# Manifold Exploration: <br> A Markov Chain Monte Carlo technique for rendering scenes with difficult specular transport 

Expanded Technical Report

Wenzel Jakob Steve Marschner<br>Cornell University

May 7, 2012


#### Abstract

It is a long-standing problem in unbiased Monte Carlo methods for rendering that certain difficult types of light transport paths, particularly those involving viewing and illumination along paths containing specular or glossy surfaces, cause unusably slow convergence. In this paper we introduce Manifold Exploration, a new way of handling specular paths in rendering. It is based on the idea that sets of paths contributing to the image naturally form manifolds in path space, which can be explored locally by a simple equation-solving iteration. This paper shows how to formulate and solve the required equations using only geometric information that is already generally available in ray tracing systems, and how to use this method in in two different Markov Chain Monte Carlo frameworks to accurately compute illumination from general families of paths. The resulting rendering algorithms handle specular, near-specular, glossy, and diffuse surface interactions as well as isotropic or highly anisotropic volume scattering interactions, all using the same fundamental algorithm. An implementation is demonstrated on a range of challenging scenes and evaluated against previous methods.


Remark: This technical report follows the structure of the main paper. The main difference is that it contains additional discussion in Sections 3.3 and 6, a second appendix with C++ code, as well as MEMLT results in Section 8.


Figure 1: Two views of an interior scene with complex specular and near-specular transport, rendered using manifold exploration path tracing: (a) Refractive, reflective, and glossy tableware, (b) A rough brass luminaire with 24 glass-enclosed bulbs used to light the previous closeup.

## 1 Introduction

Certain classes of light paths have traditionally been a source of difficulty in conducting Monte Carlo simulations of light transport. A well-known example is specular-diffuse-specular paths, such a tabletop seen through a drinking glass sitting on it, a bottle containing shampoo or other translucent liquid, or a shop window viewed and illuminated from outside. Even in scenes where these paths do not cause dramatic lighting effects, their presence can lead to unusably slow convergence in renderers that attempt to account for all transport paths.

Furthermore, wherever ideally specular paths are troublesome, nearly specular paths involving glossy materials are also troublesome. They can be more problematic, in fact, because they elude special mechanisms designed to handle specular interactions. These glossy paths have become more important as material models have evolved, and we would prefer to handle them using natural generalizations of strategies for specular surfaces, rather than generalizations of strategies for diffuse surfaces.

Finding these paths efficiently is a key problem of light transport simulations. In this paper, we show how a Markov Chain Monte Carlo method can be used to efficiently render paths with illumination and/or viewing through arbitrary chains of specular or glossy reflections or refractions. The idea is that sets of paths contributing to the image naturally form manifolds in path space, and using a simple equation-solving iteration it is easy to move around on these manifolds. We show how to formulate and solve the equations required to find specular paths, using only geometric information that is already generally available in ray tracing systems. This solution method is then used to define a Markov Chain, which has the right properties to be applied in a Metropolis-Hastings algorithm to compute lighting through very general families of paths that can involve specular, near-specular, glossy, and diffuse surface interactions as well as isotropic or highly anisotropic volume scattering interactions.

We begin by discussing prior work on which our method is built, then in Section 3 we develop the theory of the specular manifold, used to handle interactions with ideal specular (polished) surfaces, and offset specular manifolds, which provide a graceful generalization to near-specular materials. In Section 5, we derive an algorithm to move from one path to another on a specular or offset specular manifold, use it to build an algorithm that generates Markov sequences in path
space, and show how this can be used in the Metropolis Light Transport or Energy Redistribution Path Tracing frameworks to provide methods for rendering scenes with any kind of light transport. We go on in the following sections to extend the theory and algorithm to the case of participating media, and we finish by showing results and comparisons in Section 8.

## 2 Prior work

Simulating light transport has been a major effort in computer graphics for over 25 years, beginning with the complementary approaches of finite-element simulation, or radiosity [Goral et al. 1984], and ray-tracing [Whitted 1980]. The introduction of Monte Carlo methods for ray tracing [Cook et al. 1984], followed by Kajiya's formulation of global illumination in terms of the Rendering Equation [Kajiya 1986], established the field of Monte Carlo global illumination. Unbiased sampling methods, in which each pixel in the image is a random variable with an expected value exactly equal to the solution of the Rendering Equation, started with Kajiya's original path tracing method and continued with bidirectional path tracing [Lafortune and Willems 1993; Veach and Guibas 1994], in which light transport paths can be constructed partly from the light and partly from the eye, and the seminal Metropolis Light Transport [Veach and Guibas 1997] algorithm, which uses bidirectional path tracing methods in a Markov Chain Monte Carlo framework.

Various two-pass methods use a particle-tracing pass that sends energy out from light sources in the form of "photons" that are traced through the scene [Shirley et al. 1995; Jensen 1996; Walter et al. 1997]. and stored in a spatial data structure. The second pass then renders the image using ray tracing, making use of the stored particles to estimate illumination by density estimation. [Jensen and Christensen 1998; Jarosz et al. 2008; Hachisuka and Jensen 2009].

Photon mapping and other two-pass methods are characterized by storing an approximate representation of some part of the illumination in the scene, which requires assumptions about the smoothness of illumination distributions. On one hand, this enables rendering of some modes of transport that are difficult for unbiased methods, since the exact paths by which light travels do not need to be found; separate paths from the eye and light that end at nearby points suffice under assumptions of smoothness. However, this smoothness assumption inherently leads to smoothing errors in images: the results are biased, in the Monte Carlo sense. Glossy-glossy transport, without a sufficiently diffuse surface on which to store photons, is challenging to handle with photon maps, since large numbers of photons must be collected to adequately sample position-direction space. Some photon mapping variants avoid this by treating glossy materials as specular, but this means that the the resulting method increasingly resembles path tracing as the number of rough surfaces in the input scene grows.

### 2.1 Specular reflection geometry

Separate from work on global illumination algorithms, various research has examined the properties of specular reflection paths. Mitchell and Hanrahan [1992] devised a method to compute irradiance from implicitly defined reflectors, using Fermat's principle with interval Newton's method to locate all reflection paths from a source to a point. Walter et al. [2009] proposed a related method that computes the singly scattered radiance within a refractive object with triangle mesh boundaries. Like these works, our method searches for specular paths. But because it does so within the neighborhood of a given path, it avoids the complexities and constraints entailed by a full global
search. Another difference is that our manifold formalism can be used to build a fully general rendering system that is not limited to the specular paths that prompted its design.

The widely used method of ray differentials for texture antialiasing [Igehy 1999] also reasons about the local structure of a set of reflected paths - in this case, paths from the eye. Igehy's approach requires elementary local differential information only, in the form of derivatives of surface normals, and does not require global surface descriptions as Mitchell and Hanrahan's method does. Manifold exploration requires the same geometric information, and can thus be implemented in most modern ray tracing sytems.

The analysis of reflection geometry presented by Chen and Arvo [2000a; 2000b] is closest to the mathematics underlying our proposed methods. Their work relies on a characterization of specular paths via Fermat's principle, which asserts that light travels along paths whose optical length (i.e. the propagation time) constitutes a local extremum amongst neighboring paths. Using Lagrange multipliers, the authors derive a path Jacobian and path Hessian with respect to perturbations of the endpoint of a path and use it to accelerate the interactive display of reflections on curved surfaces.

Our local characterization of specular paths is equivalent to Fermat's principle, and the path Jacobian is related to the derivatives that we propose to use to define tangent spaces to the specular manifold while solving for path transitions. However, the use of this derivative and the goals of the research are entirely different: in their case, estimating changes to viewing paths, and in our case, tracking the evolution of specular paths in a very general context, as part of an unbiased rendering system.

A less related but relevant idea is integrating over continuous paths in volume rendering applications, known as the path integral formulation of radiative transfer [Premoze et al. 2003]. This work, with its implications for the concentration of transport in path space, suggests the possibility of using MCMC to integrate over multiple-scattering paths in volumes, as discussed in Section 7.

### 2.2 Markov Chain Monte Carlo in rendering

The Metropolis Light Transport algorithm mentioned above introduced the tools of Markov Chain Monte Carlo (MCMC) to rendering. A Markov chain is a sequence of points in a state space in which the probability of a state appearing at a given position in the sequence depends only on the previous state. The basic idea of MCMC, first proposed by Metropolis et al. [1953], is to define a Markov chain that has the function to be integrated as its stationary distribution, meaning that if the chain is run for a long time the distribution of states it visits will be proportional to the desired distribution. The following brief introduction follows [Liu 2001].

Defining a Markov chain amounts to defining a transition rule: a process for selecting a new state $\mathbf{x}_{+}$randomly, in a way that depends on the current state $\mathbf{x}$. Metropolis et al. provided a way to take a transition rule that may not produce the desired stationary distribution $\pi(\mathbf{x})$ and turn it into one that does. Given a method for sampling a proposal distribution $T\left(\mathbf{x}, \mathbf{x}^{\prime}\right)$, the Metropolis transition rule operates in two steps:

1. Choose $\mathbf{x}^{\prime}$ according to the probability distribution $T\left(\mathbf{x}, \mathbf{x}^{\prime}\right)$.
2. $\mathrm{x}_{+}= \begin{cases}\mathrm{x}^{\prime} & \text { with probability } \min \left(1, \pi\left(\mathrm{x}^{\prime}\right) / \pi(\mathrm{x})\right) \\ \mathrm{x} & \text { otherwise }\end{cases}$

In step 1 we say $\mathbf{x}^{\prime}$ is proposed as the next state, and in step 2 it is either accepted and becomes the next state, or it is rejected and the next state repeats the previous one. The probability $\min \left(1, \pi\left(\mathbf{x}^{\prime}\right) / \pi(\mathbf{x})\right)$ is known as the acceptance probability.

If this chain is able to pass from any state in the domain to another using a finite expected number of steps, and if the length of this journey is in a sense "irregular" (i.e. aperiodic), it is referred to as ergodic, and the chain's limiting distribution is guaranteed to converge to $\pi$. This relatively mild criterion is usually satisified by choosing a transition rule with global support.

The original Metropolis algorithm only works when $T\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=T\left(\mathbf{x}^{\prime}, \mathbf{x}\right)$. Hastings [1970] proposed a new acceptance probability:

$$
\begin{equation*}
r\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=\min \left\{1, \frac{\pi\left(\mathbf{x}^{\prime}\right) T\left(\mathbf{x}^{\prime}, \mathbf{x}\right)}{\pi(\mathbf{x}) T\left(\mathbf{x}, \mathbf{x}^{\prime}\right)}\right\} \tag{1}
\end{equation*}
$$

which relaxes the symmetry restriction to one of symmetric support: $T\left(\mathbf{x}, \mathbf{x}^{\prime}\right)$ must be nonzero exactly when $T\left(\mathbf{x}^{\prime}, \mathbf{x}\right)$ is nonzero. The Metropolis-Hastings algorithm is the starting point for MCMC rendering methods.

In the rendering context, the state space is the space of all paths through the scene, points in the space are paths, and the desired probability distribution over paths is proportional to their contribution to the rendered image (i.e. the amount of illumination they carry to the camera). The final image is the projection of the path distribution into the image plane.

At the core of an MCMC rendering algorithm is an implementation of a transition rule, and any rule with symmetric support is admissible. But to avoid very low acceptance probabilities, which lead to poor performance, it is desirable for the transition probability to approximate the contribution: that is, paths with more light flowing along them should be chosen more often. Veach's [1997] transition rule is based on a set of mutations that change the structure of the path and perturbations that move the vertices by small distances while preserving the structure, both using the building blocks of bidirectional path tracing to sample paths. Using his transition rule to run long chains leads to the Metropolis Light Tranport algorithm (MLT). Kelemen et al. [2002] later proposed a transition rule based on changing the random numbers used to sample paths by bidirectional path tracing.

Considerable research activity has extended Metropolis light transport in various ways. Pauly et al. [2000] proposed a perturbation rule for rendering participating media with single scattering. Other projects include Metropolis Instant Radiosity [Segovia et al. 2007], Population Monte Carlo rendering [Lai et al. 2007], and Replica Exchange light transport [Kitaoka et al. 2009]. Recently, two groups [Chen et al. 2011; Hachisuka and Jensen 2011] have combined the transition rule of Kelemen et al. with photon mapping to obtain robust methods based on density estimation.

However, to generate proposals, all of these algorithms ultimately rely on local path sampling strategies (i.e. path tracing). Specifically, they choose the next interaction vertex along a light path by sampling from a directional distribution associated with the current vertex, followed by an intersection search. Our method introduces a new kind of transition rule with different properties.

The original MLT algorithm and subsequent variants all render an image by running a Markov chain for a long (e.g. $>10^{6}$ ) sequence of steps, and they guarantee ergodicity using a transition rule that can generate any path in the domain with some probability. The Energy Redistribution path tracing (ERPT) technique [Cline et al. 2005], which is readily adapted to work with our method, is an interesting departure from that approach. It draws on the property that the

Metropolis-Hastings algorithm (1) preserves the stationary distribution of samples even if the underlying transition rule is not ergodic (e.g. when it cannot reach parts of path space). To obtain coverage, ERPT first samples a large set of paths via path tracing and runs Markov chains for short bursts ( $\approx 10^{3}$ steps) starting at each sample. The relaxation of the ergodicity requirement makes it possible to explore paths in a very local fashion.

## 3 Path space manifolds

Manifold exploration is a technique for integrating the contributions of sets of specular or nearspecular illumination paths to the rendered image of a scene. The general approach applies to surfaces and volumes and to ideal and non-ideal (glossy) specular surfaces. In this section we begin by examining the manifold defined by ideal specular reflection or refraction, in the setting of surfaces without participating media. In the following section we develop this theory into a rendering method for scenes combining ideal specular surfaces with fairly diffuse surfaces. We will then go on to generalize the method to glossy surfaces, then to generalize the theory to encompass participating media and to extend the method to handle media with both isotropic and highly directional scattering.

### 3.1 Path space

The resulting techniques are all based on the path integral formulation of light transport, described by Veach [1997] and others, in which the value of each pixel in the image is an integral of a contribution function over path space. Denoting the union of all surfaces in the scene by $\mathcal{M}$, each sequence $\mathbf{x}_{1} \ldots \mathbf{x}_{n}$ of at least two points in $\mathcal{M}$ is a path along which light may travel. Thus path space is

$$
\begin{gather*}
\mathcal{P}=\bigcup_{n=2}^{\infty} \mathcal{P}_{k}  \tag{2}\\
\mathcal{P}_{n}=\left\{\mathbf{x}_{1} \ldots \mathbf{x}_{n} \mid \mathbf{x}_{1}, \ldots, \mathbf{x}_{n} \in \mathcal{M}\right\} \tag{3}
\end{gather*}
$$

As detailed by Veach, the value of pixel $j$ is

$$
I_{j}=\int_{\mathcal{P}} f_{j}(\overline{\mathbf{x}}) d \mu(\overline{\mathbf{x}})
$$

where $\overline{\mathbf{x}}=\mathbf{x}_{1} \ldots \mathbf{x}_{n}$ denotes an element of $\mathcal{P}$ and $\mu$ is the product measure derived from the area measure on $\mathcal{M}$. The contribution function $f_{j}$ is a product of terms, one for each vertex and each edge of the path:

$$
\begin{align*}
f_{j}\left(\mathbf{x}_{1} \cdots \mathbf{x}_{n}\right)= & L_{e}\left(\mathbf{x}_{1} \rightarrow \mathbf{x}_{2}\right) \\
& {\left[\prod_{k=2}^{n-1} G\left(\mathbf{x}_{k-1} \leftrightarrow \mathbf{x}_{k}\right) f_{s}\left(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_{k} \rightarrow \mathbf{x}_{k+1}\right)\right] G\left(\mathbf{x}_{n-1} \leftrightarrow \mathbf{x}_{n}\right) W_{e}^{(j)}\left(\mathbf{x}_{n-1} \rightarrow \mathbf{x}_{n}\right) . } \tag{4}
\end{align*}
$$



$$
G\left(\mathbf{x}_{2} \leftrightarrow \mathbf{x}_{3}\right)=\left|\frac{d \omega_{2}^{\perp}}{d A_{3}}\right|
$$



$$
\left|\frac{d \omega_{2}^{\perp}}{d A_{6}}\right|=\left|\frac{d \omega_{2}^{\perp}}{d A_{3}} \frac{d A_{3}}{d A_{6}}\right|
$$

Figure 2: The geometry factor (left) and the generalized geometry factor (right) are both derivatives of projected solid angle at one end with respect to area at the far end.

Here $L_{e}$ is emitted radiance, $W_{e}^{(j)}$ is importance for the $j$ th pixel, $f_{s}$ is the BSDF at $\mathbf{x}_{k}$ for the geometry defined by $\mathbf{x}_{k-1}, \mathbf{x}_{k}$, and $\mathbf{x}_{k+1}$, and $G$ is the geometry factor:

$$
\begin{equation*}
G(\mathbf{x} \leftrightarrow \mathbf{y})=\frac{|N(\mathbf{x}) \cdot \overrightarrow{\mathbf{x y}}||N(\mathbf{y}) \cdot \overrightarrow{\mathbf{y}} \mathbf{x}|}{\|\mathbf{x}-\mathbf{y}\|^{2}} V(\mathbf{x} \leftrightarrow \mathbf{y}) \tag{5}
\end{equation*}
$$

where $N(\mathbf{a})$ is the surface normal at $\mathbf{a}$ and $V(\mathbf{x} \leftrightarrow \mathbf{y})$ is the visibility function.

### 3.2 Motivating examples

This formulation is the basis for several rendering methods including bidirectional path tracing and Metropolis light transport. However, in the presence of ideal specular reflection, some difficulties arise, which are normally sidestepped in the transition from theory to algorithm, but which we prefer to confront directly. When some surface interactions are specular, the entire contribution to the path space integral is from paths that obey specular reflection or refraction geometry, and the set of such paths is lower in dimension than the full path space. For instance, consider a family of paths of the form LDSDE (in Heckbert's [1990] notation) with one specular reflection vertex. These paths belong to the $\mathcal{P}_{5}$ component of $\mathcal{P}$, but the paths that contribute all have the property

$$
\left(\overrightarrow{\mathbf{x}_{3} \mathbf{x}_{2}}+\overrightarrow{\mathbf{x}_{3} \mathbf{x}_{4}}\right) \| N\left(\mathbf{x}_{3}\right)
$$

that is, the half-vector at $\mathbf{x}_{3}$ is in the direction of the normal. This places two constraints on the path, meaning that all contributing paths lie on a manifold of dimension 8 embedded in $\mathcal{P}_{5}$, which is of dimension 10. The integral is more naturally expressed as an integral over the manifold, rather than as an integral over the whole path space.

To compute illumination due to specular paths, we use a local parameterization of the manifold in terms of the positions of all nonspecular vertices on the path:

$$
\iiint \int_{\mathcal{M}^{4}} f\left(\mathbf{x}_{1} \ldots \mathbf{x}_{5}\right) d \mathbf{x}_{1} d \mathbf{x}_{2} d \mathbf{x}_{4} d \mathbf{x}_{5}
$$

Note the missing integral over $\mathbf{x}_{3}$, the specular vertex. The contribution function $f$ still has
the same form, a product of terms corresponding to vertices and edges of the path, but the BSDF values at the specular vertex are replaced by (unitless) specular reflectance values, and the geometry factors for the two edges involving the specular vertex are replaced by a single generalized geometry factor that we will denote $G\left(\mathbf{x}_{2} \leftrightarrow \mathbf{x}_{3} \leftrightarrow \mathbf{x}_{4}\right)$.

The standard geometry factor for a non-specular edge is the derivative of projected solid angle at one vertex with respect to area at the other vertex, and the generalized geometry factor is defined analogously: the derivative of solid angle at one end of the specular chain with respect to area at the other end of the chain, considering the path as a function of the positions of the endpoints. Figure 2 illustrates this for a more complex path involving a chain of three specular vertices. We will explain below how $G$ can be easily computed from the differential geometry of the specular manifold.

### 3.3 Specular manifold geometry

In the general case, each path of length $k$ belongs to a class in $\{D, S\}^{k}$ based on the classification of each of its vertices. (In this scheme point or orthographic cameras, and point or parallel lights, are denoted S, while finite-aperture cameras and area lights are D.) Each S surface vertex has an associated constraint that involves its position and the position of the preceding and following vertices:

$$
\mathbf{c}_{i}\left(\mathbf{x}_{i-1}, \mathbf{x}_{i}, \mathbf{x}_{i+1}\right)=\mathbf{0}
$$

The constraint function computes a half-vector at vertex $i$ and projects it into the shading tangent space; the resulting 2 -vector is zero when the half-vector is parallel to the normal. By making use of the generalized half-vector of Walter et al. [2007], both reflection and refraction can be handled by a single constraint function:

$$
\begin{gather*}
\mathbf{c}_{i}\left(\mathbf{x}_{i-1}, \mathbf{x}_{i}, \mathbf{x}_{i+1}\right)=T\left(\mathbf{x}_{i}\right)^{T} h\left(\mathbf{x}_{i}, \overrightarrow{\mathbf{x}_{i} \mathbf{x}_{i-1}}, \overrightarrow{\mathbf{x}_{i} \mathbf{x}_{i+1}}\right)  \tag{6}\\
h(\mathbf{x}, \mathbf{v}, \mathbf{w})= \tag{7}
\end{gather*}
$$

where $T(\mathbf{x})$ is a matrix whose columns form a basis for the shading tangent plane at $\mathbf{x}$, and $\eta(\mathbf{x}, \mathbf{v})$ denotes the refractive index associated with the ray $(\mathbf{x}, \mathbf{v})$. This generalized geometry factor is related to the "extended form factor" discussed by Sillion and Puech [1989].

Specular endpoint vertices also introduce constraints that depend on their type: for instance, the position of a vertex on a point emitter must remain fixed (e.g. $\mathbf{x}_{1}=$ const). When the emitter is directional, it is the outgoing direction that is constrained, and so on.

From now on, we implicitly identify each vertex $\mathbf{x}_{i}$ with an associated point in $\mathbb{R}^{2}$ using local parameterizations of $\mathcal{M}$. These parameterizations may be defined on arbitrarily small neighborhoods, since only their derivatives are relevant to what follows.

With this, the constraints for a length- $n$ path with $p$ specular vertices can be stacked together into a function $C: \mathbb{R}^{2 n} \rightarrow \mathbb{R}^{2 p}$, and the specular manifold is simply the set

$$
\begin{equation*}
\mathcal{S}=\{\overline{\mathbf{x}} \mid C(\overline{\mathbf{x}})=\mathbf{0}\} \tag{8}
\end{equation*}
$$

Expressing $\mathcal{S}$ using a constraint in this way makes it convenient to work with neighborhoods of a particular path. The Implicit Function Theorem [Spivak 1965] guarantees the existence of a

(c) Constraint Jacobian

$$
\begin{aligned}
& C(\overline{\mathbf{x}})=0 \quad \text { where } \\
& C=\left[\begin{array}{c}
\mathbf{c}_{2}\left(\mathbf{x}_{1}, \mathbf{x}_{2}, \mathbf{x}_{3}\right) \\
\vdots \\
\mathbf{c}_{6}\left(\mathbf{x}_{5}, \mathbf{x}_{6}, \mathbf{x}_{7}\right)
\end{array}\right]
\end{aligned}
$$

(b) Associated constraints

$$
\frac{\partial\left[\begin{array}{c}
\mathbf{x}_{2} \\
\vdots \\
\mathbf{x}_{6}
\end{array}\right]}{\partial \mathbf{x}_{1}}=-A^{-1} B_{1}
$$

$$
\frac{\partial\left[\begin{array}{c}
\mathbf{x}_{2} \\
\vdots \\
\mathbf{x}_{6}
\end{array}\right]}{\partial \mathbf{x}_{7}}=-A^{-1} B_{7}
$$

(d) Tangent space

Figure 3: The linear system used to compute the tangent space to the specular manifold, also known as the derivative of a specular chain with respect to its endpoints.
parameterization of the manifold, in the neighborhood of any path $\overline{\mathbf{x}}$ that is nonsingular (in the sense explained below). This parameterization is a function $q: \mathbb{R}^{2(n-p)} \rightarrow \mathbb{R}^{2 p}$ that determines the positions of all the specular vertices from the positions of all the nonspecular vertices. Furthermore, the derivative of $q$, which gives us the tangent space to the manifold at $\overline{\mathbf{x}}$, is simple to compute from the derivative of $C$.

For the specifics we restrict ourselves to the case of a single chain of specular vertices with non-specular vertices (surfaces, cameras, or light sources) at the ends, which suffices to cover most cases by applying it separately to mutiple specular chains along a path. Paths with specular endpoints are handled with simple variations of this scheme. Number the vertices in the chain $\mathbf{x}_{1}, \ldots, \mathbf{x}_{k}$, with $\mathbf{x}_{1}$ and $\mathbf{x}_{k}$ being the (non-specular) endpoints of the segment and the remaining $k-2$ vertices being specular. In this case $C: \mathbb{R}^{2 k} \rightarrow \mathbb{R}^{2(k-2)}$, and the derivative $\nabla C$ is a matrix of $k-2$ by $k 2$-by- 2 blocks, with a tridiagonal structure (Figure 3).

The Implicit Function Theorem gives us a parameterization of the manifold in terms of any 2 vertices, and if we pick $\mathbf{x}_{1}$ and $\mathbf{x}_{k}$ this simply says that the path, in a neighborhood of the
current path ${ }^{1}$, is a function of the two endpoints. Furthermore, it also tells us the derivative of that parameterization, which is to say, the derivative of all the specular vertices' positions with respect to the positions of the endpoints. If we block the derivative $\nabla C$, as shown in the figure, into 2-column matrices $B_{1}$ and $B_{k}$ for the first and last vertices and a square matrix $A$ for the specular chain, then this tangent space to the manifold is

$$
T_{\mathcal{S}}(\overline{\mathbf{x}})=-A^{-1}\left[\begin{array}{ll}
B_{1} & B_{k}
\end{array}\right]
$$

This matrix is $k-2$ by 2 blocks in size, and each block gives the derivative of one vertex (in terms of its own tangent frame) with respect to one endpoint.

The matrix $A$ fails to be invertible when the specular vertices are not even locally unique for given endpoints, which means that one endpoint is on a caustic for light emitting from the other endpoint.

We use $T_{\mathcal{S}}(\overline{\mathbf{x}})$ for two things: to navigate on the manifold and to compute the generalized geometry factor. The right two or left two columns of $T_{\mathcal{S}}(\overline{\mathbf{x}})$ are useful for updating the specular chain with $\mathbf{x}_{1}$ or $\mathbf{x}_{k}$ held fixed, respectively, as discussed in the next section.

The top-right or bottom-left block of $T_{\mathcal{S}}(\overline{\mathbf{x}})$ can be used to compute the generalized geometry factor as follows. Assuming orthonormal parameterizations ${ }^{2}$, the determinant of the top-right block gives the ratio of an infinitesimal area at $\mathbf{x}_{k}$ to its reflection/refraction, as observed from $\mathbf{x}_{1}$, measured on the surface at $\mathbf{x}_{2}$. To convert this to a ratio of area at $\mathbf{x}_{k}$ to solid angle at $\mathbf{x}_{1}$, we multiply this determinant by the ordinary geometry factor $G\left(\mathbf{x}_{1} \leftrightarrow \mathbf{x}_{2}\right)$; this product is the generalized geometry factor $G\left(\mathbf{x}_{1} \leftrightarrow \cdots \leftrightarrow \mathbf{x}_{k}\right)$. More succinctly,

$$
\begin{align*}
G\left(\mathbf{x}_{1} \leftrightarrow \cdots \leftrightarrow \mathbf{x}_{k}\right) & =\left|P_{2} A^{-1} B_{k}\right| G\left(\mathbf{x}_{1} \leftrightarrow \mathbf{x}_{2}\right)  \tag{9}\\
& =\left|P_{k-1} A^{-1} B_{1}\right| G\left(\mathbf{x}_{k-1} \leftrightarrow \mathbf{x}_{k}\right),
\end{align*}
$$

where $P_{i}$ is a 2 by $2(k-2)$ matrix that projects onto the two dimensions associated with vertex $i$.
A useful property of this framework is its reliance on local information that is easily provided in ray tracing-based rendering systems. To compute the blocks of the $A$ and $B$ matrices, we must have access to the partial derivatives of position and shading normal with respect to any convenient parameterization of the surfaces, along with the refractive indices of all objects. These are exactly the same quantities also needed to trace ray differentials through refractive boundaries, which is part of many mature ray tracing-based rendering systems. A consequence of the simple form of the constraint (8) is that our technique works with any object that can provide such local information, including implicitly defined shapes or triangle meshes with shading normals.

These theoretical results about the structure of the specular manifold can be used in an algorithm to solve for specular paths, which we discuss in Section 4.

A simple example: We now demonstrate how to compute the tangent space and generalized geometric term associated with a simple example path shown in Figure 4. This is a path with three vertices; $\mathbf{x}_{1}$ and $\mathbf{x}_{3}$ are planar endpoints used to parameterize the manifold, and $\mathbf{x}_{2}$ lies on a

[^0]

Figure 4: An example path with three vertices
specularly reflecting cylinder that stretches out along the $z$-axis.
Assuming that this path is encountered in a ray tracer similar to PBRT [Pharr and Humphreys 2004] or Mitsuba [Jakob 2010], the rendering system will associate an intersection data record with each vertex containing (amongst other things) the position and surface normal, as well as derivatives thereof along some parameterization. Suppose that the information provided is as follows:

$$
\begin{aligned}
& \mathbf{x}_{1}=(-1,2,0), \partial_{u} \mathbf{x}_{1}=(-1,0,0), \partial_{v} \mathbf{x}_{1}=(0,0,1), \mathbf{n}_{1}=(0,-1,0), \partial_{u} \mathbf{n}_{1}=0, \quad \partial_{v} \mathbf{n}_{1}=0 \\
& \mathbf{x}_{2}=(0,1,0), \quad \partial_{u} \mathbf{x}_{2}=(1,0,0), \quad \partial_{v} \mathbf{x}_{2}=(0,0,1), \quad \mathbf{n}_{2}=(0,1,0), \quad \partial_{u} \mathbf{n}_{2}=(1,0,0), \partial_{v} \mathbf{n}_{2}=0 \\
& \mathbf{x}_{3}=(1,2,0), \quad \partial_{u} \mathbf{x}_{3}=(1,0,0), \quad \partial_{v} \mathbf{x}_{3}=(0,0,1), \mathbf{n}_{3}=(0,-1,0), \partial_{u} \mathbf{n}_{3}=0, \quad \partial_{v} \mathbf{n}_{3}=0
\end{aligned}
$$

Note that the parameterizations above are locally orthonormal - this is generally preferable, since it simplifies many computations involving $\nabla C$. In practice, we can simply apply a suitable linear transformation to the tangents and normal derivatives to to make them correspond to a locally orthonormal chart. Let us now compute the entries of $\nabla C$.

Recall that the constraint associated with vertex $\mathbf{x}_{2}$ is

$$
C\left(\mathbf{x}_{1}, \mathbf{x}_{2}, \mathbf{x}_{2}\right)=T\left(\mathbf{x}_{2}\right)^{T}\left(\frac{\mathbf{x}_{1}-\mathbf{x}_{2}}{\left\|\mathbf{x}_{1}-\mathbf{x}_{2}\right\|}+\frac{\mathbf{x}_{3}-\mathbf{x}_{2}}{\left\|\mathbf{x}_{3}-\mathbf{x}_{2}\right\|}\right) /\left\|\frac{\mathbf{x}_{1}-\mathbf{x}_{2}}{\left\|\mathbf{x}_{1}-\mathbf{x}_{2}\right\|}+\frac{\mathbf{x}_{3}-\mathbf{x}_{2}}{\left\|\mathbf{x}_{3}-\mathbf{x}_{2}\right\|}\right\|
$$

To compute derivatives of this expression, we can parameterize all needed ingredients to first order using the above data records, i.e.

$$
\begin{gathered}
\mathbf{x}_{i}\left(u_{i}, v_{i}\right)=\mathbf{x}_{i}+u_{i}\left(\partial_{u} \mathbf{x}_{i}\right)+v_{i}\left(\partial_{v} \mathbf{x}_{i}\right), \quad \mathbf{n}_{i}\left(u_{i}, v_{i}\right)=\mathbf{n}_{i}+u_{i}\left(\partial_{u} \mathbf{n}_{i}\right)+v_{i}\left(\partial_{v} \mathbf{n}_{i}\right) \\
T\left(u_{2}, v_{2}\right)=\binom{\partial_{u} \mathbf{x}_{2}-\left\langle\partial_{u} \mathbf{x}_{2}, \mathbf{n}_{2}\left(u_{2}, v_{2}\right)\right\rangle \mathbf{n}_{2}\left(u_{2}, v_{2}\right)}{\partial_{v} \mathbf{x}_{2}-\left\langle\partial_{u} \mathbf{x}_{2}, \mathbf{n}_{2}\left(u_{2}, v_{2}\right)\right\rangle \mathbf{n}_{2}\left(u_{2}, v_{2}\right)}
\end{gathered}
$$

The rest is just a big nested application of the product rule. When differentiating $C$, we get two terms: one which accounts for changes of the tangent frame at $\mathbf{x}_{2}$, with $h$ held constant, and one which accounts for changes of $h$, while the tangent frame is held constant. For the former, we require derivatives of $T$, e.g:

$$
\frac{\partial T}{\partial u_{2}}=\binom{-\left\langle\partial_{u} \mathbf{x}_{2}, \partial_{u} \mathbf{n}_{2}\right\rangle \mathbf{n}_{2}-\left\langle\partial_{u} \mathbf{x}_{2}, \mathbf{n}_{2}\right\rangle \partial_{u} \mathbf{n}_{2}}{-\left\langle\partial_{v} \mathbf{x}_{2}, \partial_{u} \mathbf{n}_{2}\right\rangle \mathbf{n}_{2}-\left\langle\partial_{v} \mathbf{x}_{2}, \mathbf{n}_{2}\right\rangle \partial_{u} \mathbf{n}_{2}}, \quad \frac{\partial T}{\partial v_{2}}=\binom{-\left\langle\partial_{u} \mathbf{x}_{2}, \partial_{v} \mathbf{n}_{2}\right\rangle \mathbf{n}_{2}-\left\langle\partial_{u} \mathbf{x}_{2}, \mathbf{n}_{2}\right\rangle \partial_{v} \mathbf{n}_{2}}{-\left\langle\partial_{v} \mathbf{x}_{2}, \partial_{v} \mathbf{n}_{2}\right\rangle \mathbf{n}_{2}-\left\langle\partial_{v} \mathbf{x}_{2}, \mathbf{n}_{2}\right\rangle \partial_{v} \mathbf{n}_{2}}
$$

and for the second term we can repeatedly apply the following vector calculus identity:

$$
\frac{\partial}{\partial t} \frac{\mathbf{z}(t)}{\|\mathbf{z}(t)\|}=\frac{1}{\|\mathbf{z}(t)\|} \frac{\partial \mathbf{z}(t)}{\partial t}-\frac{\mathbf{z}(t)}{\|\mathbf{z}(t)\|^{3}}\left\langle\mathbf{z}(t), \frac{\partial \mathbf{z}(t)}{\partial t}\right\rangle
$$

where $\mathbf{z}$ is a vector-valued function that depends on $t$. The resulting expression is quite messy and thus we only provide numerical values here. In particular, $\nabla C$ is given by

$$
\nabla C=\left[\begin{array}{cccccc}
-\frac{1}{4} & 0 & -\frac{3}{2} & 0 & \frac{1}{4} & 0 \\
0 & \frac{1}{2} & 0 & -1 & 0 & \frac{1}{2}
\end{array}\right]
$$

and therefore

$$
T_{\mathcal{S}}(\overline{\mathbf{x}})=-A^{-1}\left[\begin{array}{ll}
B_{1} & B_{3}
\end{array}\right]=\left(\begin{array}{cccc}
-\frac{1}{6} & 0 & \frac{1}{6} & 0 \\
0 & \frac{1}{2} & 0 & \frac{1}{2}
\end{array}\right)
$$

Since the parameterizations are orthonormal, the geometric term between $\mathbf{x}_{1}$ and $\mathbf{x}_{3}$ is simply

$$
G\left(\mathbf{x}_{1} \leftrightarrow \mathbf{x}_{3}\right)=\left|P_{2} A^{-1} B_{3}\right| G\left(\mathbf{x}_{1} \leftrightarrow \mathbf{x}_{2}\right)=\left|\begin{array}{cc}
\frac{1}{6} & 0 \\
0 & \frac{1}{2}
\end{array}\right| G\left(\mathbf{x}_{1} \leftrightarrow \mathbf{x}_{2}\right)=\frac{1}{48}
$$

For a $\mathrm{C}++$ implementation that computes $\nabla C$, please refer to the appendix A 2 .

## 4 Walking on the specular manifold

Our rendering algorithms use MCMC to explore the manifold of specular paths, and for this they require some form of local parameterization. With the differential geometry of the specular manifold in hand, we are now able to develop this extremely useful building block. We propose an algorithm that moves one of the endpoints of a specular chain and takes all the intermediate vertices to a valid new configuration. Later, in Section 5, we show how to apply this algorithm to render images.

To simplify the discussion, we will focus on the case where the position of a vertex $\mathbf{x}_{n}$ of a specular chain $\mathbf{x}_{1}, \ldots \mathbf{x}_{n}$ is adjusted to a given new position $\mathbf{x}_{n}^{\prime}$, while $\mathbf{x}_{1}$ is held fixed. We shall also briefly introduce the assumption that $\mathbf{x}_{n}$ is located on planar surface of infinite extent.

Our manifold walking algorithm is based on two key insights:

1. The $A$ and $B$ matrices (Section 3.3) may be used to map an infinitesimal in-plane movement of $\mathbf{x}_{n}$ to displacements of the vertices $\mathbf{x}_{2}, \ldots, \mathbf{x}_{n-1}$. We can use these displacements to approximate a finite change to the path simply by adding an offset to each vertex, but this will move the path off the specular manifold.
2. Ray tracing provides a deterministic means of projecting an off-specular path back onto the space of valid configurations based on its first two vertices. Given $\mathbf{x}_{1}$ and $\mathbf{x}_{2}$, we can trace a sequence of rays $\mathbf{x}_{i} \rightarrow \mathbf{x}_{i+1}$, at each step performing a specular reflection or refraction, and this leads to corrected positions $\mathbf{x}_{2}^{+} \ldots \mathbf{x}_{n}^{+}$.

By combining 1. and 2., we obtain a predictor-corrector type algorithm (Figure 5) that performs a step according to a local linear model, followed by a projection that restores the specular configuration, resulting in a new path $\mathbf{x}_{1}, \mathbf{x}_{2}^{+}, \ldots, \mathbf{x}_{n}^{+}$, and repeats until convergence. As long as
the prediction step solves the linear model and moves in the tangent space to the manifold, this iteration behaves like Newton's method, exhibiting quadratic convergence near the solution. As with all Newton-like iterations, it is not guaranteed to converge when started far from the solution, since the linear model may not be accurate enough to make progress. But since the model is first-order accurate, the algorithm is guaranteed to make forward progress when the constraint function is differentiable and the partial steps are small enough. Our algorithm uses a simple heuristic to decrease the step size when progress is not made, then increase back to full steps to get quadratic convergence as it approaches the target configuration. This iteration is illustrated in Figure 5 and laid out in the following algorithm:

```
WalkManifold \(\left(\mathrm{x}_{1}, \ldots, \mathrm{x}_{n} \rightsquigarrow \mathrm{x}_{n}^{\prime}\right)\)
    Set \(i=0\) and \(\beta=1\)
    while \(\left\|\mathrm{x}_{n}-\mathrm{x}_{n}^{\prime}\right\|>\varepsilon L\)
    \(\mathbf{p}=\mathbf{x}_{2}-\beta T\left(\mathbf{x}_{2}\right) P_{2} A^{-1} B_{k} T\left(\mathbf{x}_{n}\right)^{T}\left(\mathbf{x}_{n}^{\prime}-\mathbf{x}_{n}\right)\)
    Propagate the ray \(\mathbf{x}_{1} \rightarrow \mathbf{p}\) through all specular
    interactions, producing \(\mathbf{x}_{2}^{+}, \ldots, \mathbf{x}_{n}^{+}\).
    if step 4 succeeded and \(\left\|\mathbf{x}_{n}^{+}-\mathbf{x}_{n}^{\prime}\right\|<\left\|\mathbf{x}_{n}-\mathbf{x}_{n}^{\prime}\right\|\)
        \(\mathrm{x}_{2}, \ldots, \mathrm{x}_{n}=\mathrm{x}_{2}^{+}, \ldots, \mathrm{x}_{n}^{+}\)
        \(\beta=\min \{1,2 \beta\}\)
    else
        \(\beta=\frac{1}{2} \beta\)
        Set \(i=i+1\), and fail if \(i>N\).
    return \(\mathbf{x}_{2}, \ldots, \mathbf{x}_{n-1}\)
```

Here, $i$ records the number of iterations until a specified maximum $N$ is reached, and $\varepsilon$ is a relative error threshold to a scene-scale length $L$ (we use $L=\max _{i}\left\|\mathbf{x}_{i}\right\|, N=20$ and $\varepsilon=10^{-7}$ ). The variable $\beta$ denotes a step size that is dynamically adjusted to ensure that steps reduce the target distance. The pseudocode assumes that the plane vertex $\mathbf{x}_{n}$ has an orthonormal parameterization.

To make this algorithm usable for general specular chains, we must remove the previous assumption that the endpoint $\mathbf{x}_{n}$ is located on a plane. However, the actual locations of $\mathbf{x}_{n}^{+}$along the way from $\mathbf{x}_{n}$ to $\mathbf{x}_{n}^{\prime}$ are not needed; all that matters is that $\mathbf{x}_{n}^{\prime}$ is on the surface at convergence. Therefore, we construct a plane containing $\mathbf{x}_{n}$ and $\mathbf{x}_{n}^{\prime}$, and in step 4 of WalkManifold the last propagation step computes an intersection against this hypothetical plane, ignoring the actual scene geometry. Once the algorithm converges, we must ensure that $\mathbf{x}_{n-1}$ and $\mathbf{x}_{n}$ are mutually visible before reporting the manifold walk as successful.

In our implementation, we assign the plane normal using the following symmetric orthogonalization procedure

$$
\mathbf{n}_{\text {plane }}:=\gamma\left(\gamma\left(\mathbf{n}+\mathbf{n}^{\prime}\right)-\overrightarrow{\mathbf{x}_{n} \mathbf{x}_{n}^{\prime}}\left(\gamma\left(\mathbf{n}+\mathbf{n}^{\prime}\right), \overrightarrow{\mathbf{x}_{n} \mathbf{x}_{n}^{\prime}}\right\rangle\right)
$$

where $\mathbf{n}$ and $\mathbf{n}^{\prime}$ are the surface normals at $\mathbf{x}_{n}$ and $\mathbf{x}_{n}^{\prime}$, respectively, and $\gamma(\mathbf{v}):=\mathbf{v} /\|\mathbf{v}\|$. This ensures that the same plane is used for walks $\mathbf{x}_{n} \rightsquigarrow \mathbf{x}_{n}^{\prime}$ and $\mathbf{x}_{n}^{\prime} \rightsquigarrow \mathbf{x}_{n}$.


Figure 5: One iteration of updating a path using the specular manifold. (a) The path vertices are modified according to a local linear model, which (b) corresponds to a step along the tangent plane to the manifold, then (a) a nearby valid specular path is found, which (b) corresponds to projecting back onto the manifold.

To compute the matrix $A$, we derived symbolic expressions for the manifold constraints in $C$ (Equation 8). A C++ implementation is given in the appendix. When solving the resulting linear system in step 3, it is beneficial to exploit the special structure of this matrix, which becomes important when processing specular chains with more than about ten vertices. We solve for $\mathbf{p}$ using a block tridiagonal LU factorization, and this reduces the time complexity from $O\left(n^{3}\right)$ to $O(n), n$ being the number of vertices in the chain.

There are several situations under which this algorithm may fail to converge: first, a specular path between $\mathbf{x}_{1}$ and $\mathbf{x}_{n}^{\prime}$ need not exist at all. Secondly, WalkManifold usually cannot find paths that lie on a different connected component of the manifold. Thirdly, when the local structure of the manifold is complex (e.g. due to high-frequency geometric detail of the reflectors and refractors) and $\left\|\mathrm{x}_{n}-\mathrm{x}_{n}^{\prime}\right\|$ is large, the iteration may not converge to a solution. Finally, the linear system may not be invertible, which happens when the last path lies at the fold of a wavefront, e.g. a caustic receiving an infinite power per unit area. During the $\sim 10^{12}$ manifold walks performed to produce the results of the paper, this last case only occurred $\sim 1.7 \cdot 10^{5}$ times, hence it does not appear to be an issue in practice. However, in the MCMC context it is not an problem for the iteration to fail occasionally, as explained in the next section.

The manifold walking algorithm works reliably for large chains with over 10 vertices, especially when it is used by the transition rule discussed in the next section, which only moves the endpoints of chains by a small amount. In our scenes, we observe between 92 and $98 \%$ successful walks, each taking 2-3 iterations on average to converge to the tolerance $\varepsilon=10^{-7}$. The failing $2-8 \%$ mainly contain cases where WalkManifold failed for good reasons, because it was asked to walk to a point for which there is no valid configuration on the manifold.

## 5 Manifold exploration for surfaces

In this section, we present a new transition rule that proposes steps in path space using manifold walks. In the context of MCMC, a transition rule, or perturbation, is a random process that generates a proposal state conditioned on the current state of a Markov chain. It provides the basic means of navigating through the state space, but to do this correctly the rule must satisfy two basic criteria: transitions must be reversible (i.e. return to the previous state with nonzero probability density), and the rule must also supply a function that, up to constant factors, computes the
probability of proposals conditioned on the current state.
To create an efficient sampling procedure, a transition rule should furthermore propose modifications of an appropriate scale. A rule that takes large steps will tend to leave local maxima of the target distribution $\pi$, and such steps are rejected with high probability. A rule that takes tiny steps will find most of them accepted, but it will not explore the state space well. Our perturbation is designed so that its scale naturally adapts to the scene, including the geometry and material properties.

The new perturbation supports general scenes and can be used both with the ERPT algorithm by Cline et al. and the path space MLT framework proposed by Veach and Guibas, where it replaces and generalizes the lens, caustic, and multi-chain perturbations. Depending on which combination is used, we call the resulting algorithm either Manifold Exploration Path Tracing (MEPT) or Manifold Exploration Metropolis Light Transport (MEMLT).

### 5.1 Manifold perturbation

Given an input path, the manifold perturbation finds a nearby path using a sequence of steps that can be grouped into sampling and connection phases (Figure 6). The sampling phase chooses a subpath to be modified that consists of three non-specular vertices which are potentially separated by specular chains. We shall denote these non-specular vertices as $\mathbf{x}_{a}, \mathbf{x}_{b}$, and $\mathbf{x}_{c}$. After establishing the type of perturbation to be performed, the sampling step generates a perturbed outgoing direction at the vertex $\mathbf{x}_{a}$ and propagates it through the specular chain between $\mathbf{x}_{a}$ and $\mathbf{x}_{b}$ (if any) until arriving at a new non-specular vertex $\mathbf{x}_{b}^{\prime}$ in the neighborhood of $\mathbf{x}_{b}$. If the path configuration (i.e. the arrangement of specular and nonspecular vertices) changed, the perturbation is rejected immediately.

Up to this point, the proposed scheme is similar to the set of perturbations proposed by Veach and Guibas. However, recall that in their work, perturbations must propagate through the path until arriving at a pair of adjacent non-specular vertices ("DD" in Heckbert's [1990] notation) that can be used to establish a connection edge. Any attempt to connect two sampled subpaths that involves a specular vertex must fail, since the probability of creating a valid path in this manner is zero.

In comparison, our perturbation can stop at vertex $\mathbf{x}_{b}$ and use the WALKMANIFOLD algorithm to update the configuration of the specular chain between $\mathbf{x}_{b}$ and $\mathbf{x}_{c}$ (Figure 7). This seemingly subtle difference has major repercussions on the types of scenes that can be rendered efficiently. In particular, the resulting method can systematically explore large classes of specular paths instead of having to rely on finding them by chance.


Figure 6: The manifold perturbation samples a perturbed outgoing direction from a vertex $\mathbf{x}_{a}$ and propagates it through a specular chain (if any) using ray tracing until arriving at a non-specular vertex $\mathbf{x}_{b}$. To connect the vertices $\mathbf{x}_{b}$ and $\mathbf{x}_{c}$, the perturbation performs a manifold walk to determine the positions of intermediate specular vertices (if any).


Figure 7: Manifold perturbation example: a slightly perturbed outgoing direction at $\mathbf{x}_{a}$ is propagated until encountering the non-specular vertex $\mathbf{x}_{b}^{\prime}$. Previously, it was not clear how to "connect" $\mathbf{x}_{b}^{\prime}$ to $\mathbf{x}_{c}$ through multiple specular interactions. Our method can find this connection given knowledge about the previous path.

In the following, we will discuss the perturbation in more detail; first for the ideally specular case and then in a more general form that extends to rough surfaces.

Strategy sampling: Motivated by the desire to attempt a large range of different types of path modifications with high probability, the sampling step first chooses among the possible perturbation strategies for a given path by selecting three vertices as follows. Given a path $\mathbf{x}_{1}, \ldots, \mathbf{x}_{k}$, uniformly select a non-specular initial vertex $\mathbf{x}_{a}$, as well as a perturbation direction (i.e. towards the light source or towards the camera). Walk along the path in this direction until the first non-specular vertex is encountered, and continue until a second non-specular vertex is found. This path traversal may fail by walking past the end of the path, in which case the strategy sampling phase is simply restarted from scratch. This determines $\mathbf{x}_{a}, \mathbf{x}_{b}$ and $\mathbf{x}_{c}$. For notational convenience, assume that $a<b<c$.

Perturbation sampling: With the overall strategy established, the sampling phase now perturbs the path segment $\mathbf{x}_{a+1}, \ldots, \mathbf{x}_{b}$. The goal here is to produce a new subpath $\mathbf{x}_{a+1}^{\prime}, \ldots, \mathbf{x}_{b}^{\prime}$ that is "nearby". When the vertex $\mathbf{x}_{a}$ denotes a surface scattering event with incident and exitant directions $\omega_{i}=\overrightarrow{\mathbf{x}_{a} \mathbf{x}_{a-1}}$ and $\omega_{o}=\overrightarrow{\mathbf{x}_{a} \mathbf{x}_{a+1}}$, the perturbation determines $\mathbf{x}_{a+1}^{\prime}$ by tracing a ray in a direction $\omega_{o}^{\prime}$ that is sampled from a suitable spherical distribution $D\left(\omega_{o}^{\prime}\right)$ concentrated around $\omega_{o}$. It is absolutely critical that this distribution generates direction changes of the appropriate scale: for instance, when $\mathbf{x}_{a}$ is a diffuse material, relative large perturbations are in order. On the other hand, when $\mathbf{x}_{a}$ is a glossy material that only reflects into a small cone of directions, large perturbations will almost always be rejected, reducing performance.

Observe that a useful hint about the right scale can be obtained directly from the scattering model at $\mathbf{x}_{a}$, in particular from the associated importance sampling density $p\left(\omega_{i} \rightarrow \omega_{o}\right)$. When this sampling density is high, $\omega_{o}$ is likely located on a sharp peak of the scattering function, and small steps are appropriate.

Our strategy is to sample from a distribution $D\left(\omega_{o}^{\prime}\right)$ centered at $\omega_{o}$ whose concentration is set so that $D\left(\omega_{o}\right)$ equals $\lambda^{2} p\left(\omega_{i} \rightarrow \omega_{o}\right)$. For $D\left(\omega_{o}^{\prime}\right)$, we use the spherical von Mises-Fisher distribution.


Figure 8: The effects of the $\lambda$ parameter. This scene was rendered with relatively short Markov Chains (100 steps), hence the range of "good" parameter values is lower than in our other examples scenes. Modeled after a scene by Eric Veach.

The parameter $\lambda$ (generally set between 50 and 500 ) specifies how large the perturbations are relative to standard BSDF sampling. This is the main parameter of our technique, and it affects how far perturbations will move in path space. When $\lambda$ is set to an inappropriately low or high value, the amount of noise present in the output renderings increases. In the first case, too few mutations are accepted, causing the chain to become "stuck" in certain paths for many iterations. In the latter case, the steps taken by the chain are too small to effectively explore path space, and this results in the typical coherent noise patterns that are known from other MLT-type algorithms. A comparison involving different settings is shown in Figure 8. We currently set this parameter manually to achieve a desired acceptance ratio, but this could in theory be automated using adaptive MCMC.

When $\mathbf{x}_{a}$ is a camera or light source, we choose a new outgoing direction in much the same way, but query the underlying model for the directional density of the associated sampling method (e.g. uniformly choosing pixels in screen-space). To further enlarge the space of possible perturbations, following Veach, we split camera and light source endpoints into two separate vertices corresponding to the position and direction components. When $\mathbf{x}_{a}$ is such a position vertex, we perturb its location on the aperture or light source by sampling a tangential displacement from a 2D normal distribution with variance $\rho /\left(2 \pi \lambda^{2}\right)$, where $\rho$ is the surface area.

After $\mathbf{x}_{a+1}^{\prime}$ has been determined in this manner, the perturbation is propagated through the specular chain until reaching $\mathbf{x}_{b}^{\prime}$. This process is deterministic.

Connection: When there is no specular chain between $\mathbf{x}_{b}^{\prime}$ and $\mathbf{x}_{c}$, the connection step only entails checking that the vertices are mutually visible, and that their scattering models carry illumination along the connection edge. When there is a manifold, we first set

$$
\mathbf{x}_{c-1}^{\prime}, \ldots \mathrm{x}_{b+1}^{\prime}=\text { WaLkMANIFOLD }\left(\mathbf{x}_{c}, \ldots \mathrm{x}_{b} \rightarrow \mathrm{x}_{b}^{\prime}\right)
$$

and then perform the same verification.
Recall that a key requirement of the Metropolis-Hastings algorithm discussed in Section 2.2 was that a nonzero transition probability $T\left(\mathbf{x}, \mathbf{x}^{\prime}\right)>0$ also implies that $T\left(\mathbf{x}^{\prime}, \mathbf{x}\right)>0$. This creates a potential issue when walking on the manifold, because WALKMANIFOLD can be non-reversible. It might succeed in moving from $\mathbf{x}$ to $\mathbf{x}^{\prime}$ but fail to move from $\mathbf{x}^{\prime}$ to $\mathbf{x}$. Even when the reverse iteration converges, the manifold can contain bifurcations so that it may converge to a different solution. Therefore, we always perform another manifold walk in the reverse direction and reject the perturbation if the path did not return to its original configuration. In the example scenes, we observed between $0 \%$ and $0.4 \%$ non-reversible walks.

Transition probability computation Finally, the change in the contribution function is computed and, together with the transition probabilities, used to randomly accept or reject the proposal with probability

$$
\begin{equation*}
r\left(\bar{x}, \bar{x}^{\prime}\right)=\min \left\{1, \frac{f_{j}\left(\overline{\mathbf{x}}^{\prime}\right) T\left(\overline{\mathbf{x}}^{\prime}, \overline{\mathbf{x}}^{\prime}\right)}{f_{j}(\overline{\mathbf{x}}) T\left(\overline{\mathbf{x}}, \overline{\mathbf{x}}^{\prime}\right)}\right\} \tag{10}
\end{equation*}
$$

where $f_{j}$ is the contribution function (Equation 4) and $\overline{\mathbf{x}}$ and $\overline{\mathbf{x}}^{\prime}$, denote the original and proposal path respectively. Note that many factors cancel in the above ratio, particularly all of those in $f_{j}$ that are associated with the unchanged path segment, or common terms in the transition
probability. For instance, the probability of choosing a particular sampling strategy cancels, since it only depends on the (unchanged) path configuration.

We require that $T\left(\overline{\mathbf{x}}, \overline{\mathbf{x}}^{\prime}\right)$ and $T\left(\overline{\mathbf{x}}^{\prime}, \overline{\mathbf{x}}\right)$ express the density of forward and reverse proposals in a common measure. Observe that sampling an outgoing direction $\omega_{o}^{\prime}$ from $D\left(\omega_{o}^{\prime}\right)$ at $\mathbf{x}_{a}$, and propagating it through the first specular chain, produces area density $D^{\perp}\left(\omega_{o}^{\prime}\right) G\left(\mathbf{x}_{a} \leftrightarrow \cdots \leftrightarrow \mathbf{x}_{b}^{\prime}\right)$ on the surface at $\mathbf{x}_{b}^{\prime}\left(\right.$ where $D^{\perp}\left(\omega_{o}^{\prime}\right)=D\left(\omega_{o}^{\prime}\right) /\left|\cos \left(\mathbf{n}_{a}, \omega_{o}\right)\right|$ denotes probability with respect to the projected solid angle measure at $\mathbf{x}_{a}$ ). This is the needed transition probability $T\left(\overline{\mathbf{x}}, \overline{\mathbf{x}}^{\prime}\right)$.

In a MCMC-based rendering system, we will generally want to sample paths based on their contribution to the entire image rather than to a single pixel. This is accomplished by replacing $W_{e}^{(j)}$ in (4) with an importance function that measures the overall luminance received on the image plane.

### 5.2 Extension to glossy materials

The method just presented can be used for scenes with both specular and non-specular transport, and the two classes are treated separately. This is unfortunate, since a near-specular chain through an almost-smooth dielectric object cannot be explored as effectively as a perfectly specular one. However, it turns out that a simple generalization suffices to encompass these materials as well.

The path space integrand corresponding to a chain of glossy interactions has its energy concentrated in a thin "band" near the specular manifold, and the key idea of how we handle glossy materials is to take steps along a family of offset manifolds that are parallel to the specular manifold, so that path space near the specular manifold can be explored without stepping out of this thin band of near-specular transport. In this section, we add a simple extension that endows the perturbation with the ability to walk on offset manifolds and to recognize when this is appropriate.

For this, we first replace Equation (8) with the offset manifold

$$
\begin{equation*}
\mathcal{S}_{\mathbf{o}}=\{\overline{\mathbf{x}} \mid C(\overline{\mathbf{x}})=\mathbf{o}\}, \tag{11}
\end{equation*}
$$

where o captures the offset from ideal specular transport. Inspecting $C$ (Equation 6) yields an intuitive explanation for the contents of the vector $\mathbf{o}$ in terms of microfacet theory. The two entries associated with each vertex $\mathbf{x}_{i}$ record the microsurface normal $\mathbf{m}_{i}$ responsible for the reflection or refraction (which is now different from the shading normal $\mathbf{n}_{i}$ ), projected into the shading tangent space $T\left(\mathbf{x}_{i}\right)$. Our extended perturbation then preserves the projection of this microgeometry normal $\mathbf{m}_{i}$ as an invariant during the manifold traversal.

Changes to the manifold perturbation: Since the differential geometry of the offset manifold is identical to that of the ordinary manifold, the only required change in WalkManifold affects the ray tracing step on line four, where the algorithm now reflects and refracts using $\mathbf{m}_{i}$ instead of $\mathbf{n}_{i}$. Similarly, the deterministic phase of the manifold perturbation responsible for propagating the sampled direction at $\mathbf{x}_{a}$ to a position $\mathbf{x}_{b}^{\prime}$ uses these normals instead. Note that it is straightforward to handle both cases, near-specular and specular perturbations, using the same implementation.

Recognizing near-specular transport: An important aspect about the treatment of general scattering is the decision of whether the surface associated with a scattering event is "smooth


Figure 9: Perturbation of path with near-specular surface interactions: instead of requiring that the half-vectors agree with the surface normals, their direction is preserved in the surface frame.
enough" to be classified as part of a specular chain. We make this decision randomly by assigning a specular probability $\psi\left(\mathbf{x}_{i}\right)$ to each vertex that takes on values 0 and 1 when $\mathbf{x}_{i}$ is diffuse and specular, respectively, and values in $(0,1)$ when $\mathbf{x}_{i}$ is at a rough interface. This avoids the issues of "hard" classifications that are commonly used in rendering algorithms. For specifics on our specular probability function $\psi$, please refer to the appendix.

Transition probability: In the purely specular case discussed earlier, the proposal and target distributions were both supported on the same space $\mathcal{S}$. This property permitted computing transition probabilities under an arbitrary projection (e.g. onto the vertex $\mathbf{x}_{b}$ ), since the determinant of the associated change of variables canceled when considering ratios of the form $f_{j} / T$.

In the glossy case, this does not hold anymore: $T$ is a distribution on an offset manifold $\mathcal{S}_{\mathbf{o}}$, whereas $f_{j}$ is defined on the higher-dimensional space $\bigcup_{\mathbf{o}} \mathcal{S}_{\mathbf{0}}$. To compute transition probabilities, we must perform a change of variables to separate out these "perpendicular" dimensions in $f_{j}$,


Figure 10: Consistently classifying glossy materials as non-specular or specular produces unsatisfactory results. Instead, our method makes this decision randomly whenever encountering a rough object (modeled after a scene by Cline et al.)


Figure 11: An illustration of the matrices involved in the glossy transition probability computation.
and this causes a determinant to appear in the final transition probability:

$$
\begin{equation*}
\left.T\left(\overline{\mathbf{x}}, \overline{\mathbf{x}}^{\prime}\right)=T_{\text {spec }}\left(\overline{\mathbf{x}}, \overline{\mathbf{x}}^{\prime}\right)\left|\frac{\partial\left[\mathbf{x}_{b}, \mathbf{x}_{i_{1}}, \ldots, \mathbf{x}_{i_{k}}\right]}{\partial\left[\mathbf{x}_{b}, \mathbf{o}_{i_{1}}, \ldots, \mathbf{o}_{i_{k}}\right]}\right|\left(1-\psi\left(\mathbf{x}_{b}\right)\right)\left(1-\psi\left(\mathbf{x}_{c}\right)\right)\right) \prod_{i=a+1, i \neq b}^{c-1} \psi\left(\mathbf{x}_{i}\right) . \tag{12}
\end{equation*}
$$

Here, $T_{\text {spec }}$ is the transition probability of the purely specular case and the indices $i_{1}, \ldots, i_{k} \in$ $\{a+1, \ldots, c-1\} \backslash\{b\}$ refer to glossy vertices that were classified as specular. The terms involving $\psi$ represent the discrete probability of performing this classification; roughness values may change during a perturbation, and hence we must account for them to maintain detailed balance.

The determinant in (12) is not hard to compute. Recall that the $A$-matrix (Section 3.3) maps perturbations of the vertex positions to changes of the half-vectors. Here, we seek the opposite: how the non-specular vertices move as a function of $\mathbf{o}$ and $\mathbf{x}_{b}$. We therefore compute the matrix $A$ over the vertex range $\mathbf{x}_{a+1}, \ldots, \mathbf{x}_{c-1}$ with one small adjustment: the two rows associated with vertex $\mathbf{x}_{b}$ are set so that there is a $2 \times 2$ identity matrix on the diagonal and zeroes everywhere else. (Figure 11) We invert this matrix and discard all rows and columns in this inverse that are associated with specular vertices, and the determinant of the result is the sought change of variables factor. Note that when the two chains only consists of glossy vertices, it is equivalent (and considerably faster) to compute the inverse determinant of the block tridiagonal matrix.

## 6 A unified path space framework for surfaces and volumes

We have presented Manifold Exploration in the context of surfaces, but the extension to volumes poses no fundamental difficulties. Before we can proceed, however, we require a description of how volumetric interactions are modeled in the path-space setting.

The purpose of this section is to find an operator-based description of surface and volume light transport similar in style to the frameworks presented by Arvo [1995] and Veach [1997], which is later used to derive a path-space integral. Such a unified surface and volume path space has been proposed in prior work by Pauly et al. [2000], but the context of that work was computational, and hence it only definitions were supplied.

The goal here is to provide a complete derivation that provides solid theoretical underpinnings
for the development of new path-space rendering algorithms. We roughly follow the approach of Chapter 4 in Eric Veach's Ph.D. thesis [Veach 1997] and conclude with an extended set of operators that can describe interactions with surfaces and volumes.

### 6.0.1 Definitions

It will be convenient to establish some notation: as before, $\mathcal{M}$ denotes the set of surfaces. We define the distance to the next surface as seen by the ray $(\mathbf{x}, \omega) \subseteq \mathbb{R}^{3} \times S^{2}$ as

$$
d_{\mathcal{M}}(\mathbf{x}, \omega):=\inf \{d>0 \mid \mathbf{x}+d \omega \in \mathcal{M}\}
$$

where $\inf \emptyset=\infty$. Based on this distance, we can define a ray-casting function $\mathbf{x}_{\mathcal{M}}$ :

$$
\mathbf{x}_{\mathcal{M}}(\mathbf{x}, \omega)=\mathbf{x}+d_{\mathcal{M}}(\mathbf{x}, \omega) \omega
$$

We shall also make use of the following notation to express integration along rays

$$
\int_{\mathbf{x}}^{\mathbf{x}_{\mathcal{M}}(\mathbf{x}, \omega)} f(\mathbf{z}) \mathrm{d} \mathbf{z}:=\int_{0}^{d_{\mathcal{M}}(\mathbf{x}, \omega)} f(\mathbf{x}+t \omega) \mathrm{d} t .
$$

Note that that while $\mathbf{x}_{\mathcal{M}}(\mathbf{x}, \omega)$ is generally undefined for "points at infinity" $\left(d_{\mathcal{M}}(x, \omega)=\infty\right)$, it has a well-defined meaning when encountered in integrals of the above form. We will also write

$$
\int_{\mathbf{x}}^{\mathbf{y}} f(\mathbf{z}) \mathrm{d} \mathbf{z}:=\int_{0}^{\|\mathbf{x}-\mathbf{y}\|} f(\mathbf{x}+t \overrightarrow{\mathbf{x}} \overrightarrow{\mathbf{y}}) \mathrm{d} t, \quad \text { where } \quad \overrightarrow{\mathbf{x}} \mathbf{y}:=\frac{\mathbf{y}-\mathbf{x}}{\|\mathbf{y}-\mathbf{x}\|}
$$

### 6.1 Light transport model

The radiative transfer equation (RTE) is commonly used to describe the steady-state light transport on a three-dimensional domain $\mathcal{V} \subseteq \mathbb{R}^{3}$. It is given by the following energy balance equation:

$$
\begin{align*}
& (\omega \cdot \nabla) L(\mathbf{x}, \omega)+\sigma_{t}(\mathbf{x}) L(\mathbf{x}, \omega)= \\
& \sigma_{s}(\mathbf{x}) \int_{S^{2}} f_{p}\left(\mathbf{x}, \omega^{\prime} \rightarrow \omega\right) L\left(\mathbf{x}, \omega^{\prime}\right) \mathrm{d} \sigma_{\mathbf{x}}\left(\omega^{\prime}\right)+L_{e}(\mathbf{x}, \omega), \quad \mathbf{x} \in \mathcal{V}^{\circ} . \tag{13}
\end{align*}
$$

where $\sigma_{s}$ and $\sigma_{t}$ are the scattering and extinction coefficients, $f_{p}$ is the phase function, $L$ denotes the volume radiance, and $L_{e}$ represents emitted radiance. $d \sigma_{\mathbf{x}}$ denotes differential solid angles at $\mathbf{x}$. Please note that due to the volumetric setting, the emission term $L_{e}$ has units of radiance per unit length at interior points $x \in \mathcal{V}^{\circ}$.

We will later specify boundary conditions on the set of surfaces $\mathcal{M} \subseteq \mathcal{V}$, hence the above equation only holds on the interior of the domain, i.e. $\mathcal{V}^{\circ}:=\mathcal{V} \backslash \mathcal{M}$.

Restricted to a line of direction $\omega$ parameterized by $s \in \mathbb{R}$, Equation (13) can be rewritten as a one-dimensional ordinary differential equation of the form

$$
L^{\prime}(s)+\sigma_{t}(s) L(s)=Z(s)
$$

where $Z(s)$ represents the emitted and in-scattered radiance at $s$. Given an initial condition at
$s=0$, this ODE has the solution

$$
\begin{align*}
L(s) & =\exp \left(-\int_{0}^{s} \sigma_{t}(r) \mathrm{d} r\right)\left[L(0)+\int_{0}^{s} \exp \left(\int_{0}^{t} \sigma_{t}(r) \mathrm{d} r\right) Z(t) \mathrm{d} t\right] \\
& =\exp \left(-\int_{0}^{s} \sigma_{t}(r) \mathrm{d} r\right) L(0)+\int_{0}^{s} \exp \left(-\int_{t}^{s} \sigma_{t}(r) \mathrm{d} r\right) Z(t) \mathrm{d} t \tag{14}
\end{align*}
$$

where the first term corresponds to attenuated radiance from surfaces and the second term accounts for in-scattered radiance.

In this solution, the parameterization and boundary condition at $s=0$ were arbitrarily chosen. In general, the boundary condition will be specified by the nearest surface visible along the ray $(\mathbf{x},-\omega)$. To formalize this constraint, we will first define the reduced surface radiance at $(\mathbf{x}, \omega)$ as

$$
L_{\mathrm{red}}(\mathbf{x}, \omega):= \begin{cases}L_{o}\left(\mathbf{x}_{\mathcal{M}}(\mathbf{x},-\omega), \omega\right) \tau\left(\mathbf{x}_{\mathcal{M}}(\mathbf{x},-\omega) \leftrightarrow \mathbf{x}\right), & d_{\mathcal{M}}(\mathbf{x},-\omega)<\infty \\ 0, & \text { otherwise }\end{cases}
$$

where

$$
\tau(\mathbf{x} \leftrightarrow \mathbf{y}):=\exp \left(-\int_{\mathbf{x}}^{\mathbf{y}} \sigma_{t}(\mathbf{z}) \mathrm{d} \mathbf{z}\right)
$$

and $L_{o}$ represents the surface radiance scattered into direction $\omega$ (we will consider this term in more detail later). Now, we can substitute a suitable parameterization into Equation (14) and replace the first term by the reduced surface radiance, which yields the integral form of the radiative transfer equation:

$$
\begin{aligned}
& L(\mathbf{x}, \omega)=L_{\mathrm{red}}(\mathbf{x}, \omega)+\int_{0}^{d_{\mathcal{M}}(\mathbf{x},-\omega)} \tau(\mathbf{x} \leftrightarrow \mathbf{x}-t \omega) \\
& \\
& \quad\left(\sigma_{s}(\mathbf{x}-t \omega) \int_{S^{2}} f_{p}\left(\mathbf{x}-t \omega, \omega^{\prime} \rightarrow \omega\right) L\left(\mathbf{x}-t \omega, \omega^{\prime}\right) \mathrm{d} \sigma_{\mathbf{x}-t \omega}\left(\omega^{\prime}\right)+L_{e}(\mathbf{x}-t \omega, \omega)\right) \mathrm{d} t
\end{aligned}
$$

which can be written more compactly using the introduced notation:

$$
\begin{equation*}
L(\mathbf{x}, \omega)=L_{\mathrm{red}}(\mathbf{x}, \omega)+\int_{\mathbf{x}}^{\mathbf{x}_{\mathcal{M}}(\mathbf{x},-\omega)} \tau(\mathbf{x} \leftrightarrow \mathbf{y})\left(\sigma_{s}(\mathbf{y}) \int_{S^{2}} f_{p}\left(\mathbf{y}, \omega^{\prime} \rightarrow \omega\right) L\left(\mathbf{y}, \omega^{\prime}\right) \mathrm{d} \sigma_{\mathbf{y}}\left(\omega^{\prime}\right)+L_{e}(\mathbf{y}, \omega)\right) \mathrm{d} \mathbf{y} \tag{15}
\end{equation*}
$$

### 6.2 Boundary conditions

By specifying boundary conditions on the surfaces $\mathcal{M}$, we will invariably also introduce discontinuities in the volumetric radiance function $L$. These discontinuities are problematic when referring to the surface radiance field, since it will generally not be identical on both sides.

Analogous to a one-sided limit in the 1D case, we shall therefore distinguish between the frontand back-facing components $L^{+}(\mathbf{x}, \omega)$ and $L^{-}(\mathbf{x}, \omega)$ as determined by the surface normal $N(\mathbf{x})$ for $\mathbf{x} \in \mathcal{M}$. Based on this separation, the more intuitive incident and exitant radiance functions
can then be defined as

$$
L_{i}(\mathbf{x}, \omega):=\left\{\begin{array}{ll}
L^{+}(\mathbf{x},-\omega), & \omega \cdot N(\mathbf{x})>0 \\
L^{-}(\mathbf{x},-\omega), & \omega \cdot N(\mathbf{x})<0
\end{array} \quad \text { and } \quad L_{o}(\mathbf{x}, \omega):= \begin{cases}L^{+}(\mathbf{x}, \omega), & \omega \cdot N(\mathbf{x})>0 \\
L^{-}(\mathbf{x}, \omega), & \omega \cdot N(\mathbf{x})<0\end{cases}\right.
$$

With the help of these definitions, we can introduce the usual boundary condition known as the rendering equation [Kajiya 1986]:

$$
\begin{equation*}
L_{o}(\mathbf{x}, \omega)=\int_{S^{2}} f_{s}\left(\mathbf{x}, \omega^{\prime} \rightarrow \omega\right) L_{i}\left(\mathbf{x}, \omega^{\prime}\right) \mathrm{d} \sigma_{\mathbf{x}}^{\perp}\left(\omega^{\prime}\right)+L_{e}(\mathbf{x}, \omega), \quad \mathbf{x} \in \mathcal{M} \tag{16}
\end{equation*}
$$

where $f_{s}$ is the BSDF of the surface.
Away from surfaces, the distinction between $L_{i}$ and $L_{o}$ is technically not necessary (involving only a change of direction) but we introduce it here for consistency. Specifically, we will set $L_{i}(\mathbf{x}, \omega):=L(\mathbf{x},-\omega)$ and $L_{o}(\mathbf{x}, \omega):=L(\mathbf{x}, \omega)$ for $\mathbf{x} \in \mathcal{V}^{\circ}$ and to analogously express the integral form of the RTE (15) in terms of $L_{i}$ and $L_{o}$ :

$$
\begin{align*}
& L_{o}(\mathbf{x}, \omega)=L_{\mathrm{red}}(\mathbf{x}, \omega)+ \\
& \qquad \int_{\mathbf{x}}^{\mathbf{x} \mathcal{M}(\mathbf{x},-\omega)} \tau(\mathbf{x} \leftrightarrow \mathbf{y})\left(\sigma_{s}(\mathbf{y}) \int_{S^{2}} f_{p}\left(\mathbf{y},-\omega^{\prime} \rightarrow \omega\right) L_{i}\left(\mathbf{y}, \omega^{\prime}\right) \mathrm{d} \sigma_{\mathbf{y}}+L_{e}(\mathbf{y}, \omega)\right) \mathrm{d} \mathbf{y} . \tag{17}
\end{align*}
$$

Note the negated $\omega^{\prime}$ term in the spherical integral, which has a positive correspondence in the surface boundary condition (16). This apparent discrepancy is due to the the different parameter conventions of surface and volume scattering models. ${ }^{3}$

### 6.3 Operator notation

Volumetric light transport can be seen as the alternation of two steps: scattering on a surface or within the medium, followed by propagation and attenuation. Analogous to [Veach 1997] and [Arvo 1995], these two steps can be formulated as linear operators defined on the space of radiance functions. The goal of this section is to partition Equations (16) and (17) so that they can be expressed in this manner.

We define the scattering operator $\mathbf{K}$ as one of the in-scattering integrals in Equations (16) or (17) dependent on whether or not $\mathbf{x} \in \mathcal{M}$ :

$$
(\mathbf{K} h)(\mathbf{x}, \omega):= \begin{cases}\int_{S^{2}} f_{s}\left(\mathbf{x}, \omega^{\prime} \rightarrow \omega\right) h\left(\mathbf{x}, \omega^{\prime}\right) \mathrm{d} \sigma_{\mathbf{x}}^{\perp}\left(\omega^{\prime}\right), & \mathbf{x} \in \mathcal{M}  \tag{18}\\ \sigma_{s}(\mathbf{x}) \int_{S^{2}} f_{p}\left(\mathbf{x},-\omega^{\prime} \rightarrow \omega\right) h\left(\mathbf{x}, \omega^{\prime}\right) \mathrm{d} \sigma_{\mathbf{x}}\left(\omega^{\prime}\right), & \text { otherwise }\end{cases}
$$

On surfaces, $\mathbf{K}$ turns incident radiance into outgoing radiance. On the interior of the domain, it turns incident radiance into outgoing radiance per unit length.

In comparison, the transport operator $\mathbf{G}$ has a single definition on the whole domain. Its role is to transform outgoing radiance per unit length from the volume and outgoing radiance from

[^1]surfaces into incident radiance.
\[

$$
\begin{equation*}
(\mathbf{G} h)(\mathbf{x}, \omega):=\int_{\mathbf{x}}^{\mathbf{x}_{\mathcal{M}}(\mathbf{x}, \omega)} \tau(\mathbf{x} \leftrightarrow \mathbf{y}) h(\mathbf{y},-\omega) \mathrm{d} \mathbf{y}+\tau\left(\mathbf{x} \leftrightarrow \mathbf{x}_{\mathcal{M}}(\mathbf{x}, \omega)\right) h\left(\mathbf{x}_{\mathcal{M}}(\mathbf{x}, \omega),-\omega\right) \tag{19}
\end{equation*}
$$

\]

This leads to interesting differences in comparison to previous work: to match equations (16) and (17), we arrive at an equilibrium equation of the form

$$
\begin{equation*}
L_{i}=\mathbf{G}\left(\mathbf{K} L_{i}+L_{e}\right) \tag{20}
\end{equation*}
$$

whereas $L_{o}=L_{e}+\mathbf{K G} L_{o}$ was used by [Veach 1997]. The differences arise since volume light transport forces us to deal with both radiance and radiance per unit length. To express the equilibrium condition purely in terms of radiance, the order of the $\mathbf{G}$ and $\mathbf{K}$ operators must thus be reversed. Assuming invertibility, the solution operator $\mathbf{S}$ can now be found:

$$
\begin{aligned}
& L_{i}=\mathbf{G}\left(\mathbf{K} L_{i}+L_{e}\right) \\
\Leftrightarrow & (\mathbf{I}-\mathbf{G K}) L_{i}=\mathbf{G} L_{e} \\
\Leftrightarrow & L_{i}=\underbrace{(\mathbf{I}-\mathbf{G K})^{-1} \mathbf{G}}_{=: \mathbf{S}} L_{e} .
\end{aligned}
$$

When the following Neumann series converges in operator norm, the solution operator can also be expressed as

$$
\begin{equation*}
\mathbf{S}=\sum_{k=0}^{\infty}(\mathbf{G K})^{k} \mathbf{G} . \tag{21}
\end{equation*}
$$

### 6.4 Unified path integral formulation

Having defined the necessary operators, we can now proceed to find the desired path integral formulation. Note that in the following, we will often switch between parameterizing functions in terms of directions and positions. The arrow direction indicates the flow of light, which leads to the following conventions:

$$
\begin{aligned}
L_{e}(\mathbf{x} \rightarrow \mathbf{y}) & :=L_{e}(\mathbf{x}, \overrightarrow{\mathbf{x y}}) \\
W_{e}(\mathbf{x} \leftarrow \mathbf{y}) & :=W_{e}(\mathbf{x}, \overrightarrow{\mathbf{x y}}) \\
f_{s}(\mathbf{x} \rightarrow \mathbf{y} \rightarrow \mathbf{z}) & :=f_{s}(\mathbf{y}, \overrightarrow{\mathbf{y x}}, \overrightarrow{\mathbf{y z}}) \\
f_{p}(\mathbf{x} \rightarrow \mathbf{y} \rightarrow \mathbf{z}) & :=f_{p}(\mathbf{y}, \overrightarrow{\mathbf{x y}}, \overrightarrow{\mathbf{y z}})
\end{aligned}
$$

Here, $W_{e}$ denotes emitted importance. When dealing with rays that do not intersect any surfaces, we consider the ray-casting operator $\mathbf{x}_{\mathcal{M}}(\mathbf{x}, \omega)$ to return special 'points at infinity'. This ensures that the following is well-defined even if $d_{\mathcal{M}}(\mathbf{x}, \omega)=\infty$ :

$$
L_{e}\left(\mathbf{x} \rightarrow \mathbf{x}_{\mathcal{M}}(\mathbf{x}, \omega)\right)=L_{e}(\mathbf{x}, \omega) .
$$

### 6.4.1 Change of variables

When dealing with light transport on surfaces, it is often convenient to switch between integration over a sphere and integration over all surfaces of the scene, e.g.

$$
\int_{S^{2}} f\left(\mathbf{x}_{\mathcal{M}}(\mathbf{x}, \omega)\right) \mathrm{d} \sigma_{\mathbf{x}}^{\perp}(\omega)=\int_{\mathcal{M}} f(\mathbf{y}) G_{\text {surf }}(\mathbf{x} \leftrightarrow \mathbf{y}) \mathrm{d} A(\mathbf{y}) .
$$

This change of variables involves the geometric term

$$
G_{\text {surf }}(\mathbf{x} \leftrightarrow \mathbf{y})=V(\mathbf{x} \leftrightarrow \mathbf{y}) \cdot \frac{\left|\cos \theta_{1} \cos \theta_{2}\right|}{\|\mathbf{x}-\mathbf{y}\|^{2}}
$$

where $\theta_{1}$ and $\theta_{2}$ denote the angles that $\overrightarrow{\mathbf{x y}}$ makes with the surface normals $N(\mathbf{x})$ and $N(\mathbf{y})$, respectively, and $V$ is a visibility function defined as

$$
V(\mathbf{x} \leftrightarrow \mathbf{y}):= \begin{cases}1, & \text { if }\{\alpha \mathbf{x}+(1-\alpha) \mathbf{y} \mid \alpha \in(0,1)\} \cap \mathcal{M}=\emptyset \\ 0, & \text { otherwise }\end{cases}
$$

In the volume setting, a similar change of variables is possible: we consider a volume integral over rays radially emanating from a point $\mathbf{x}$. Fubini's theorem and some manipulation leads to

$$
\begin{aligned}
& \int_{S^{2}} \int_{\mathbf{x}}^{\mathbf{x} \mathcal{M}(\mathbf{x}, \omega)} f(\mathbf{y}) \mathrm{d} \mathbf{y} \mathrm{~d} \omega \\
= & \int_{S^{2}} \int_{0}^{\infty} f(\mathbf{x}+r \omega) V(\mathbf{x} \leftrightarrow \mathbf{x}+r \omega) \mathrm{d} r \mathrm{~d} \omega \\
= & \int_{0}^{\infty} \frac{1}{r^{2}} \int_{S^{2}(\mathbf{x}, r)} f(\mathbf{y}) V(\mathbf{x} \leftrightarrow \mathbf{y}) \mathrm{d} \mathbf{y} \mathrm{~d} t \\
= & \int_{\mathbb{R}^{3}} \frac{V(\mathbf{x} \leftrightarrow \mathbf{y})}{\|\mathbf{x}-\mathbf{y}\|^{2}} f(\mathbf{y}) \mathrm{d} \mathbf{y} .
\end{aligned}
$$

To handle all possible combinations of volume and surface endpoints, we redefine the basic geometry term as as follows:

$$
G(\mathbf{x} \leftrightarrow \mathbf{y}):=V(\mathbf{x} \leftrightarrow \mathbf{y}) \cdot \frac{D_{\mathbf{x}}(\overrightarrow{\mathrm{x}} \mathbf{y}) D_{\mathbf{y}}(\overrightarrow{\mathrm{yx}})}{\|\mathrm{x}-\mathbf{y}\|^{2}}
$$

where

$$
D_{\mathbf{a}}(\omega):= \begin{cases}|N(\mathbf{a}) \cdot \omega|, & \mathbf{a} \in \mathcal{M} \\ 1, & \text { otherwise }\end{cases}
$$

### 6.4.2 Path space measurement integral

Based on the generalized operators $\mathbf{G}$ and $\mathbf{K}$, we can proceed to find a path integral formulation of light transport similar to what is done in Chapter 8 of Eric Veach's thesis [Veach 1997].

We begin by considering a measurement $I_{j}$, which is defined as the ray-space convolution of the
importance $W_{e}^{(j)}$ emitted by sensor $j$ with the incident radiance $L_{i}$ (see [Veach 1997] for details):

$$
\begin{align*}
I_{j} & =\left\langle W_{e}^{(j)}, L_{i}\right\rangle \\
& =\int_{\mathcal{M}} \int_{S^{2}} W_{e}^{(j)}(\mathbf{x}, \omega) L_{i}(\mathbf{x}, \omega) \mathrm{d} \sigma_{\mathbf{x}}^{\perp}(\omega) \mathrm{d} A(\mathbf{x}) \tag{22}
\end{align*}
$$

For simplicity, we restrict the ray origins to the boundary, hence volumetric sensors are not supported. The goal of the path space approach is to express the measurement integral (22) over the set of transport paths $\mathcal{P}$

$$
I_{j}=\int_{\mathcal{P}} f_{j}(\bar{x}) \mathrm{d} \mu(\bar{x}),
$$

where $\mu$ is a measure on $\mathcal{P}$, and $f_{j}$ is a contribution weighting function specific to the measurement.
To render the definitions precise, let us define the path space $\mathcal{P}$ as the union of all fixed-length paths formed by concatenating vertices from $\mathcal{M}$ and $\mathcal{V}^{\circ}$ according to a configuration vector $\mathbf{c} \in\{0,1\}^{k}$, i.e.

$$
\mathcal{P}:=\bigcup_{k=1}^{\infty} \bigcup_{\mathbf{c} \in\{0,1\}^{k}} \mathcal{P}_{k}^{\mathrm{c}}
$$

where

$$
\mathcal{P}_{k}^{\mathrm{c}}:=\left\{\mathbf{x}_{0} \mathbf{x}_{1} \cdots \mathbf{x}_{k} \mid\left(\mathbf{x}_{i} \in \mathcal{M} \text { if } \mathbf{c}_{i}=0\right) \text { and }\left(\mathbf{x}_{i} \in \mathcal{V}^{\circ} \text { if } \mathbf{c}_{i}=1\right)(i=1, \ldots, k)\right\} .
$$

Using the Lebesgue measures for area and volume on $\mathcal{M}$ and $\mathcal{V}^{\circ}$, we can define a combined product measure on $\mathcal{P}$ :

$$
\mu(D):=\sum_{k=1}^{\infty} \sum_{\mathbf{c} \in\{0,1\}^{k}} \mu_{k}^{\mathbf{c}}\left(D \cap \mathcal{P}_{k}^{\mathbf{c}}\right) \quad \text { where } \quad \mu_{k}^{\mathbf{c}}(D):=\int_{D} \prod_{i=1}^{k} \begin{cases}\mathrm{~d} A\left(\mathbf{x}_{i}\right), & \mathbf{c}_{i}=0 \\ \mathrm{~d} V\left(\mathbf{x}_{i}\right), & \mathbf{c}_{i}=1\end{cases}
$$

To find the path space formulation, let us insert the operator form of the equilibrium equation (20) into the measurement integral (22):

$$
I_{j}=\int_{\mathcal{M}} \int_{S^{2}} W_{e}^{(j)}(\mathbf{x}, \omega)\left(\mathbf{G}\left(\mathbf{K} L_{i}+L_{e}\right)\right)(\mathbf{x}, \omega) \mathrm{d} \sigma_{\mathbf{x}}^{\perp}(\omega) \mathrm{d} A(\mathbf{x}) .
$$

Expanding the G operator leads to

$$
\begin{aligned}
=\int_{\mathcal{M}} \int_{S^{2}} W_{e}^{(j)}(\mathbf{x}, \omega) & {\left[\int_{\mathbf{x}}^{\mathbf{x} \mathcal{M}(\mathbf{x}, \omega)} \tau(\mathbf{x} \leftrightarrow \mathbf{y})\left[\mathbf{K} L_{i}+L_{e}\right](\mathbf{y},-\omega) \mathrm{d} \mathbf{y}\right.} \\
& \left.+\tau\left(\mathbf{x} \leftrightarrow \mathbf{x}_{\mathcal{M}}(\mathbf{x}, \omega)\right)\left[\mathbf{K} L_{i}+L_{e}\right]\left(\mathbf{x}_{\mathcal{M}}(\mathbf{x}, \omega),-\omega\right)\right] \mathrm{d} \sigma_{\mathbf{x}}^{\perp}(\omega) \mathrm{d} A(\mathbf{x}) .
\end{aligned}
$$

Using a change of variables (c.f. Section 6.4.1), we can turn the two summands into a volume and a surface integral, respectively:

$$
\begin{aligned}
=\int_{\mathcal{M}}[ & \int_{\mathcal{V}^{\circ}} W_{e}^{(j)}(\mathbf{x} \leftarrow \mathbf{y}) \tau(\mathbf{x} \leftrightarrow \mathbf{y}) G(\mathbf{x} \leftrightarrow \mathbf{y})\left[\mathbf{K} L_{i}+L_{e}\right](\mathbf{y} \rightarrow \mathbf{x}) \mathrm{d} V(\mathbf{y}) \\
& \left.\quad+\int_{\mathcal{M}} W_{e}^{(j)}(\mathbf{x} \leftarrow \mathbf{y}) \tau(\mathbf{x} \leftrightarrow \mathbf{y}) G(\mathbf{x} \leftrightarrow \mathbf{y})\left[\mathbf{K} L_{i}+L_{e}\right](\mathbf{y} \rightarrow \mathbf{x}) \mathrm{d} A(\mathbf{y})\right] \mathrm{d} A(\mathbf{x}) .
\end{aligned}
$$

Using the fact that $\mathbf{K} L_{i}=\mathbf{K} \mathbf{S} L_{e}=\sum_{k=1}^{\infty}(\mathbf{K G})^{k} L_{e}$, this can be rewritten as

$$
\begin{align*}
=\sum_{k=0}^{\infty} \int_{\mathcal{M}} & {\left[\int_{\mathcal{V}^{\circ}} W_{e}^{(j)}(\mathbf{x} \leftarrow \mathbf{y}) \bar{G}(\mathbf{x} \leftrightarrow \mathbf{y})\left[(\mathbf{K G})^{k} L_{e}\right](\mathbf{y} \rightarrow \mathbf{x}) \mathrm{d} V(\mathbf{y})\right.} \\
& \left.+\int_{\mathcal{M}} W_{e}^{(j)}(\mathbf{x} \leftarrow \mathbf{y}) \bar{G}(\mathbf{x} \leftrightarrow \mathbf{y})\left[(\mathbf{K G})^{k} L_{e}\right](\mathbf{y} \rightarrow \mathbf{x}) \mathrm{d} A(\mathbf{y})\right] \mathrm{d} A(\mathbf{x}) . \tag{23}
\end{align*}
$$

where we have defined the attenuated geometric term $\bar{G}(\mathbf{x} \leftrightarrow \mathbf{y}):=G(\mathbf{x} \leftrightarrow \mathbf{y}) \tau(\mathbf{x} \leftrightarrow \mathbf{y})$.
In the above geometric sum, the operators $\mathbf{G}$ and $\mathbf{K}$ are reversed in comparison previous encounters (e.g. Equation (21)). This makes it possible to perform additional simplifications: consider the concatenated operator $\mathbf{K G}$ for $\mathbf{y} \in \mathcal{M}$ :

$$
\begin{aligned}
(\mathbf{K G} h)(\mathbf{y} \rightarrow \mathbf{x})=\int_{S^{2}} f_{s}\left(\mathbf{x}_{\mathcal{M}}\left(\mathbf{y}, \omega^{\prime}\right) \rightarrow \mathbf{y} \rightarrow \mathbf{x}\right) & {\left[\int_{\mathbf{y}}^{\mathbf{x} \mathcal{M}\left(\mathbf{y}, \omega^{\prime}\right)} \tau(\mathbf{y} \leftrightarrow \mathbf{z}) h\left(\mathbf{z},-\omega^{\prime}\right) \mathrm{d} \mathbf{z}\right.} \\
& \left.+\tau\left(\mathbf{y} \leftrightarrow \mathbf{x}_{\mathcal{M}}\left(\mathbf{y}, \omega^{\prime}\right)\right) h\left(\mathbf{x}_{\mathcal{M}}\left(\mathbf{y}, \omega^{\prime}\right),-\omega^{\prime}\right)\right] \mathrm{d} \sigma_{\mathbf{y}}^{\perp}\left(\omega^{\prime}\right) .
\end{aligned}
$$

As before, we can perform a change of variables to volume and surface integrals:

$$
\begin{aligned}
= & \int_{\mathcal{V}^{\circ}} f_{s}(\mathbf{z} \rightarrow \mathbf{y} \rightarrow \mathbf{x}) \bar{G}(\mathbf{y} \leftrightarrow \mathbf{z}) h(\mathbf{z} \rightarrow \mathbf{y}) \mathrm{d} V(\mathbf{z}) \\
& +\int_{\mathcal{M}} f_{s}(\mathbf{z} \rightarrow \mathbf{y} \rightarrow \mathbf{x}) \bar{G}(\mathbf{y} \leftrightarrow \mathbf{z}) h(\mathbf{z} \rightarrow \mathbf{y}) \mathrm{d} A(\mathbf{z})
\end{aligned}
$$

For $\mathbf{y} \in \mathcal{V}^{\circ}$, we get

$$
\begin{aligned}
(\mathbf{K G} h)(\mathbf{y} \rightarrow \mathbf{x})= & \sigma_{s}\left[\int_{\mathcal{V}^{\circ}} f_{p}(\mathbf{z} \rightarrow \mathbf{y} \rightarrow \mathbf{x}) \bar{G}(\mathbf{y} \leftrightarrow \mathbf{z}) h(\mathbf{z} \rightarrow \mathbf{y}) \mathrm{d} V(\mathbf{z})\right. \\
& \left.+\int_{\mathcal{M}} f_{p}(\mathbf{z} \rightarrow \mathbf{y} \rightarrow \mathbf{x}) \bar{G}(\mathbf{y} \leftrightarrow \mathbf{z}) h(\mathbf{z} \rightarrow \mathbf{y}) \mathrm{d} A(\mathbf{z})\right]
\end{aligned}
$$

To unify the two cases, we define the generalized scattering function $\bar{f}$ as

$$
\bar{f}(\mathbf{z} \rightarrow \mathbf{y} \rightarrow \mathbf{x}):= \begin{cases}\sigma_{s} f_{p}(\mathbf{z} \rightarrow \mathbf{y} \rightarrow \mathbf{x}), & \mathbf{y} \in \mathcal{V}^{\circ}  \tag{24}\\ f_{s}(\mathbf{z} \rightarrow \mathbf{y} \rightarrow \mathbf{x}), & \mathbf{y} \in \mathcal{M}\end{cases}
$$

which leads to a single definition on the whole domain:

$$
\begin{aligned}
(\mathbf{K G} h)(\mathbf{y} \rightarrow \mathbf{x})= & \int_{\mathcal{V}^{\circ}} \bar{f}(\mathbf{z} \rightarrow \mathbf{y} \rightarrow \mathbf{x}) \bar{G}(\mathbf{y} \leftrightarrow \mathbf{z}) h(\mathbf{z} \rightarrow \mathbf{y}) \mathrm{d} V(\mathbf{z}) \\
& +\int_{\mathcal{M}} \bar{f}(\mathbf{z} \rightarrow \mathbf{y} \rightarrow \mathbf{x}) \bar{G}(\mathbf{y} \leftrightarrow \mathbf{z}) h(\mathbf{z} \rightarrow \mathbf{y}) \mathrm{d} A(\mathbf{z})
\end{aligned}
$$

Inserting increasing powers of KG into Equation (23) leads to a nested integral, which has the following compact representation on path space:

$$
\begin{equation*}
I_{j}=\int_{\mathcal{P}} f_{j}(\bar{x}) \mathrm{d} \mu(\bar{x}) \tag{25}
\end{equation*}
$$

where

$$
\begin{align*}
f_{j}(\overline{\mathbf{x}})=f_{j}\left(\mathbf{x}_{1} \cdots \mathbf{x}_{n}\right)= & L_{e}\left(\mathbf{x}_{1} \rightarrow \mathbf{x}_{2}\right)\left[\prod_{k=2}^{n-1} \bar{f}\left(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_{k} \rightarrow \mathbf{x}_{k+1}\right) \bar{G}\left(\mathbf{x}_{k-1} \leftrightarrow \mathbf{x}_{k}\right)\right] \\
& \cdot \bar{G}\left(\mathbf{x}_{n-1} \leftrightarrow \mathbf{x}_{n}\right) W_{e}^{(j)}\left(\mathbf{x}_{n-1} \rightarrow \mathbf{x}_{n}\right) \tag{26}
\end{align*}
$$

The seeming direction reversal in the importance emission function is due to the convention that " $\rightarrow$ " indicates the flow of light.

## 7 Manifold exploration for volumes

There are no specular reflections per se in the volume, but from a purely mathematical standpoint the phase function of a strongly forward-scattering volume is not unlike reflection from a rough mirror. The fast-varying part of the mirror BRDF is a function of the half-vector, and hence our method preserves it during manifold walks. In the medium case, we are interested in being able to handle highly peaked phase functions that are a function of the scattering angle. For this purpose we treat


Figure 12: Medium constraint the scattering angle as analogous to the half vector, introducing the specular manifold constraint

$$
\begin{equation*}
\mathbf{c}\left(\mathbf{x}_{i-1}, \mathbf{x}_{i}, \mathbf{x}_{i+1}\right)=T\left(\overrightarrow{\mathbf{x}_{i-1} \mathbf{x}_{i}}\right)^{T} \overrightarrow{\mathbf{x}_{i} \mathbf{x}_{i+1}} \quad\left(=\mathbf{o}_{i}\right) \tag{27}
\end{equation*}
$$

where $T(\mathbf{v})$ is a basis for the plane orthogonal to the direction $\mathbf{v}$ (Figure 12). This still leaves one DOF per vertex, which we remove by preserving the distance to the scattering event (i.e. $\left\|\mathbf{x}_{i-1}-\mathbf{x}_{i}\right\|=$ const.). While computing the entries of the constraint Jacobian $\nabla C$, we use these two constraints in place of the previous definition (Section 3.3) whenever a vertex describes a medium interaction.


Figure 13: It is possible to extend the path space specular integration framework to volumes, for instance to render caustics from refractive objects, such as this dodecahedron-shaped luminaire with tinted glass inlays.

### 7.1 Medium manifold perturbation

From an algorithmic perspective, manifold exploration for volumes is almost identical to the surface case. Our implementation handles both cases jointly, which permits constructing specular chains that contain both surface and medium interactions.

Apart from the new type of constraint (27), the computation of offset manifold (11) tangent vectors is unchanged. In the ray tracing step 4 of WalkManifold when encountering a vertex $\mathbf{x}_{i-1}$ that is followed by a medium interaction vertex $\mathbf{x}_{i}$, we set $\mathbf{x}_{i}^{\prime}=\mathbf{x}_{i-1}+\left\|\mathbf{x}_{i-1}-\mathbf{x}_{i}\right\| \mathbf{d}$, where $\mathbf{d}$ is the outgoing direction at $\mathbf{x}_{i-1}$ (this enforces the length constraint mentioned earlier). Afterwards, the manifold offset $\mathbf{o}_{i}$ is transformed into an outgoing direction in the new frame at $\mathbf{x}_{i}^{\prime}$ (Figure 12).

As in the glossy surface case, we require a criterion that clarifies when treating a medium vertex $\mathbf{x}_{i}$ as non-specular is in order, and when it is better handled by the manifold. Again, this decision is made probabilistically, based on a modified specular probability function $\psi\left(\mathbf{x}_{i}\right)$ described in the appendix.

When computing the generalized geometric term through a specular chain with medium endpoints, we consider them to be parameterized on a surface perpendicular to the direction facing the chain.

## 8 Results

We have implemented the proposed technique and prior work as extension modules to the Mitsuba renderer [Jakob 2010]. All techniques operate on top of a newly added bidirectional abstraction layer that exposes cameras, light sources, scattering models, and participating media as generalized path vertex and edge objects with a common basic interface. This greatly simplified the implementation effort, as bidirectional rendering algorithms can usually be stated much more succinctly in terms of operations on vertices that are oblivious to whether they contain e.g. a camera model or a medium scattering event.

| Scene | Seed generator |  | Perturb. | Mutations |  | Manifold walks |  |  | Manif. size |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | samples/px. | chains/px. | $\lambda$ | total | accepted | total | conv. | avg. iter. | avg. | max. |
| Torus (MEPT) | 32 | 2 | 100 | 1178 M | 78.3\% | 1002 M | 96.7\% | 2.3 | 3.4 | 7 |
| Chandelier (MEPT) | 64 | 3 | 160 | 1216 M | 73.6\% | 975 M | 97.6\% | 2.4 | 4.6 | 13 |
| Chandelier (MEMLT) | - | - | 160 | 3626 M | 34.1\% | 376 M | 98.4\% | 2.1 | 4.5 | 13 |
| Table (MEPT) | 32 | 1 | 300 | 1074 M | 77.5\% | 868 M | 95.5\% | 2.8 | 4.4 | 14 |
| Table (MEMLT) | - | - | 300 | 1921 M | 35.9\% | 772.4 M | 94.3\% | 3.2 | 4.4 | 14 |
| GlassEgG (MEPT) | 128 | 2 | 90 | 1533 M | 72.4\% | 1246 M | 92.2\% | 3.4 | 4.1 | 14 |
| GlassEgg (MEMLT) | - | - | 90 | 2774 M | 40.3\% | 1101 M | 92.5\% | 3.3 | 4.1 | 14 |

Table 1: Listing of seed generator and perturbation parameters, as well as performance statistics.
We compare the following algorithms:

- Primary sample space MLT by Kelemen et al., implemented on top of bidirectional path tracing (PSSMLT).
- Path space MLT by Veach and Guibas (MLT).
- An extended form of energy redistribution path tracing by Cline et al. (ERPT), which is seeded by bidirectional rather than unidirectional path tracing. The ERPT implementation shares the caustic, lens, and multi-chain perturbation with the previous algorithm. Since they introduce bias, we did not use the post-processing filters proposed in the original paper.
- Manifold exploration path tracing (MEPT), which is structured similarly to ERPT. We modified the original algorithm by replacing its highly specialized caustic, lens, and multichain perturbations with the manifold perturbation. Due to its general design, the new perturbation subsumes and extends the capabilities of the original set.
- Manifold Metropolis Light Transport (MEMLT), which corresponds to MLT with our perturbation (i.e. the bidirectional mutation and manifold perturbation, but none of the original perturbations from MLT).

Due to the aforementioned abstraction layer, all techniques transparently support participating media even if this was originally not part of their description. We found that MEPT generally performs better that MEMLT due to certain limitations of the bidirectional mutation that are discussed later.

The rendering of result images was conducted on Amazon EC2 cc1.4xlarge instances, which are eight-core Intel Xeon X5570 machines. A single machine was used per image. To exploit the local parallelism, our implementation runs a separate Markov chain on each core, and the resulting buffers are averaged together when exposing the image.

We have rendered three views of a challenging interior scene containing approximately 2 million triangles with shading normals and a mixture of glossy, diffuse, and specular surfaces and some scattering volumes. One hour of processing time was allocated to each rendering technique, and a comparison of the resulting images is shown in Figures 14, 15, and 16. The converged images in Figure 1 were rendered using 48 hours. The one hour renderings are intentionally unconverged to permit a visual analysis of the convergence behavior. Table 1 lists parameters and statistics collected during these renderings. The path generator columns refer to the seeding scheme used by ERPT and MEPT, which samples and subsequently resamples a number of paths
per pixel before launching Markov Chains. The statistics include the total number of mutations and acceptance ratio, as well as the the convergence behavior of the manifold walks and vertex count of encountered manifolds.

Chandelier: In this set of results, the poor performance of MLT is most apparent and is caused by the ineffectiveness of the bidirectional mutation in finding long specular paths. Because it must decide up front on the configuration of a path before generating it, most of the time the mutation fails, resulting in acceptance rates under $1 \%$. Consequently, too few jumps between disjoint connected components of path space occur, causing parts of the image have an incorrect relative brightness. This weakness is inherited by MEMLT, which also builds upon the bidirectional mutation. It is more successful at exploring some of the diffuse-specular-diffuse paths through the bulbs but overall does not work well on this scene. Seeding the same perturbations using bidirectional path tracing, which does not suffer from this disadvantage, performs much better, as can be seen in the ERPT and MEPT renderings.

Table: This scene is lit by the chandelier, with its glass-enclosed sources, so all illumination is by specular paths. By reasoning about the geometry of the specular and offset specular manifolds for the paths it encounters, our perturbation rule is more successful at rendering paths-such as illumination that refracts from the bulbs into the butter dish, then to the camera ( 6 specular vertices) - that the other methods struggle with. The MLT rendering looks too dark, because it did not find enough of these paths and mainly captures diffuse illumination from the walls. The noise in the ERPT result reveals that the underlying bidirectional path tracer encountered some of those paths but the Veach-Guibas perturbations are not able to explore path space around them effectively. The primary sample space MLT variant also has difficulties rendering this scene, because it has no knowledge about the underlying path geometry. MEMLT produces a clean result, but the relative brightness of different parts of the image is incorrect due to bidirectional mutation's difficulty in performing sufficiently many jumps between them.

GlassEgG: In this scene, our technique's ability to create a specular chain containing both medium and surface interactions leads to fast convergence when rendering the forward-scattering medium inside the glass egg. MLT and ERPT perform poorly here, since they do not have suitable perturbations for exploring this space. Because the MLT perturbations treat glossy and diffuse materials identically, they have difficulty rendering the near-specular tabletop, producing streak-like artifacts in the output rendering.

## 9 Conclusions

We have presented a new type of Markov Chain Monte Carlo rendering method that models the space of valid specular and near-specular light paths using high-dimensional manifolds. The differential geometry of these manifolds provides a powerful tool to efficiently explore these paths, which can be very hard for previous methods to find. Our technique applies in the frameworks of MLT or ERPT, producing rendering algorithms with support for specular paths fundamentally built in at the core. In equal-time comparisons on very challenging scenes, the new Manifold Exploration Path Tracing algorithm compares favorably to previous work in Monte Carlo and MCMC rendering.

This new algorithm does still share certain limitations with its predecessors. Most important, it needs well distributed seed paths, because it can only explore connected components of the


Figure 14: Chandelier: This view contains a brass chandlier with 24 light bulbs, each surrounded by a glass enclosure. The chandlier uses a realistic metal material based on microfacet theory and is attached to the ceiling using specular metal cylinders. This scene is challenging, as certain important light paths are found with low probability, particularly those involving interreflection between the bulbs and the body of the chandelier. In this and the following comparisons, one hour of processing time was allocated to each rendering technique.


Figure 14: Chandelier (continued)


Figure 15: TABLE: This view of our room scene shows chinaware (using a BRDF with both diffuse and specular components), a teapot containing an absorbing medium, and a butter dish on a glossy silver tray. Illumination comes from the chandelier in Figure 14.


Figure 15: TABLE (continued)


Figure 16: GlassEgG: This view of a different part of the room, now lit through windows using a spherical environment map surrounding the scene, contains a forward-scattering participating medium ( $g=0.8$ ) inside the glass egg.
manifold for which seed paths are provided. Bidirectional Path Tracing is reasonably effective but still has trouble finding many components, and this problem fundamentally becomes more and more difficult as the number of path types increases. Ultimately, as the number of different path types exceeds the number of samples that can be generated, local exploration of path space becomes ineffective; future algorithms could be designed to attempt exploration only in sufficiently large path space components.

Unlike many methods for caustics and other specular phenomena, we have shown how to generalize Manifold Exploration almost trivially to handle glossy surfaces and volumes. Similar refinements can let the same method handle perfectly anisotropic reflections, strongly oriented volume scattering media, and other kinds of problems with exactly or approximately constrained paths.

While MCMC rendering is a natural match for our methods of dealing with specular paths, our predictor-corrector iteration can be used in other kinds of algorithms as well, including deterministic ones to map out specular paths, for instance in design of luminaires or optical systems.


Figure 16: GlassEgG (continued)

## 10 Acknowledgements

This research was conducted in conjunction with the Intel Science and Technology Center for Visual Computing. Additional funding was provided by the National Science Foundation under grant IIS-1011919. The authors are indebted to Olesya Isaenko, who meticulously crafted the example scenes used in the evaluation.

## Appendix A1

Specular probability function for surfaces: We found the following following heuristic based on microfacet theory to work well: when the microsurface normals at $\mathbf{x}_{i}$ follow a distribution $D_{\alpha}\left(\mathbf{m}_{i}\right)$ with roughness parameter $\alpha$, the BSDF at $\mathbf{x}_{i}$ will take on small values when $\mathbf{m}_{i}$ moves into a region where $D_{\alpha}\left(\mathbf{m}_{i}\right)$ has low density. We thus set $\psi\left(\mathbf{x}_{i}\right)$ by computing the expected probability that treating vertex $\mathbf{x}_{i}$ as non-specular during a manifold perturbation would move its microsurface normal $\mathbf{m}_{i}$ from a region of high density to one of low density, and we choose the 90th-percentile to classify the support of $D_{\alpha}$ into such regions.


To obtain the specular probability, our implementations must know the expected angular change $\Delta \theta$ of microsurface normals during a perturbation, which is found by briefly running the Markov chain before rendering starts. During rendering, $\psi\left(\mathbf{x}_{i}\right)$ is computed as the area ratio of the two highlighted regions on the sphere:

$$
\begin{equation*}
\psi\left(\mathbf{x}_{i}\right)=\frac{1-\cos \theta_{q}\left(\alpha\left(\mathbf{x}_{i}\right)\right)}{1-\cos \left(\theta_{q}\left(\alpha\left(\mathbf{x}_{i}\right)\right)+\Delta \theta\right)}, \tag{28}
\end{equation*}
$$

where $\theta_{q}$ is the aforementioned percentile (with $q$ set to 0.9 ). For the Beckmann distribution, this is given by

$$
\theta_{q}(\alpha):=\tan ^{-1}\left(-\alpha^{2} \log (1-q)\right) .
$$

Implementation-wise, this heuristic requires material models to be able to compute their Beckmann distribution-equivalent roughness or provide a custom quantile function.

Specular probability function for participating media: In the medium case, we use the same probability (28), but now with a percentile that is suitable for volumetric scattering. We use

$$
\theta_{q}(g)=\cos ^{-1} \frac{(1+|g|)^{2}-2(1+|g|)\left(1+g^{2}\right) q+2|g|\left(1+g^{2}\right) q^{2}}{(1+|g|-2|g| q)^{2}}
$$

where $g$ is the mean cosine of the phase function, $q$ is set to 0.5 , and $\theta_{q}$ was derived from the Henyey-Greenstein phase function. In $\Psi\left(\mathbf{x}_{i}\right)$ (Equation 28) we must also replace $\Delta \theta$ with the average change in scattering angle at medium vertices, again determined in a brief phase before rendering.

## Appendix A2

Let us assume the existence of a Vertex data structure with the following contents:

```
struct Vertex {
    Point p; // Vertex position
    Vector dpdu, dpdv; // Tangent vectors
    Normal n; // Normal vector
    Vector dndu, dndv; // Normal derivatives
    float eta; // Relative index of refraction
    Matrix2x2 A, B, C; // Matrix blocks of \nablaC in the row
    // associated with the current vertex
};
```

Then the following $\mathrm{C}++$ code computes the derivatives of $\nabla C$ associated with a single surface interaction vertex. It assumes that the function is called with a pointer into a list of vertices.

```
void computeDerivatives(Vertex *v) {
    /* Compute relevant directions and a few useful projections */
    Vector wi = v[-1].p - v[0].p;
    Vector wo = v[ 1].p - v[0].p;
    float ili = 1/wi.length();
    float ilo = 1/wo.length();
    wi *= ili; wo *= ilo;
    Vector H = wi + v[0].eta * wo;
    float ilh = 1/H.length();
    H *= ilh;
    float dot_H_n = dot(v[0].n, H),
            dot_H_dndu = dot(v[0].dndu, H),
            dot_H_dndv = dot(v[0].dndv, H),
            dot_u_n = dot(v[0].dpdu, v[0].n),
            dot_v_n = dot(v[0].dpdv, v[0].n);
    /* Local shading tangent frame */
    Vector s = v[0].dpdu - dot_u_n * v[0].n;
    Vector t = v[0].dpdv - dot_v_n * v[0].n;
    ilo *= v[0].eta * ilh; ili *= ilh;
    /* Derivatives of C with respect to x_{i-1} */
    Vector
        dH_du = (v[-1].dpdu - wi * dot(wi, v[-1].dpdu)) * ili,
        dH_dv = (v[-1].dpdv - wi * dot(wi, v[-1].dpdv)) * ili;
    dH_du -= H * dot(dH_du, H);
    dH_dv -= H * dot(dH_dv, H);
    v[0].A = Matrix2x2(
        dot(dH_du, s), dot(dH_dv, s),
        dot(dH_du, t), dot(dH_dv, t));
    /* Derivatives of C with respect to x_i */
    dH_du = -v[0].dpdu * (ili + ilo) + wi * (dot(wi, v[0].dpdu) * ili)
        + wo * (dot(wo, v[0].dpdu) * ilo);
    dH_dv = -v[0].dpdv * (ili + ilo) + wi * (dot(wi, v[0].dpdv) * ili)
        + wo * (dot(wo, v[0].dpdv) * ilo);
    dH_du -= H * dot(dH_du, H);
    dH_dv -= H * dot(dH_dv, H);
    v[0].B = Matrix2x2(
        dot(dH_du, s) - dot(v[0].dpdu, v[0].dndu) * dot_H_n - dot_u_n * dot_H_dndu,
        dot(dH_dv, s) - dot(v[0].dpdu, v[0].dndv) * dot_H_n - dot_u_n * dot_H_dndv,
        dot(dH_du, t) - dot(v[0].dpdv, v[0].dndu) * dot_H_n - dot_v_n * dot_H_dndu,
```

```
    dot(dH_dv, t) - dot(v[0].dpdv, v[0].dndv) * dot_H_n - dot_v_n * dot_H_dndv);
    /* Derivatives of C with respect to x_{i+1} */
    dH_du = (v[1].dpdu - wo * dot(wo, v[1].dpdu)) * ilo;
    dH_dv = (v[1].dpdv - wo * dot(wo, v[1].dpdv)) * ilo;
    dH_du -= H * dot(dH_du, H);
    dH_dv -= H * dot(dH_dv, H);
    v[0].C = Matrix2x2(
    dot(dH_du, s), dot(dH_dv, s),
    dot(dH_du, t), dot(dH_dv, t));
}
```


## References

Arvo, J. R. 1995. Analytic Methods for Simulated Light Transport. PhD thesis, Yale University.
Chen, M., and Arvo, J. 2000. Perturbation methods for interactive specular reflections. IEEE Trans. Vis. and Comp. Graph. 6, 3 (July/Sept.), 253-264.

Chen, M., and Arvo, J. 2000. Theory and application of specular path perturbation. ACM Trans. Graph. 19, 4 (Oct.), 246-278.

Chen, J., Wang, B., and Yong, J.-H. 2011. Improved stochastic progressive photon mapping with metropolis sampling. Computer Graphics Forum 30, 4, 1205-1213.

Cline, D., Talbot, J., and Egbert, P. 2005. Energy redistribution path tracing. ACM Trans. Graph. 24, 3 (Aug.), 1186-1195.

Cook, R. L., Porter, T., and Carpenter, L. 1984. Distributed ray tracing. In Computer Graphics (Proceedings of SIGGRAPH 84), 137-145.

Goral, C. M., Torrance, K. E., Greenberg, D. P., and Battaile, B. 1984. Modeling the interaction of light between diffuse surfaces. In Computer Graphics (Proceedings of SIGGRAPH 84), 213-222.

Hachisuka, T., and Jensen, H. W. 2009. Stochastic progressive photon mapping. ACM Trans. Graph. 28, 5 (Dec.).

Hachisuka, T., and Jensen, H. W. 2011. Robust adaptive photon tracing using photon path visibility. ACM Trans. Graph. 30, 5 (Oct.), 114:1-114:11.

Hastings, W. K. 1970. Monte carlo sampling methods using markov chains and their applications. Biometrika 57, 1, 97-109.

Heckbert, P. S. 1990. Adaptive radiosity textures for bidirectional ray tracing. In Computer Graphics (Proceedings of SIGGRAPH 90), 145-154.

Igehy, H. 1999. Tracing ray differentials. In Computer Graphics (Proceedings of SIGGRAPH 99), 179-186.

Jakob, W., 2010. Mitsuba renderer. http://www.mitsuba-renderer.org.
Jarosz, W., Zwicker, M., and Jensen, H. W. 2008. The beam radiance estimate for volumetric photon mapping. Computer Graphics Forum 27, 2 (Apr.), 557-566.

Jensen, H. W., and Christensen, P. H. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In Computer Graphics (Proceedings of SIGGRAPH 98), 311-320.

Jensen, H. W. 1996. Global illumination using photon maps. In Eurographics Rendering Workshop 1996, 21-30.

Kajiya, J. T. 1986. The rendering equation. In Computer Graphics (Proceedings of SIGGRAPH 86), 143-150.

Kelemen, C., Szirmay-Kalos, L., Antal, G., and Csonka, F. 2002. A simple and robust mutation strategy for the metropolis light transport algorithm. Computer Graphics Forum 21, 3, 531-540.

Kitaoka, S., Kitamura, Y., and Kishino, F. 2009. Replica exchange light transport. Computer Graphics Forum 28, 8 (Dec.), 2330-2342.

Lafortune, E. P., and Willems, Y. D. 1993. Bi-directional path tracing. In Proceedings of Compugraphics 93.

Lai, Y.-C., Fan, S. H., Chenney, S., and Dyer, C. 2007. Photorealistic image rendering with population monte carlo energy redistribution. In Rendering Techniques 2007: 18th Eurographics Workshop on Rendering, 287-296.

Liu, J. S. 2001. Monte Carlo strategies in scientific computing. Springer.
Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. 1953. Equation of state calculations by fast computing machines. J. Chem. Phys. 21, 6.

Mitchell, D. P., and Hanrahan, P. 1992. Illumination from curved reflectors. In Computer Graphics (Proceedings of SIGGRAPH 92), 283-291.

Pauly, M., Kollig, T., and Keller, A. 2000. Metropolis light transport for participating media. In Rendering Techniques 2000: 11th Eurographics Workshop on Rendering, 11-22.

Pharr, M., and Humphreys, G. 2004. Physically Based Rendering: From Theory to Implementation. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

Premoze, S., Ashikhmin, M., and Shirley, P. 2003. Path integration for light transport in volumes. In Eurographics Symposium on Rendering: 14th Eurographics Workshop on Rendering.

Segovia, B., Iehl, J., and Péroche, B. 2007. Metropolis instant radiosity. Computer Graphics Forum 26, 3 (Sept.), 425-434.

Shirley, P. S., Wade, B., Hubbard, P., Zareski, D., Walter, B., and Greenberg, D. P. 1995. Global illumination via density estimation. In Eurographics Rendering Workshop 1995, 219-231.

Sillion, F. X., and Puech, C. 1989. A general two-pass method integrating specular and diffuse reflection. In Computer Graphics (Proceedings of SIGGRAPH 89), 335-344.

Spivak, M. 1965. Calculus on Manifolds. Addison-Wesley.
Veach, E., and Guibas, L. 1994. Bidirectional estimators for light transport. In Fifth Eurographics Workshop on Rendering.

Veach, E., and Guibas, L. J. 1997. Metropolis light transport. In Computer Graphics (Proceedings of SIGGRAPH 97), 65-76.

Veach, E. 1997. Robust Monte Carlo Methods for Light Transport Simulation. PhD thesis, Stanford University.

Walter, B., Hubbard, P. M., Shirley, P. S., and Greenberg, D. F. 1997. Global illumination using local linear density estimation. ACM Trans. Graph. 16, 3 (July), 217-259.

Walter, B., Marschner, S. R., Li, H., and Torrance, K. E. 2007. Microfacet models for refraction through rough surfaces. In Rendering Techniques 2007: 18th Eurographics Workshop on Rendering, 195-206.

Walter, B., Zhao, S., Holzschuch, N., and Bala, K. 2009. Single scattering in refractive media with triangle mesh boundaries. ACM Trans. Graph. 28, 3 (July), 92:1-92:8.

Whitted, T. 1980. An improved illumination model for shaded display. Communications of the ACM 23, 6 (June), 343-349.


[^0]:    ${ }^{1}$ Because it is possible to have several separated specular paths joining two points, the parameterization cannot be global.
    ${ }^{2}$ If the parameterizations are not orthonormal, two additional determinants are required to account for the change in area.

[^1]:    ${ }^{3}$ As a somewhat unfortunate consequence of notations in different fields, in volume scattering models, the incident direction arguments points towards the scattering location, whereas it points away in the surface case.

