the density estimation requires searching for photons within a global data structure, which is quite expensive. This formulation is suboptimal, firstly because it may gather the same photons more than once if the spherical neighborhoods overlap and, secondly, because it can lead to noise if the step size is too large and photons are omitted (see Figure 2).

3. Reformulation of Volumetric Photon Mapping

Our technique solves these shortcomings by querying once for photons along the length of an entire ray, instead of multiple times near points along the ray (see Figure 2). More formally, whereas regular photon mapping estimates L_i at discrete points using Equation 4, our main contribution is to directly estimate

$$\int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t)\sigma_s(\mathbf{x}_t)L_i(\mathbf{x}_t, \vec{\omega})\, dt \tag{5}$$

along rays.

Though the explanation of photon mapping from the previous section is appealing at an intuitive level, it does not rigorously present the algorithm as a numerical solution to the RTE. Furthermore, this explanation is heavily tied to the geometric interpretation of gathering photons within a disc (on surfaces) or within a sphere (in participating media). In order to avoid these limitations and use the photon map to estimate general radiometric quantities in the volume, such as Equation 5, we use a more flexible derivation of particle tracing methods presented by Veach [Vea98]. We extend this derivation to handle participating media (Section 3.1) and show how to represent particle tracing algorithms like volumetric photon mapping in terms of the measurement equation (Sections 3.2 and 3.3). Finally, we show how to use the same photon maps to estimate more general quantities of radiance (Section 3.4).

3.1. Volumetric Path Integral Formulation

We use the path integral formulation of the RTE, which arises by recursively expanding the right hand side of Equation 1. Instead of expressing the radiance equilibrium *recursively*, the resulting path integral formulation is a *sum* over light-carrying paths of different lengths. For conciseness we restrict the following derivation to include only scattering within the volume. A full derivation incorporating scattering at surfaces is available in a technical report [JZJ08].

A path of length k, $\bar{\mathbf{x}}_k$, with $k+1$ vertices is defined as

$$\bar{\mathbf{x}}_k = \mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_k, \tag{6}$$

where light starts on a light source at $\mathbf{x}_k$ and scatters $k-1$ times within the medium before reaching $\mathbf{x}_0$. The outgoing radiance L_o at $\mathbf{x}_1$ towards $\mathbf{x}_0$ can be expressed as a sum over all possible light paths, of any length, arriving at $\mathbf{x}_1$,

$$L_o(\mathbf{x}_1 \to \mathbf{x}_0) = \sum_{k=1}^{\infty} \bar{L}(\bar{\mathbf{x}}_k). \tag{7}$$

$\bar{L}(\bar{\mathbf{x}}_k)$ measures the amount of radiance transported to $\mathbf{x}_1$ from all paths of length k,

$$\bar{L}(\bar{\mathbf{x}}_k) = \int_{\mathcal{A}(\mathbf{x}_k)} \cdots \int_{\mathcal{V}(\mathbf{x}_2)} L_e(\mathbf{x}_k \to \mathbf{x}_{k-1})\cos(\theta_k) \left(\prod_{j=1}^{k-1} p(\mathbf{x}_j) G(\mathbf{x}_{j+1} \to \mathbf{x}_j)\right) d\mathcal{V}(\mathbf{x}_2)\cdots d\mathcal{A}(\mathbf{x}_k), \tag{8}$$

where $\mathcal{V}$ and $\mathcal{A}$ are the domains of the volume and surfaces respectively (all integrals except at $\mathbf{x}_k$ are over volume), θ_k is the angle between the surface normal at $\mathbf{x}_k$ and vertex $\mathbf{x}_{k-1}$, and we use the shorthand $p(\mathbf{x}_j) = p(\mathbf{x}_{j+1} \to \mathbf{x}_j \to \mathbf{x}_{j-1})$ for the phase function. The geometry term, G, is defined as

$$G(\mathbf{x} \to \mathbf{y}) = \frac{V(\mathbf{x} \leftrightarrow \mathbf{y})\sigma_s(\mathbf{x})T_r(\mathbf{x} \leftrightarrow \mathbf{y})}{\|\mathbf{x}-\mathbf{y}\|^2}. \tag{9}$$

The radiance transported along an example path is shown in Figure 3.

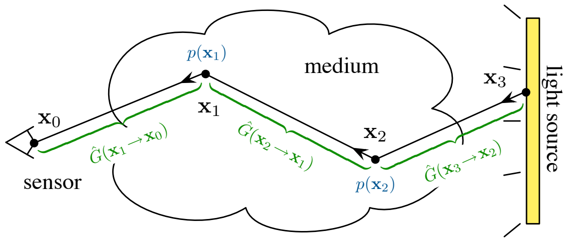

Figure 3: An example path, $\bar{\mathbf{x}}_3$, with length $k = 3$. The radiance transported along the path, $\bar{L}(\bar{\mathbf{x}}_3)$, is the emitted radiance at the light multiplied by a series of scattering events (blue) and geometry terms (green).

3.2. The Measurement Equation

Many global illumination algorithms can be described in terms of the *measurement equation*. The measurement equation describes an abstract measurement of incident radiance taken over some set of rays in a scene

$$I = \langle W_e, L \rangle = \int_{\mathcal{V}} \int_{\Omega_{4\pi}} W_e(\mathbf{x}, \vec{\omega}) L(\mathbf{x}, \vec{\omega})\, d\vec{\omega}\, d\mathcal{V}(\mathbf{x}). \quad (10)$$

The importance function W_e represents an abstract measuring sensor and is defined over the whole ray space $\mathcal{V} \times \Omega_{4\pi}$ (though typically W_e is non-zero for only a small subset of this domain).

Path tracing, for instance, measures the contribution of radiance arriving over the area of a pixel. Radiosity algorithms integrate the contribution of radiance over basis functions defined on the scene geometry. Both of these approaches can be described using Equation 10 with an appropriate importance function.

In his dissertation, Veach [Vea98] showed how particle tracing methods for surface illumination can also be expressed as a solution to the measurement equation by using the path integral form of the rendering equation. We extend on this idea and use the volumetric path integral formulation to describe volumetric photon tracing in the same way.

3.3. Volumetric Photon Tracing

Photon tracing methods can be thought of as a way of generating samples from the scene's equilibrium radiance distribution and then using this single collection of samples to render the entire image. The photon tracing stage generates N weighted sample rays, or photons, $(\alpha_i, \mathbf{x}_i, \vec{\omega}_i)$, where each $(\mathbf{x}_i, \vec{\omega}_i)$ is a ray and α_i is a corresponding weight. Our goal is to use these samples to take *unbiased* estimates of the radiance as a weighted sum,

$$E\left[\frac{1}{N}\sum_{i=1}^{N} W_e(\mathbf{x}_i, \vec{\omega}_i)\alpha_i\right] = \langle W_e, L \rangle, \quad (11)$$

for an arbitrary importance function W_e. We must therefore determine the proper distribution of samples for Equation 11 to hold. By expanding the measurement equation (10) in terms of the *outgoing* radiance using the path integral formulation (7 and 8), it can be shown that distributing the samples using a random-walk satisfies the necessary requirements if

$$\alpha_i = \frac{L_e(\mathbf{x}_{i,k_i} \to \mathbf{x}_{i,k_i-1})}{pdf(\mathbf{x}_{i,k_i}, \mathbf{x}_{i,k_i-1})} \prod_{j=1}^{k_i-1} \left(\frac{1}{q_{i,j}} \frac{p(\mathbf{x}_{i,j}) G(\mathbf{x}_{i,j+1} \to \mathbf{x}_{i,j})}{pdf(\mathbf{x}_{i,j-1})} \right) G(\mathbf{x}_{i,1} \to \mathbf{x}_{i,0}), \quad (12)$$

where $q_{i,j}$ is the probability of terminating the walk at the j^{th} vertex. A detailed derivation is provided in the supplemental technical report [JZJ08].

Connection to Conventional Photon Tracing. Though derived in a different fashion, Equation 12 is exactly how conventional photon mapping distributes photons within the scene. For instance, for a diffuse area light, photons are emitted using a cosine distribution with the power of the light source. In Equation 12 photons are emitted with the radiance of the light source divided by the *pdf* of choosing a position and direction on the light. These quantities are equivalent. Hence the particles generated above represent differential flux. The correspondence between the photon powers [JC98] and the sample weights is $\Delta\Phi_i = \frac{\alpha_i}{N}$.

3.4. Radiance Estimation Using the Measurement Equation

The main advantage of the reformulation in Section 3.3 is that it naturally accommodates computation of any measurement of radiance within the scene simply by using an appropriately defined importance function W_e. In this section, we first show how the conventional estimate for in-scattered radiance can be expressed as a measurement. We then go one step further and show how to derive a beam radiance estimate which approximates Equation 5 along rays directly.

Conventional Radiance Estimate. The conventional radiance estimate approximates the value of the in-scattered radiance L_i at fixed points within the scene. To express this using the theory from Section 3, we need to transform Equation 2 into the measurement equation. Since the measurement equation is an integral over all of ray space $(\mathcal{V} \times \Omega_{4\pi})$, we artificially expand L_i to also integrate over the volume

$$\begin{aligned} L_i(\mathbf{x}_t, \vec{\omega}) &= \int_{\Omega_{4\pi}} p(\mathbf{x}_t, \vec{\omega}, \vec{\omega}_t) L(\mathbf{x}_t, \vec{\omega}_t)\, d\vec{\omega}_t \quad (13) \\ &= \int_{\mathcal{V}} \int_{\Omega_{4\pi}} \delta(\|\mathbf{x}' - \mathbf{x}_t\|) p(\mathbf{x}', \vec{\omega}, \vec{\omega}_t) \\ &\qquad L(\mathbf{x}', \vec{\omega}_t)\, d\vec{\omega}_t\, d\mathcal{V}(\mathbf{x}'). \end{aligned}$$

In order to keep the expressions equivalent when we add the integration over volume, we also introduce a Dirac delta function δ.

The bottom row of the above equation is now the measurement equation where $W_e = \delta(\|\mathbf{x}' - \mathbf{x}_t\|) p(\mathbf{x}', \vec{\omega}, \vec{\omega}_t)$. Hence we can compute an *unbiased* estimate using the photon map by evaluating Equation 11 with this importance function. However, in order to obtain a useful estimate of radiance at all points in the scene a normalized kernel function is used in place of the delta function. This is where bias is introduced in the photon mapping method. Another interpretation is that by replacing the delta function with a kernel, photon mapping computes an *unbiased* estimate of *blurred* radiance. Jensen and Christensen [JC98] use a constant three-dimensional kernel with a radius based on the n^{th} nearest neighbor. This results in the following radiance estimate by applying Equation 11

$$L_i(\mathbf{x}_t,\vec{\omega}) \approx \int_{\mathcal{V}}\int_{\Omega_{4\pi}} K_n(\|\mathbf{x}'-\mathbf{x}_t\|)p(\mathbf{x}',\vec{\omega},\vec{\omega}_t) L(\mathbf{x}',\vec{\omega}_t)\, d\vec{\omega}_t\, d\mathcal{V}(\mathbf{x}'), \tag{14}$$

$$\approx \frac{1}{N}\sum_{i=1}^{N} K_n(\|\mathbf{x}_i-\mathbf{x}_t\|)p(\mathbf{x}_i,\vec{\omega},\vec{\omega}_i)\alpha_i \tag{15}$$

where the kernel K_n is defined as

$$K_n(r) = \begin{cases} \frac{3}{4\pi d_n^3} & \text{if } r \in [0,d_n] \\ 0 & \text{otherwise} \end{cases}, \tag{16}$$

and d_n is the distance to the n^{th} photon. Note that this is equivalent to the conventional volumetric radiance estimate in Equation 4.

Beam Radiance Estimate. A similar procedure can be used to derive an estimate for the accumulated in-scattered radiance along an entire ray. To accomplish this, we first expand out L_i in Equation 5 and then artificially inflate the resulting expression to integrate over the whole volume,

$$\int_0^s\int_{\Omega_{4\pi}} T_r(\mathbf{x}\leftrightarrow\mathbf{x}_t)\sigma_s(\mathbf{x}_t)p(\mathbf{x}_t,\vec{\omega},\vec{\omega}_t)L(\mathbf{x}_t,\vec{\omega}_t)\, d\vec{\omega}_t\, dt = \tag{17}$$

$$\int_{\mathbb{R}}\int_0^{2\pi}\int_{\mathbb{R}}\int_{\Omega_{4\pi}} \delta(r)(H(t)-H(t-s))T_r(\mathbf{x}\leftrightarrow\mathbf{x}')\sigma_s(\mathbf{x}') p(\mathbf{x}',\vec{\omega},\vec{\omega}_t)L(\mathbf{x}',\vec{\omega}_t)\, d\vec{\omega}_t\, dr\, d\theta\, dt. \tag{18}$$

$\mathbb{R}$ is the set of real numbers and $\mathbf{x}'$ is expressed in cylindrical coordinates, (t,θ,r), about $(\mathbf{x},\vec{\omega})$, where r is the radius to the ray (see Figure 4). We have added a Dirac delta function δ as before, and the Heaviside step functions ($H(x)=1$ when $x>0$ and 0 otherwise) constrain the computation to $t\in(0,s)$. Equation 18 is now equivalent to the measurement equation where the integral over volume has been converted into cylindrical coordinates and where $W_e=\delta(r)(H(t)-H(t-s))T_r(\mathbf{x}\leftrightarrow\mathbf{x}')\sigma_s(\mathbf{x}')p(\mathbf{x}',\vec{\omega},\vec{\omega}_t)$.

Since the probability of photons landing exactly on the ray $(\mathbf{x},\vec{\omega})$ is zero, we introduce bias by blurring the radiance and replacing the delta and step functions with a smooth kernel, K. This integral can then be estimated with the measurement

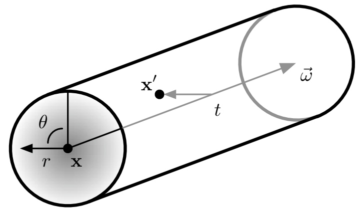

Figure 4: In the beam radiance estimate, $\mathbf{x}'$ is parameterized in cylindrical coordinates, (t,θ,r), about the ray $(\mathbf{x},\vec{\omega})$. An unbiased estimate would only consider points directly on the ray, while a biased version uses a kernel (shown in grey) to blur the radiance within a beam.

equation using the photons as:

$$\int_{\mathbb{R}}\int_0^{2\pi}\int_{\mathbb{R}}\int_{\Omega_{4\pi}} K(t,\theta,r)T_r(\mathbf{x}\leftrightarrow\mathbf{x}')\sigma_s(\mathbf{x}')p(\mathbf{x}',\vec{\omega},\vec{\omega}_t) L(\mathbf{x}',\vec{\omega}_t)\, d\vec{\omega}_t\, dr\, d\theta\, dt = \tag{19}$$

$$\frac{1}{N}\sum_{i=1}^{N} K(t_i,\theta_i,r_i)T_r(\mathbf{x}\leftrightarrow\mathbf{x}_i)\sigma_s(\mathbf{x}_i)p(\mathbf{x}_i,\vec{\omega},\vec{\omega}_i)\alpha_i, \tag{20}$$

where $(t_i,\theta_i,r_i)=\mathbf{x}_i$ are the cylindrical coordinates of photon i about the ray.

The blurring in the conventional radiance estimate is spherical and so the kernel needs to be normalized for 3D. However, with the beam radiance estimate, we blur in two dimensions (perpendicular to the ray) since the radiance we are computing already includes the integration along the ray itself. Therefore, the kernel in the beam estimate is normalized for 2D.

3.5. Kernel Radiance Estimation

For both the conventional and beam radiance estimates the characteristics of the bias and blur are determined by the smoothing function chosen. Several options exist for applying a smoothing kernel to the photon map data.

The *kernel method* uses a fixed-radius smoothing kernel and results in a uniform blur of radiance within the scene. In practice, using a fixed-width circular kernel implies that in order to evaluate the beam radiance estimate (Equation 19) using the photon map (Equation 11) we only need to consider photons which are located within a fixed-radius cylinder about the ray $(\mathbf{x},\vec{\omega})$. Alternatively, the equivalent dual interpretation considers each photon as the center of an oriented disc *facing the ray* and all *photon-discs* that intersect the ray need to be found.

If the density of photons varies significantly it can be difficult to choose a single radius that works well for all regions of the scene. This can be solved by allowing the size and shape of the blurring kernel to vary spatially. In conventional photon mapping, the *nearest neighbor method* (k-NN) is used to adapt the kernel width to the local density. Generalizing point-based k-NN to a visually comparable range search along rays is challenging. However, spatial variation can eas-

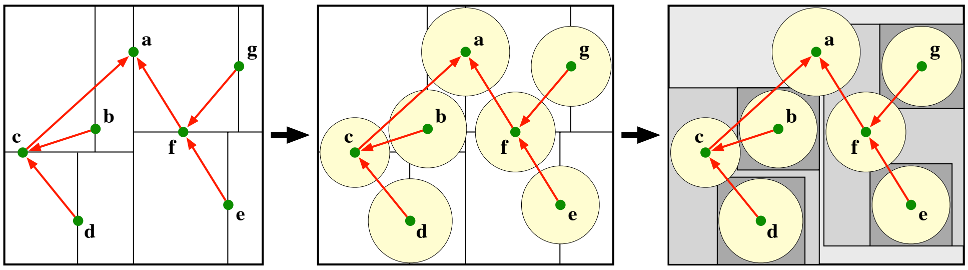

Construct balanced kd-tree — Compute radius per photon — Construct BBH over photon-discs

Figure 5: After photons have been distributed in the scene, our algorithm constructs a balanced kd-tree (left). We assign a valid radius to each photon by querying in the kd-tree (middle). Finally, we rapidly construct a bounding-box hierarchy over the photon-discs (right) by reusing the same hierarchical structure (shown in red) as the kd-tree.

ily be applied to the dual *photon-discs* interpretation using the *variable kernel method* (VK) [BMP77]. A smoothing kernel is "attached" to each photon and the radius of the kernel is allowed to vary from one photon to another, based on the local density. In contrast to k-NN estimation, where the kernel widths vary based on the distance from the *evaluation* location to the data points, in the VK method the kernel widths only depend on the data points themselves. In order to facilitate this, the method relies on a *pilot estimate* of the local density at each data point in order to assign the kernel widths. This is the approach we take.

4. Algorithm

In order to use the dual interpretation to evaluate the beam radiance estimate (Equation 19), we need an efficient way of locating all photon-discs that intersect an arbitrary ray. Additionally, to use variable width kernels we need to efficiently compute a radius for each photon in the photon map. At a high level, our volumetric photon mapping technique performs the following five steps:

1. Shoot photons from light sources.
2. Construct a balanced kd-tree for the photons.
3. Assign a radius for each photon.
4. Construct a bounding-box hierarchy over the photon-discs.
5. Use the photon BBH to render the image.

Steps 1 and 2 are identical to conventional photon mapping while 3–5 are unique to our approach and are explained in more detail below.

Photon Radius Computation. We augment the traditional photon map by associating a radius with each photon. For fixed width kernels the radius is a constant for all photons and does not need to be explicitly stored. For variable width kernels using the VK method, we perform a density estimation at each photon to assign a radius. At each photon we compute the local density by estimating the distance to the n^{th} nearest photon and use this as the photon-disc's radius. This pilot estimate is performed using the photon map kd-tree. The value n plays the same role as in the conventional radiance estimate, it controls the amount of blur.

As an optimization, we only search for the nearest $m \ll n$ photons and estimate the necessary radius for n photons. By assuming locally uniform photon density, if $d_{i,m}$ is the distance to the m^{th} photon from photon i, we estimate the radius as $r_i = d_{i,m}\sqrt[3]{\frac{n}{m}}$. The m parameter controls the sensitivity of the computed radius to the local variation in density. Very small values of m, $m < 5$, can produce noisy radii, which change drastically between neighboring photons, while large values are more expensive to compute. In practice, we have found that $m = \sqrt{n}$ works well as a default value and this value was used for all our scenes, significantly accelerating the preprocessing step.

Algorithm 1 ConstructBBH(p)

Require: p is a node in a balanced photon map.
Ensure: The subtree at p contains a valid BBH.

$B \Leftarrow$ BoundingBox($p.position$, $p.radius$)
if $p.leftChild$ **then**
  $B \Leftarrow B \cup$ ConstructBBH($p.leftChild$)
end if
if $p.rightChild$ **then**
  $B \Leftarrow B \cup$ ConstructBBH($p.rightChild$)
end if
$p.bbox \Leftarrow B$
return B

Bounding Box Hierarchy Construction. In order to efficiently locate all photons-discs which intersect a ray, we construct a bounding box hierarchy. Heuristics for constructing efficient BBHs have been extensively studied within the context of ray tracing [WMG*07]. However, the performance characteristics of our ray intersections are different than for regular ray tracing since we are interested in all intersections, not just the first intersection along a ray. Furthermore, the best heuristics tend to induce long construction times. We employ a rapid construction scheme by exploiting the information in the photon map kd-tree and re-using that hierarchy for our BBH.

For each photon in the photon map, we compute the bounding box of all descendent photon-discs. The bounding box of each node encloses the node's photon radius and the bounding boxes of its two child nodes. The computation starts at

Scene	Photon Tracing				Fixed Radius Comparison				Adaptive Radius Comparison				
	N	Shoot (s)	Balance (s)	Radius (s)	r	Δt	C (m)	O (m)	r_+	n	Δt	C (m)	O (m)
Cornell	0.4M	1.50	0.30	2.0	0.4	0.40	3:19	3:00	0.6	1.5K	0.80	4:03	3:35
Stage	1M	3.25	0.76	5.0	0.3	0.20	4:21	4:15	0.5	0.5K	0.70	6:38	6:22
Cars	2M	19.0	1.50	2.0	0.4	1.25	1:30	1:30	0.5	1K	1.25	2:02	1:53
Lighthouse	1M	2.83	0.78	6.0	0.4	0.25	1:07	0:59	0.5	0.4K	1.00	1:12	1:05

Table 2: Rendering parameters and timings (in seconds (s) and minutes (m)) for all example scenes. Statistics relating to the photon tracing preprocess are shown in the first table. The middle and right tables compare our method (O) to conventional photon mapping (C) using a fixed width kernel and adaptive width kernel. The r column represents the fixed-width kernel radius, while r_+ is the maximum search radius and the number of nearest neighbors is n.

the leaves and propagates upwards through the tree. This procedure results in a balanced median-split-style BBH, but unlike traditional BBHs our hierarchy contains photons at interior nodes, not just at the leaves. Figure 5 illustrates the relationship between the kd-tree and the BBH. The BBH can be constructed by passing the root of the photon map kd-tree to Algorithm 4.

Given a balanced kd-tree, this linear-time construction procedure is extremely fast and produces BBHs which can be efficiently traversed for nearby photons during rendering. After the BBH is constructed the photon map kd-tree is no longer used and can be freed from memory. Using a BBH and a per-photon radius, an additional 7 floating-point values need to be stored, increasing the size of each photon from 20 bytes to 46 bytes.

The Beam Radiance Estimate. During rendering we estimate the accumulated in-scattered radiance (Equation 5) along viewing rays by locating all photons whose bounding spheres intersect the ray. These photons are found using a standard ray-BBH intersection traversal. The contribution from each photon $(\alpha_i, \mathbf{x}_i, \vec{\omega}_i)$ is accumulated using Equation 19; however, with the variable kernel method, a kernel K_i is defined per photon. This leads to the following radiance estimate

$$\frac{1}{N}\sum_{i=1}^{N} K_i(\mathbf{x},\vec{\omega},s,\mathbf{x}_i,r_i)T_r(\mathbf{x}\leftrightarrow\mathbf{x}'_i)\sigma_s(\mathbf{x}'_i)p(\mathbf{x}_i,\vec{\omega},\vec{\omega}_i)\alpha_i, \quad (21)$$

where $\mathbf{x}'_i = \mathbf{x} + t_i\vec{\omega}$ is the projection of the photon location $\mathbf{x}_i$ onto the ray's direction $\vec{\omega}$, and $t_i = (\mathbf{x}_i - \mathbf{x})\cdot\vec{\omega}$. We define the kernel as

$$K_i(\mathbf{x},\vec{\omega},s,\mathbf{x}_i,r_i) = \begin{cases} r_i^{-2}K_2\left(\frac{d_i}{r_i}\right) & \text{if } d_i \in [0, r_i] \\ 0 & \text{otherwise} \end{cases}, \quad (22)$$

where r_i is the pre-computed radius for photon i. We use Silverman's two-dimensional biweight kernel [Sil86] along the ray, $K_2(x) = 3\pi^{-1}(1-x^2)^2$, where d_i is the shortest distance from photon i to the ray. We chose this kernel because it is smooth, efficient to evaluate, and has local support. Equation 21 is the beam radiance estimate, and it replaces the ray marching computation from conventional photon mapping (second row of Equation 3).

Scene	σ_s	σ_a	g
Cornell	0.225	0.225	0.00
Stage	0.225	0.225	0.75
Cars	0.06	0.015	0.00
Lighthouse	0.24	0.010	0.75

Table 3: Scattering properties of the participating media used in each of the example scenes.

Heterogeneous Media. For homogeneous media, the transmission terms, $T_r(\mathbf{x}\leftrightarrow\mathbf{x}'_i)$, can be computed explicitly for each photon during gathering. Beam gathering in heterogeneous media can also be handled efficiently by first marching along the ray and saving the transmission values in a 1D lookup table. Then, during gathering, each photon's transmission is computed by interpolating within the lookup table. This preprocess needs to be performed independently for each ray, just prior to gathering, so the lookup table can be reused. Furthermore, if single scattering is simulated separately by directly sampling light sources the lookup table can be populated during this marching step.

5. Results

We compared our beam gathering technique against conventional volumetric photon mapping using ray marching. In order to isolate just the performance of the photon gathering methods, we use the photon map for both single and multiple scattering. We compared results on four test scenes: Cars, Lighthouse, Stage, and a Cornell box. For each scene we compare using a fixed gathering radius for both types of estimates, and we also compare the conventional estimate using k-NN to the beam estimate using the VK method. The images were all rendered with a maximum dimension of 1024 pixels with up to 4 samples per pixels on an Intel Core 2 Duo 2.4 GHz machine using one core.

In our experimental setup, we first choose suitable gathering parameters (search radius and number of nearest neighbors n) and render the scenes using our method. We then use the same parameters using conventional photon mapping but tune the minimum step-size Δt to obtain approximately equal render times. Note that Δt is the *minimum* step-size and that exponential stepping is used to sample the ray according to transmission. Finally, we render a high quality result with conventional photon mapping using a very small step size as a "reference."

Figure 6: A comparison between the convention radiance estimate and our beam radiance estimate on the Stage scene with render times provided as (hours:minutes:seconds). Our method (right) produces images with much less noise than an equal time rendering using conventional volumetric photon mapping (middle) for both a fixed radius and an adaptive radius gathering approach. Our method does not require stepping but matches the quality of conventional photon mapping if a very small step size is used (left).

We show visual comparisons of the methods in Figures 6–7. All images of each scene are rendered using the *same photon map*. The only differences in quality and performance are due to the gathering method used. We report the render times and gathering parameters, as well as timings for constructing the photon maps in Table 2. We used the Henyey-Greenstein phase function for all scenes with the medium parameters specified in Table 3.

In all cases, our method produces significantly higher quality images than conventional photon mapping. This is because querying once for all photons along a ray is more efficient than repeatedly querying for photons near numerous samples along the ray. Not surprisingly, we see that the reduced blurring of the adaptive kernel gathering methods is essential for scenes like the Stage and Lighthouse, where concentrated beams of light are visible. However, at the same render time this advantage is difficult to discern in the conventional photon mapping images.

Though the k-NN and VK methods both adapt the width of the kernel to the local photon density, they are distinct approaches which result in similar, but not identical, density estimates. This is what accounts for the small differences in blurring between our adaptive results and the k-NN "reference" images. However, as our results show, the same value of n produces visually comparable renderings using the two methods.

6. Discussion and Future Work

Specific trade-offs between the k-NN and VK methods have been extensively studied within the density estimation literature [Sil86]. For computer graphics, a visual advantage of the VK method is that the estimated function inherits all the differential properties of the smoothing kernel. In contrast, even with smooth kernels, the k-NN method results in estimates with a discontinuous first derivative. The VK method does, unfortunately, require a pre-process to compute the kernel width for each photon. However, the amount of time to compute adaptive radii for each photon using our method is relatively inexpensive (typically less than 1–2% of the total render time). On the other hand, k-NN gathering involves maintaining a priority queue and is much more costly than a fixed radius query. With our method, the gathering procedure is identical whether an adaptive or fixed radius kernel is used and no priority queue is needed.

Though adaptive kernel widths can be a distinct advantage, in uniformly illuminated scenes a fixed radius may be sufficient. While assigning the radii is inexpensive, each resulting BBH node consumes more memory than a kd-tree node due to the additional 7 floating-point values needed to store the bounding box and radius. If memory is limited and a fixed radius search is acceptable, then the regular photon map kd-tree can be used to perform beam gathering by tracing a cylinder through the kd-tree. This approach could be implemented in the spirit of ray-bundle traversals [WMG*07] by bounding the cylinder using 6 planes. This offers the additional benefit of re-using the exact same data structure as conventional photon mapping.

Our photon-disc approach using a BBH is not inherently tied to participating media. VK density estimation could also be applied to photon mapping at surfaces. It would be interesting to see whether a similar technique could be beneficial for surfaces by querying the BBH for all photons that overlap with a surface location. The pilot estimation preprocess may

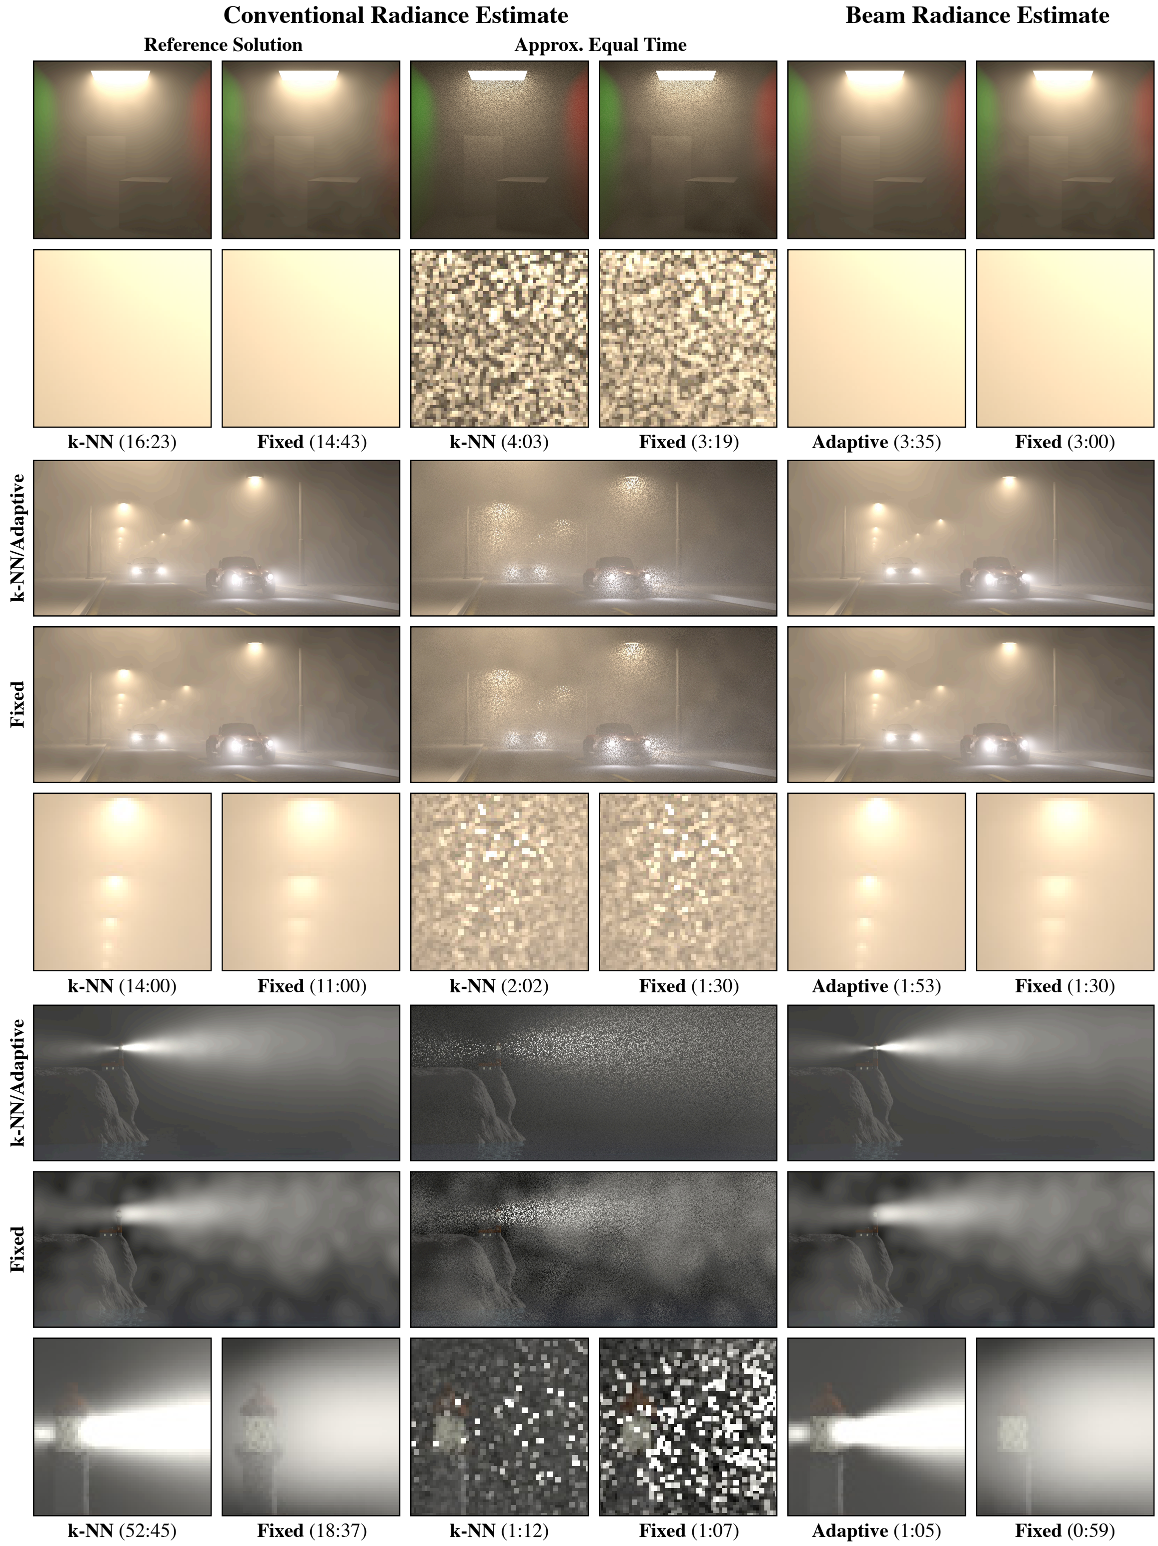

Figure 7: The Cornell box, Cars, and Lighthouse scenes. Render times are shown as (minutes:seconds). For both the fixed and adaptive gathering approaches our method produces noise-free results while conventional photon mapping suffers from significant noise, especially around distant light sources.

also be beneficial at detecting and reducing boundary bias in photon mapping for surfaces.

The advantages of beam gathering are more pronounced in large open environments where ray marching needs to be performed over large distances and where potentially interesting lighting is visible far away. One advantage of the conventional radiance estimate, however, is that a fast preview render can be obtained by using a large step-size. Since no step-size is used for beam gathering, all photons intersecting a ray need to be considered. In fact, these two approaches estimate the lighting integral using different pdfs. The ray marching computation employs exponential stepping, which distributes photon queries with a pdf proportional to the *transmission* term T_r. In contrast, our approach visits every photon that intersects a line and so concentrates effort where the *lighting* is important. Since the radiance seen by the eye is a product of the lighting *and* the transmission, optimally we should concentrate effort according to this product. One way of exploiting this could be to prune the BBH ray traversal using Russian roulette based on some upper-bound on the transmission term and the number of photons in each subtree. Such a scheme could further reduce render times by not considering photons which are too distant and faint to contribute much to the image. We leave this optimization as future work.

It would be interesting to explore what other useful radiance estimate could be formed using the measurement equation formulation. For instance, Cammarano and Jensen [CJ02] considered the problem of estimating the density of photons within four dimensional "cylinders" to properly simulate motion blurred caustics. This process could easily be formulated as a measurement by defining the importance function over space-time.

Photon splatting for participating media was presented by Boudet et al. [BPPP05], where the conventional photon mapping method needed to be meticulously transformed into a splatting algorithm. An extra benefit of our reformulation is that Equation 21 can immediately be seen as a splatting operation and could therefore be efficiently implemented in graphics hardware. Furthermore, whereas Boudet et al. only considered fixed-size splats, our VK approach could easily be used to adapt the splat size to the local photon density. A splatting approach could also be used in software to accelerate the computation for eye rays; however, the more general beam gathering would need to be used for any secondary rays.

7. Conclusion

In this paper, we showed how volumetric photon mapping can be expressed as a solution to the measurement equation. This formulation showed that *any* measurement of radiance within participating can be estimated using the photon map. We applied this formulation by using the photon map to directly estimate accumulated in-scattered radiance along rays. This approach was implemented using an efficient beam gathering method which can be used for both fixed and adaptive width kernels. The resulting algorithm produces images with significantly less noise than conventional volumetric photon mapping using the same render time.

8. Acknowledgements

This work was supported in part by NSF grant CPA 0701992.

References

[BMP77] Breiman L., Meisel W., Purcell E.: Variable kernel estimates of multivariate densities. *Technometrics 19*, 2 (May 1977), 135–144.

[BPPP05] Boudet A., Pitot P., Pratmarty D., Paulin M.: Photon splatting for participating media. In *GRAPHITE '05: Proceedings of the 3rd international conference on Computer graphics and interactive techniques in Australasia and South East Asia* (New York, NY, USA, 2005), ACM Press, pp. 197–204.

[Cha60] Chandrasekar S.: *Radiative Transfer*. Dover Publications, New York, 1960.

[CJ02] Cammarano M., Jensen H. W.: Time dependent photon mapping. In *EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2002), Eurographics Association, pp. 135–144.

[CPP*05] Cerezo E., Pérez F., Pueyo X., Seron F. J., cois X. Sillion F.: A survey on participating media rendering techniques. *The Visual Computer 21*, 5 (June 2005), 303–328.

[JC98] Jensen H. W., Christensen P. H.: Efficient simulation of light transport in scences with participating media using photon maps. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1998), ACM Press, pp. 311–320.

[JMLH01] Jensen H. W., Marschner S. R., Levoy M., Hanrahan P.: A practical model for subsurface light transport. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), ACM Press, pp. 511–518.

[JZJ08] Jarosz W., Zwicker M., Jensen H. W.: *The Beam Radiance Estimate for Volumetric Photon Mapping*. Tech. Rep. CS2008-0914, University of California, San Diego, 2008.

[Kaj86] Kajiya J. T.: The rendering equation. In *SIGGRAPH '86: Proceedings of the 13th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1986), ACM Press, pp. 143–150.

[KH84] Kajiya J. T., Herzen B. P. V.: Ray tracing volume densities. In *SIGGRAPH '84: Proceedings of the 11th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1984), ACM Press, pp. 165–174.

[LW96] Lafortune E. P., Willems Y. D.: Rendering Participating Media with Bidirectional Path Tracing. In *Proceedings of the eurographics workshop on Rendering techniques '96* (London, UK, 1996), Springer-Verlag, pp. 91–100.

[PARN04] Premoze S., Ashikhmin M., Ramamoorthi R., Nayar S. K.: Practical rendering of multiple scattering effects in participating media. In *Rendering Techniques* (2004), pp. 363–373.

[PKK00] Pauly M., Kollig T., Keller A.: Metropolis light transport for participating media. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000* (London, UK, 2000), Springer-Verlag, pp. 11–22.

[PM93] Pattanaik S. N., Mudur S. P.: Computation of global illumination in a participating medium by Monte Carlo simulation. *The Journal of Visualization and Computer Animation 4*, 3 (July–Sept. 1993), 133–152.

[RT87] Rushmeier H. E., Torrance K. E.: The zonal method for calculating light intensities in the presence of a participating medium. *Computer Graphics (SIGGRAPH '87 Proceedings) 21*, 4 (July 1987), 293–302.

[Sil86] Silverman B.: *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, New York, NY, 1986.

[SRNN05] Sun B., Ramamoorthi R., Narasimhan S. G., Nayar S. K.: A practical analytic single scattering model for real-time rendering. *ACM Trans. Graph. 24*, 3 (2005), 1040–1049.

[Sta95] Stam J.: Multiple Scattering as a Diffusion Process. In *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)* (New York, NY, 1995), Hanrahan P. M., Purgathofer W., (Eds.), Springer-Verlag, pp. 41–50.

[Vea98] Veach E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, 1998.

[WMG*07] Wald I., Mark W. R., Günther J., Boulos S., Ize T., Hunt W., Parker S. G., Shirley P.: State of the art in ray tracing animated scenes. In *STAR Proceedings of Eurographics 2007* (Prague, Czech Republic, September 2007), Eurographics Association, pp. 0–0.