

Path Integration for Light Transport in Volumes

Simon Premože,¹ Michael Ashikhmin² and Peter Shirley¹

¹ Computer Science Department, University of Utah

² Computer Science Department, SUNY at Stony Brook

Abstract

Simulating the transport of light in volumes such as clouds or objects with subsurface scattering is computationally expensive. We describe an approximation to such transport using path integration. Unlike the more commonly used diffusion approximation, the path integration approach does not explicitly rely on the assumption that the material within the volume is dense. Instead, it assumes the phase function of the volume material is strongly forward scattering and uniform throughout the medium, an assumption that is often the case in nature. We show that this approach is useful for simulating subsurface scattering and scattering in clouds.

1. Introduction

The appearance of many materials (e.g., skin, fruits, snow, clouds) cannot be described by a simple BRDF-style reflectance model. The main reason for this is volumetric scattering which manifests itself in important lighting effects for materials^{4, 5, 19} and scenes³⁶. The radiative transfer equation and propagation of light in a scattering medium have been both analytically and numerically studied in astrophysics, atmospheric optics, and more recently medical applications. For most problems with non-trivial boundary conditions, phase functions and initial conditions there are no analytic solutions. Most solutions are based on the diffusion approximation which assumes that enough scattering events have occurred for light to be uniformly scattered in all directions. This approximation has proven useful for generating images with subsurface scattering¹⁸. Monte Carlo methods are also often used to compute radiative transport within a medium. Although simple and powerful, these methods suffer from slow convergence. Finite element methods are also used, but they require large amounts of storage to capture discontinuities and strong directional light distributions.

Because the diffusion approximation is only appropriate for dense uniform media¹⁹, there is a gap in the computer graphics literature when accurate approximations are desired for sparse or non-uniform media. This paper attempts to fill that gap using an alternative to the diffusion approximation based on Feynman's path integral approach to solving quantum mechanics problems¹⁰. Path integral formulations of physical processes have been used in physics to solve a

wide variety of problems including energy propagation in random media³⁸ and transfer equation³⁹.

The radiative transfer equation describing light propagation can be viewed as a collection of paths taken by radiation as it travels through space. A *path integral* is an integral over all such possible paths traveled by a photon. Radiative transfer is decomposed into a series of smaller problems formulated by the Green function propagator. Scattered optical fields are described using the concept of an ensemble of effective optical paths of partial contributions. This physical picture for treating light transport in multiple scattering media as a collection of most probable paths provides insight into the light propagation in a medium. Unlike the randomized approach to using paths of Metropolis Light Transport⁴⁵, path integral methods analytically find the most important paths and develop analytical estimates based on them. Once the most important path is found, the multiple scattering contributions are only computed along the most probable paths and the rest of the paths are dealt with implicitly via analytic integration of multiple scattering using well-known approximations. We therefore avoid computationally expensive direct numerical simulation of multiple scattering in the medium. We provide some intuition behind path integrals and demonstrate that some useful results of the theory can be obtained without any heavy mathematical tools. We discuss solutions to light propagation as a path integral (a formal sum) and how they can be used for rendering arbitrary scattering materials and media. Because many common phenomena cannot be described using only single scattering, we discuss some observable consequences of multiple scat-

tering and how they can be exploited for rendering applications.

We show results for inhomogeneous media that cannot be achieved using the diffusion approximation. The restrictions on our method are different; while sparse and inhomogeneous media are allowed, we assume that the density and therefore scattering coefficients vary smoothly, the phase function is arbitrary but constant within the medium. This expands the class of problems that can be attacked efficiently. We note that this paper serves mainly as an introduction to path integral methods and many improvements in efficiency have not yet been explored in our implementation.

2. Background and Previous Work

Light Transport Approximations

There has been much work in approximating radiative transfer in arbitrary media in many fields including computer graphics. Perez *et al.*³⁰ survey and classify global illumination algorithms in participating media in detail. Pharr and Hanrahan³¹ and Premože³⁵ also provide an extensive list of existing methods and background. Here we briefly review recent methods proven practical and robust.

Jensen and Christensen¹⁶ presented a two pass photon density estimation method. This method is simple, robust and efficient but suffers from large memory requirements to store photons if the extent of the scene is large or the lighting configuration is very complex. The method is very practical and it has been used for many phenomena including smoke⁹, fire²⁶, stone⁸, and wet materials¹⁷ proving its generality. On the other hand, the photon map becomes rather inefficient in highly scattering media¹⁸.

Veach and Guibas⁴⁵ presented a global illumination algorithm that found important paths, and then explored the path space locally because it was likely that other important paths would be nearby. Pauly *et al.*²⁷ extended the method for participating media and proposed suitable mutation strategies for paths. Although extremely general and robust, as it could handle any lighting condition and configuration, it still suffers from the classical Monte Carlo problems of noise and slow convergence.

Stam³⁷ presented a solution to multiple scattering in non-homogeneous materials by solving the diffusion equation using a multigrid method. Jensen *et al.*¹⁸ introduced an analytical solution to the diffusion approximation to multiple scattering, which is especially applicable for translucent materials that exhibit considerable subsurface light transport. Their method relies on the assumption that the multiply-scattered light is nearly isotropic and cannot be easily extended to inhomogeneous materials. Lensch *et al.*²² implemented this method in graphics hardware and Jensen and Buhler¹⁵ extended this diffusion approximation to be computationally more efficient by precomputing and storing illu-

mination in a hierarchical grid. Narasimhan and Nayar²⁵ described a physically-based multiple scattering model for simulating weather effects such as fog, haze, mist and rain.

An alternative description of light propagation was done by Pharr and Hanrahan³¹ who described a mathematical framework for solving the scattering equation in the context of a variety of rendering problems and a numerical solution in terms of Monte Carlo sampling. The scattering equation describes all scattering events inside the object and it does not depend on the incoming illumination. Unfortunately, there has been no other work exploiting this interesting paradigm of scattering objects and interactions between objects on larger scales. Lafortune²⁰ also described the *global reflectance distribution function* (GRDF) which corresponds to the scattering equation idea.

Path Integral Methods

Path integral techniques and functional integration have been widely used in statistical and quantum mechanics to solve propagators for Fokker-Planck and Schrödinger equations. The concept of photon paths has been well-known in the theory of energy propagation in random media³⁸. Tessendorf^{39, 40} used the path integral approach to study the propagation of light in weakly scattering media such as water. Perelman *et al.*^{28, 29} described energy transport in a turbid medium using a quasi-particle Lagrangian, from which the most probable paths could be found. Wilson and Wang⁴⁶ constructed a Lagrangian through a turbid medium using local path descriptors. Miller²⁴ also constructed a stochastic Lagrangian path integral representation for Green's evolution operator. Constantinou and Demetrescu⁷ showed the equivalence of the path integral formulation and virtual rays. Gross¹² studied multiple scattering of a wave in a system of random and uncorrelated scattering particles. The path integral methods have also been studied in optical tomography and medical imaging. Near-diffusive scattering regimes are very important for obtaining diagnostic information about multilayer biological tissues where standard diffusion approximation fails^{32, 33}. Jacques and Wang¹⁴ presented a basic introduction to the path integral description of photon transport and discussed classical path for describing the most probable path of a photon.

3. Mathematical Preliminaries

MEDIUM PROPERTIES AND LIGHT PATHS. In an arbitrary medium, the underlying optical properties depend on bulk material properties such as density $\rho(\mathbf{x})$, temperature $T(\mathbf{x})$, and the particle absorption and scattering cross-sections, σ_a and σ_s . Optical properties of the medium are then described in terms of the scattering coefficient $b(\mathbf{x}) = \sigma_s \rho(\mathbf{x})$, the absorption coefficient $a(\mathbf{x}) = \sigma_a \rho(\mathbf{x})$, the extinction coefficient $c(\mathbf{x}) = a(\mathbf{x}) + b(\mathbf{x})$, and the phase function $P(\mathbf{x}, \vec{\omega}, \vec{\omega}')$. The phase function P describes the probability density of light




Figure 1: Scattering in a highly scattering medium. Original radiance undergoes a series of scattering events that result in angular, spatial and temporal spreading of the original radiance distribution.

\mathbf{x}	Generic location in \mathbb{R}^3
$\vec{\omega}$	Generic direction
$a(\mathbf{x})$	Absorption coefficient at a point
$b(\mathbf{x})$	Scattering coefficient at a point
$c(\mathbf{x})$	Extinction coefficient at a point
g	Mean cosine of the scattering angle
Q	Volume source distribution
$P(\mathbf{x}, \vec{\omega}, \vec{\omega}')$	Phase function
G	Propagator (Green function)
$\vec{\gamma}_p(s)$	Light point path
$\vec{\beta}(s)$	Light propagation direction on the path
$\mathbf{d}(s)$	Displacement along the path
s	Distance along the path (arclength)
S	Arclength of the path $\vec{\gamma}$
ℓ	Spatial variability
$\langle \theta^2 \rangle$	Mean square scattering angle

Table 1: Notations for important terms used in the paper.

coming from incident direction $\vec{\omega}$ scattering into direction $\vec{\omega}'$ upon scattering event at point \mathbf{x} . The phase function is normalized so that $\int_{4\pi} P(\vec{\omega}, \vec{\omega}') d\omega' = 1$ and in most practical settings only depends on the phase angle $\cos \theta = \vec{\omega} \cdot \vec{\omega}'$. The mean cosine g of the scattering angle is defined:

$$g = \int_{4\pi} P(\vec{\omega}, \vec{\omega}') (\vec{\omega} \cdot \vec{\omega}') d\omega'.$$

These optical properties are *inherent*, because they depend solely on the medium and not on the structure of the incoming light field. Upon entering the medium, incoming light undergoes a series of scattering and absorption events that modify both the directional structure of the incoming light field and its intensity. As a result of multiple scattering events, the original radiance distribution undergoes angular, spatial and temporal spreading which result in different radiance distribution. Figure 1 shows spreading effects in an arbitrary highly scattering medium. Table 1 summarizes important terms and quantities used in the paper.

Consider a photon that originates at point \mathbf{x}' with direction $\vec{\omega}'$ and traveling in a medium along a curved arclength pa-




Figure 2: Transfer geometry in an enclosed scattering medium. A photon originating at point \mathbf{x}' with direction $\vec{\omega}'$ travels along curved path $\vec{\gamma}_p$ of length s until it reaches the final point \mathbf{x} with direction $\vec{\omega}$. $\vec{\beta}(s')$ is the direction of propagation at arclength parameter s' on the curved path.

rameterized path $\vec{\gamma}_p(s)$ until it reaches the final point \mathbf{x} with direction $\vec{\omega}$, (see Figure 2). This path results from an accumulated random walk of propagation directions governed by scattering and absorption events along the path. Because the path $\vec{\gamma}_p$ is in general curved, its total length S is greater than the distance between starting and ending points: $|\mathbf{x}' - \mathbf{x}| \leq S$. The direction of propagation along the path is defined by $\vec{\beta}(s) = \frac{d\vec{\gamma}_p(s)}{ds}$. The path therefore satisfies the “two-sided” boundary conditions:

$$\begin{aligned} \vec{\gamma}_p(0) &= \mathbf{x}' & \vec{\beta}(0) &= \vec{\omega}' \\ \vec{\gamma}_p(S) &= \mathbf{x} & \vec{\beta}(S) &= \vec{\omega}. \end{aligned} \quad (1)$$

The displacement relative to the point \mathbf{x}' at distance s is obtained by integrating $\vec{\beta}$:

$$\mathbf{d}(s) = \int_0^s \vec{\beta}(s') ds'.$$

Light undergoes a series of scattering and absorption events along the path. Note that if we ignore exact backscattering which returns photons back into the same path, the intensity of “original” light will be only diminishing because of these events since in-scattering will be due to photons traveling a *different* path in the medium. Therefore, if we introduce the effective attenuation τ which determines how much will the light intensity be reduced along the length of the path, the radiance L in the medium will be proportional to

$$L \sim \sum_{\text{all paths}} e^{-\tau(\text{path})}.$$

The path integral formulation of light transport is essentially a mathematically rigorous expression of this simple idea.

RADIATIVE TRANSFER. Light transport in arbitrary media is described by the radiative transport equation^{2, 13}:

$$\begin{aligned} (\vec{\omega} \cdot \nabla + c(\mathbf{x}))L(\mathbf{x}, \vec{\omega}) = \\ b(\mathbf{x}) \int_{4\pi} P(\vec{\omega}, \vec{\omega}') L(\mathbf{x}, \vec{\omega}') d\omega' + Q(\mathbf{x}, \vec{\omega}), \end{aligned} \quad (2)$$

where $Q(\mathbf{x}, \vec{\omega})$ is the source term. In computer graphics, the source terms $Q(\mathbf{x}, \vec{\omega})$ is often due to light emitted by the

medium itself ($L_e(\mathbf{x}, \vec{\omega})$). It is often convenient to split the total radiance within the medium into components and write it as the sum of *unscattered* (direct) radiance L_{un} , the *emission* L_e and the *scattered* radiance L_{sc} :

$$L(\mathbf{x}, \vec{\omega}) = L_{un}(\mathbf{x}, \vec{\omega}) + L_{sc}(\mathbf{x}, \vec{\omega}) + L_e(\mathbf{x}, \vec{\omega}).$$

Here L_{un} is the radiance which intensity has been reduced due to absorption and outscattering along the length S . L_{sc} is the radiance that has undergone a series of scattering events and finally scattered into a small cone around the observation direction $\vec{\omega}$.

PROPAGATOR FOR RADIATIVE TRANSFER. The solution of equation 2 is the limit of the corresponding solution for the time-dependent problem where radiance L varies over time t . It is convenient to express time in units of length s as $t = s/v$ where v is the speed of light in the medium. With this notation, the time-dependent radiative transfer (TDRT) equation is

$$\left(\frac{\partial}{\partial s} + \vec{\omega} \cdot \nabla + c(\mathbf{x})\right)L(s, \mathbf{x}, \vec{\omega}) = b(\mathbf{x}) \int_{4\pi} P(\vec{\omega}, \vec{\omega}')L(s, \mathbf{x}, \vec{\omega}')d\omega' + Q(s, \mathbf{x}, \vec{\omega}). \quad (3)$$

The solution of the TDRT equation can be formulated in terms of the Green evolution operator G which is also called the propagator, the Green function, or the *point spread function* (PSF). It is defined as the solution of homogeneous equation

$$\left(\frac{\partial}{\partial s} + \vec{\omega} \cdot \nabla + c(\mathbf{x})\right)G(s, \mathbf{x}, \vec{\omega}, \mathbf{x}', \vec{\omega}') = b(\mathbf{x}) \int_{4\pi} P(\vec{\omega}, \vec{\omega}')G(s, \mathbf{x}, \vec{\omega}'', \mathbf{x}', \vec{\omega}'')d\omega'', \quad (4)$$

with the initial condition

$$G(s=0, \mathbf{x}, \vec{\omega}, \mathbf{x}', \vec{\omega}') = \delta(\mathbf{x} - \mathbf{x}')\delta(\vec{\omega} - \vec{\omega}').$$

Physically, the Green propagator $G(s, \mathbf{x}, \vec{\omega}, \mathbf{x}', \vec{\omega}')$ represents the radiance at point \mathbf{x} in direction $\vec{\omega}$ at time s due to light emitted at time zero by a point directional light source located at \mathbf{x}' shining in direction $\vec{\omega}'$. For example, in the absence of scattering ($b = 0$), the solution for the propagator G is

$$G(s, \mathbf{x}, \vec{\omega}, \mathbf{x}', \vec{\omega}') = \delta(\mathbf{x} - \vec{\omega}s - \mathbf{x}')\delta(\vec{\omega} - \vec{\omega}') \exp\left(-\int_0^s a(\mathbf{x} - \vec{\omega}(s-s'))ds'\right).$$

Here the light travels in a straight line and is attenuated by the absorption coefficient $a(\mathbf{x})$. One can see that in this case, the formulation using the propagator is generally equivalent to simple raytracing. The Green propagator $G(s, \mathbf{x}, \vec{\omega}, \mathbf{x}', \vec{\omega}')$ represents the angular distribution and density of rays at point \mathbf{x} in direction $\vec{\omega}$ generated by point \mathbf{x}' in direction $\vec{\omega}'$ and is therefore equivalent to raytracing.

The concept of the Green function has been used in neutron transport theory³ providing an approach to solving radiative transfer problems with arbitrary boundary conditions

by first finding solution G of the elementary transfer problems stated above and then forming the complete solution of equation 3 by using the superposition principle, i.e. integrating G with the initial radiance distribution:

$$L(s, \mathbf{x}, \vec{\omega}) = \int G(s, \mathbf{x} - \mathbf{x}', \vec{\omega}, \vec{\omega}')L_0(\mathbf{x}', \vec{\omega}')d\mathbf{x}'d\omega',$$

where $G(s, \mathbf{x} - \mathbf{x}', \vec{\omega}, \vec{\omega}')$ is the evolution operator and $L_0(\mathbf{x}', \vec{\omega}')$ is the initial distribution. Conceptually, the notion of the evolution operator is equivalent to the idea of *radiative process* introduced by Preisendorfer³⁴. Note that the boundary value problem for the Green function is actually adjoint to the radiative transfer problem, but due to reciprocity (time-reversal invariance) we can solve the light transport problem by reversing the direction of light propagation¹.

PATH INTEGRALS. Consider the problem of finding the probability that a point particle at the initial position \mathbf{x}_i and time t_i will reach a final position \mathbf{x}_f and time t_f , i.e. will “propagate” (in the quantum mechanical sense) from \mathbf{x}_i to \mathbf{x}_f . This quantity can be expressed using quantum mechanical propagator $G(t_f - t_i, \mathbf{x}_f, \mathbf{x}_i)$ which for this problem is the solution of the Schrödinger equation subject to appropriate initial conditions. If the problem is broken down into a series of shorter time steps with propagators $G(t, \mathbf{x}, \mathbf{x}')$, the full propagator is expressed as a superposition of “smaller” Green function:

$$G(t, \mathbf{x}_i, \mathbf{x}_f) = \lim_{N \rightarrow \infty} \int \dots \int G(t/N, \mathbf{x}_i, \mathbf{x}_1) \dots G(t/N, \mathbf{x}_{N-1}, \mathbf{x}_f) d\mathbf{x}_1 \dots d\mathbf{x}_{N-1}.$$

The object on the right-hand side is called a *path integral*¹⁰ and showed that with an appropriate definition of differential measure $\mathcal{D}x$ in the infinite-dimensional path space one can write for the quantum mechanical amplitude of a propagating point particle

$$\langle x_f, t_f | x_i, t_i \rangle = \int \mathcal{D}x(t) e^{iA[x(t)]/\hbar}$$

where the weight factor contains the classical action $A(\mathbf{x}_i, \mathbf{x}_f, t_f - t_i)$ for each path. The *classical action* is the integral of the Lagrangian over the time the trajectory traverses. The path taken by a classical trajectory could be none other than the one that minimizes the classical action.

RADIATIVE TRANSFER AS A SUM OVER PATHS. The path integral (PI) approach provides a particular way to express the propagator $G(s, \mathbf{x}, \vec{\omega}, \mathbf{x}', \vec{\omega}')$. It is based on the simple observation that the full process of energy transfer from one point to another can be thought of as a sum over transfer events taking place along many different paths connecting points \mathbf{x} and \mathbf{x}' , each subjected to boundary conditions restricting path directions at these points to $\vec{\omega}$ and $\vec{\omega}'$, respectively. The full propagator is then just an integral of individual path contribution over all such paths. This object is the path integral defined above. Note that this is different than the terminology used in Veach and Guibas⁴⁵.

Because the integration is performed over the infinite-dimensional path space using the not very intuitive differential measure defined for it, the mathematics of path integrals is quite complex^{10, 21}. Tessendorf⁴¹ derived a path integral expression for the propagator G in homogeneous materials. Interested readers are referred to his further work^{42, 43} for detailed derivations of the path integral formulation. We present here a more intuitive approach sufficient for our purposes and state results from the literature without derivations.

4. Practical Path Integrals

4.1. The Most Probable Path

We first attempt to formulate conditions to find the most probable path (MPP) among all possible ones which satisfy the necessary boundary conditions. We will then assume that the full propagator can be sufficiently approximated by accounting only for contributions from paths “close to” this special one. Formally, this approach corresponds to evaluating the path integral using WKB (*Wentzel-Kramers-Brillouin*) expansion⁶, which is a well-established method from perturbation theory. In practice, this means that once the MPP (or its approximation) is identified, we simply consider its blurred contribution which approximates the effect of surrounding paths. The details of how this is done are presented in Section 5 below, but one can already see some potential advantages of this approach over more traditional methods. Compared with Monte-Carlo techniques which perform statistical sampling of random paths, the path integral approach attempts to find the most important ones directly and can therefore be considered an extreme form of variance reduction. The PI formulation also does not explicitly rely on further assumptions about the character of radiance distribution which are needed, for example, in diffusion approximation¹¹. However, if warranted, we can take direct computational advantage of the fact that the radiance distribution becomes more and more blurred as one travels along the MPP.

Consider an inhomogeneous medium with position-dependent scattering and absorption coefficients $b(\mathbf{x})$ and $a(\mathbf{x})$. Let $\tilde{\gamma}_p(s)$ be some arclength parameterized path from \mathbf{x}' to \mathbf{x} . The probability density for a photon to reach \mathbf{x} while traveling exactly along this path and not by any other possible one can be written as a product of two terms: the probability density of experiencing a series of scattering events which results in this particular path taken and the probability that the photon will “stay alive” at the end of its journey (i.e. not be absorbed along the path). We implicitly assume here that all absorption and scattering events are independent. This allows the second term to be expressed directly from the radiative transfer equation with no scattering which is written along the path in a trivial form $dL(s)/ds = -a(\mathbf{x})L(s)$ with initial conditions $L(0) = 1$ where we use loose notations $L(s)$ for the radiance along the path. In this case, the

fraction of initial radiance which reaches the end of the path is exactly the probability of a photon not being absorbed. Therefore, using the solution of the equation above we can write

$$p(\text{not absorbed}) = \exp\left(-\int_0^S a(\tilde{\gamma}_p(s')) ds'\right),$$

where the argument to the integral is the absorption coefficient along the path and S is total length of the path.

To deal with the scattering term, we adopt an approach similar to Wilson and Wang⁴⁶ which is inspired by physics, rather than attempting to follow the more mathematically rigorous treatment of Tessendorf⁴². Homogeneous media have been considered so far in the literature and in many cases the rigorous mathematical procedures will break down if scattering/absorption coefficients are allowed to vary across the medium.[†] A notable exception is the path construction using random walks by Pauly *et al.*²⁷. We first split the path into a number of straight line segments connecting positions of scattering events that change photon propagation direction. Then the probability density of a particular path is a product of probability densities that individual scattering events will change the propagation direction “just right” to steer the photon all the way along the path. Using the phase function $P(\Delta\theta)$ where $\Delta\theta$ is the change of propagation direction, we get for the total probability density by writing :

$$p(\text{path shape})\mathcal{D}x = \prod_{i \in \tilde{\gamma}_p(s)} P(\Delta\theta_i) d\omega_i, \quad (5)$$

where individual factors correspond to the different (i^{th}) scattering event along the path. Differential solid angles will eventually become a part of the full differential measure $\mathcal{D}x$ in the path space and are not of interest for finding the MPP since they do not affect relative probabilities of different paths. We now make further assumption that the phase function is strongly peaked in the forward direction which is true for many important media⁵. In this case, $P(\theta)$ can be approximated with the first terms of its Taylor expansion as $1 - \alpha\theta^2$ (we drop irrelevant constants here). It can be shown that if we want to maintain the mean cosine of the scattering angle g , α has to be set to $\alpha = 1/4(1 - g) = 1/(2\langle\theta^2\rangle)$ where $\langle\theta^2\rangle$ is the mean square scattering angle. Note that although the phase function is strongly forward peaked, this does not mean that the path itself has to deviate by only a small angle from its original propagation direction, which is an assumption often used to simplify derivations. We would also like to treat path $\tilde{\gamma}_p(s)$ as a continuous object by taking a limit in equation 5 as scattering events occur often enough along the path. Each scattering event changes the propagation direction by a small amount $\delta\theta$. In the case of forward peaked phase functions, only such scattering events generally have

[†] For example, the Fourier transformed RT equation will have a much more complex form in this case, containing convolution over frequencies.

significant probability. The expression of interest is

$$\prod_{i \in \tilde{\gamma}_p(s)} P(\Delta\theta_i) \approx \prod_{i \in \tilde{\gamma}_p(s)} (1 - \alpha(\delta\theta_i)^2) = \prod_{i \in \tilde{\gamma}_p(s)} \left(1 - \alpha \delta s_i \left[\left(\frac{\delta\theta_i}{\delta s_i} \right)^2 \delta s_i \right] \right), \quad (6)$$

where we introduced the lengths of path elements between scattering events δs_i . Note that by construction these segments are physically constrained to be of *finite* length of the order of $1/b(\tilde{\gamma}_p(s_i))$ and can not be simply treated as infinitely short in the limit since no scattering can physically occur on an infinitesimal path interval. We also intentionally re-arranged the last expression to highlight its part in square brackets which indeed can be treated as a full differential of some function in our approximation. Taking a logarithm of equation 6, using Taylor series, and replacing one of the δs_i with its physical value we obtain

$$\sum_{i \in \tilde{\gamma}_p(s)} \ln \left(1 - \alpha \delta s_i \left[\left(\frac{\delta\theta_i}{\delta s_i} \right)^2 \delta s_i \right] \right) \approx \sum_{i \in \tilde{\gamma}_p(s)} \frac{-\alpha}{b(\tilde{\gamma}_p(s_i))} \left[\left(\frac{\delta\theta_i}{\delta s_i} \right)^2 \delta s_i \right].$$

If the scattering scale $\delta s = 1/b$ is much smaller than the macroscopic scale of the path, the expression in square brackets can be taken to give $|d\theta/ds|^2 ds$ in this limit and the discrete sum will become an integral along the path. Taking the exponent, we get:

$$p(\text{path shape}) \sim \exp \left(- \int_0^s \frac{\alpha}{b(\tilde{\gamma}_p(s'))} \left| \frac{d\theta}{ds'} \right|^2 ds' \right).$$

We can write this because of a general property of function limits: if the limit exists, its value does not depend on the specific way to take the limit and therefore the particular subdivision of the path we use does not affect the final result. It is comforting to note this expression gives exactly the result of rigorous treatment of Tessendorf³⁹ when applied to a homogeneous medium. The full expression for path probability density is now

$$p(\text{path}) \sim \exp \left(- \int_0^s \left[a(\tilde{\gamma}_p(s')) + \frac{\alpha}{b(\tilde{\gamma}_p(s'))} \left| \frac{d\theta}{ds'} \right|^2 \right] ds' \right) = \exp(-A(\text{path})), \quad (7)$$

where A is analogous to classical action along the path.

To find the MPP, we need to determine the path which minimizes the effective attenuation in equation 7. It is not possible to write an analytic expression for arbitrary functions $a(\mathbf{x})$ and $b(\mathbf{x})$, but some important general trends can nevertheless be established. Since the expression under the integral is non-negative, equation 7 favors shortest paths with low curvature $d\theta/ds$. For example, if we are interested in the MPP connecting two points without specifying any extra conditions, this will be the straight line connecting these points. If a path has to turn to satisfy boundary conditions, it

will tend to curve more in regions with high scattering coefficient $b(\mathbf{x})$. Finally, for a homogeneous medium, the MPP will be completely determined by the applied boundary conditions. Explicit expressions can be obtained in some cases using the standard Euler-Lagrange minimization procedure applied to the integral in equation 7⁴² and we will use such results below.

4.2. Spatial Variability

Optical properties in a scattering medium can vary arbitrarily spatially. A spatial variability in a medium can be measured by the number of scattering events that occur along the path $\tilde{\gamma}_p$:

$$\ell(s) = \int_0^s b(\tilde{\gamma}_p(s')) ds'.$$

In uniform media, the spatial variability is just a constant multiple of the distance $s: \ell = bs$. Given the spatial variability ℓ of a path in inhomogeneous media, we can “invert” this equation and write displacement of the ray from its origin \mathbf{x}' with respect to ℓ as:

$$\mathbf{d}(\ell) = \int_0^\ell \frac{1}{b(\mathbf{x}' + \mathbf{d}(\ell'))} \vec{\beta}(\ell') d\ell'. \quad (8)$$

This expression suggests a practical way of constructing actual path in inhomogeneous medium by stitching together straight line segments with lengths given by the local scattering coefficient. The only extra information we need is local propagation direction which is discussed in the next subsection.

4.3. Finding the MPP

Tessendorf⁴³ described the propagation direction $\vec{\beta}(\ell)$ with Euler rotation angles and satisfying boundary conditions (equation 2) using the Fourier series expansion of the angles. We follow a simplified version of his formulation to construct the *stationary path* $\vec{\beta}_0$ which is the path that minimizes attenuation along its length. We mentioned before that due to reciprocity we can construct the path by reversing the direction of light propagation¹. Through the rest of the paper we take advantage of this property and construct the MPP starting from the initial (viewing) direction $\vec{\omega}$ and ending in the final (light) direction $\vec{\omega}'$, although the light is actually moving in the opposite direction.

Let R be a rotation matrix that rotates initial direction $\vec{\omega}$ to the z-axis vector $\vec{z} = (0 \ 0 \ 1)^T$. If the final direction $\vec{\omega}'$ is written in spherical coordinates as

$$\vec{\omega}' = \begin{pmatrix} \sin(\theta) \cos(\phi) \\ \sin(\theta) \sin(\phi) \\ \cos(\theta) \end{pmatrix},$$

then the stationary path $\vec{\beta}_0(s, \vec{z}, \vec{\omega}')$ that uniformly rotates $\vec{\omega}'$




Figure 3: Construction of the Most Probable Path in a homogeneous medium. The path is constructed by stitching together piecewise linear path segments and integrating the propagation direction $\vec{\beta}$. The propagation direction of path segments between starting point \mathbf{x} and ending point \mathbf{x}' uniformly rotates from the initial direction $\vec{\omega}$ to the final direction $\vec{\omega}'$, therefore matching the boundary conditions.

into \vec{z} over the physical path length S is

$$\vec{\beta}_0(s, \vec{z}, \vec{\omega}') = \begin{pmatrix} \sin(\theta(s)) \cos(\phi) \\ \sin(\theta(s)) \sin(\phi) \\ \cos(\theta(s)) \end{pmatrix}, \quad (9)$$

where

$$\theta(s) = \theta - \left(\frac{s}{S}\right) \theta.$$

Note that we still need to apply rotation R^{-1} that rotates \vec{z} back to $\vec{\omega}$ such that:

$$\vec{\beta}_0(s, \vec{\omega}, \vec{\omega}') = R^{-1} \vec{\beta}_0(s, \vec{z}, R\vec{\omega}'). \quad (10)$$

One can show that such a “uniformly turning” path is exactly the MPP among all paths of fixed length for the homogeneous medium, i.e. it minimizes effective attenuation given by equation 7.

If the stationary path $\vec{\beta}_0$ needs to be constructed in inhomogeneous medium with respect to path spatial variability ℓ then we will simply replace $\theta(s)$ in equation 9 by $\theta(\ell')$ where the light path $\vec{\gamma}_p$ is now parameterized according to the number of scattering events and not the physical path distance:

$$\theta(\ell') = \theta - \left(\frac{\ell'}{\ell}\right) \theta. \quad (11)$$

Given the starting point \mathbf{x} , initial direction $\vec{\omega}$, and final

direction $\vec{\omega}'$, the MPP is constructed by integrating the velocity function $\vec{\beta}_0$. The locations on the MPP are found in terms of displacements $\mathbf{d}(\ell)$ (equation 8) from the starting point \mathbf{x} . The displacements $\mathbf{d}(\ell)$ are defined implicitly and at first it appears that the spatial variability ℓ along the entire path is needed. Note, however, that each displacement $\mathbf{d}(\ell)$ along the path only depends on the value of ℓ up to this point and not on the spatial variability of the entire path, allowing an “incremental” construction of the path.

The path is constructed by stitching together piecewise linear path segments. We march along the path in steps of size $d\ell$ (step size in spatial variability not in distance). This step size $d\ell$ is arbitrary and is analogous to selecting a step size ds for direct lighting computation when marching along straight line. A sensible value for $d\ell$ can also be estimated from the optical properties and density of the volume. At every step we first obtain the propagation direction $\vec{\beta}(\ell)$ along the existing portion of the path using equations 11 and 10 with the accumulated total spatial variability ℓ_{acc} on the path so far substituted for ℓ . We then use equation 8 to compute the displacement point $\mathbf{d}(\ell_{acc})$ along the path. Finding the next point on the path involves increasing the spatial variability of the path so far by $d\ell$ ($\ell_{acc} = \ell_{acc} + d\ell$) and finding corresponding displacement $\mathbf{d}(\ell_{acc})$ from the initial point \mathbf{x} . Note that the full path is reconstructed from scratch at each step. So, for every sampling point on the curved path, the MPP is reconstructed from the initial point \mathbf{x} and not just from the previous sampling point on the path.

Note that at the starting point \mathbf{x} the spatial variability ℓ is zero (no scattering events encountered so far) which therefore causes the first path segment to be in the initial direction $\vec{\omega}$. Similarly, the very last segment of the path will by construction be in the direction $\vec{\omega}'$, matching the boundary conditions. The propagation direction of path segments between starting point \mathbf{x} and ending point \mathbf{x}' uniformly (in ℓ) rotates from the initial direction $\vec{\omega}$ to the final direction $\vec{\omega}'$. The segment length in physical space depends on the scattering coefficient b at previous displacement point $\mathbf{d}(\ell)$.

The total spatial variability ℓ along the path is just the sum of spatial variabilities $d\ell$ along each segment on the path. This is again analogous to computing the physical length s by summing segments ds along the straight line. As expected, the stationary path $\vec{\beta}_0$ is relatively flat in regions where scattering coefficient $b(\mathbf{x})$ is small and highly curved where density is high. Figure 3 illustrates construction of the most important path using the described method. Quaternions could provide an alternative and more rigorous approach to uniform rotations.

4.4. Multiple Scattering Phase Function

Tessendorf and Wasson⁴⁴ observe that the width of the phase function grows with the number of scattering events ℓ . When the number of scattering events ℓ grows large, the probabil-

ity of scattering in any direction is equal and the phase function essentially becomes isotropic. It follows from the WKB approximation that the average scattering angle Θ_{MS} after ℓ scattering events is:

$$\langle \Theta_{MS} \rangle = \frac{\ell}{1 - \exp(-\ell)}.$$

Tessendorf and Wasson⁴⁴ also introduce the idea of the multiply-scattered phase function which is defined as the probability of light scattering through an angle θ after ℓ scattering events:

$$P_{MS}(\theta, \ell) = \frac{1}{N} P \left(\theta \sqrt{\frac{\ell}{1 - \exp(-\ell)}} \right), \quad (12)$$

where N is the normalization constant such that $\int_{4\pi} P_{MS}(\theta, \ell) d\omega = 1$. Intuitively, equation 12 says that the probability of scattering into an arbitrary angle increases with the number of scattering events. When the number of scattering events is large, the phase function P_{MS} becomes isotropic. Note that equation 12 holds for arbitrary phase function.

4.5. WKB Approximation For Multiple Scattering

Tessendorf⁴¹ derived a path integral expression for the propagator G in homogeneous materials:

$$G(s, \mathbf{x}, \vec{\omega}, \mathbf{x}', \vec{\omega}') = \int [d\vec{\beta}] [dp] \delta(\vec{\beta}(0) - \vec{\omega}') \delta(\vec{\beta}(s) - \vec{\omega}) \delta \left(\mathbf{x} - \int_0^s ds' \vec{\beta}(s') \right) \exp(-cs) \exp \left(b \int_0^s ds' Z(\vec{p}(s')) \right) \exp \left(i \int_0^s ds' \vec{p}(s') \cdot \frac{d\vec{\beta}(s')}{ds'} \right), \quad (13)$$

where $Z(\vec{p})$ is the Fourier transform of the normalized Gaussian phase function P_G and \vec{p} is the Fourier transform variable. Interested readers are referred to Tessendorf^{42, 43} for detailed derivations of the path integral formulation. If the phase function P is not spatially varying and the single scattering albedo $\omega_0 = \frac{b}{c}$ is smoothly varying in space then the propagator G can be extended to inhomogeneous materials⁴⁴, but no details are provided. The WKB approximation done by Tessendorf^{42, 43} consists of approximating the path integral in equation 13 by finding the most probable (stationary) path and integrating all paths fluctuating around the MPP. The path integration approximation they obtained can be expressed in terms of the Green propagator G :

$$G(s, \mathbf{x}, \vec{\omega}, \mathbf{x}', \vec{\omega}') \sim \exp \left\{ - \int_0^s ds' c(\vec{\gamma}_p(s')) \right\} (e^\ell - 1) P_{MS}(\theta, \ell) \quad (14)$$

The propagator in equation 14 is valid for spatially varying materials as long as the density of the material is smoothly varying and the phase function is uniform within the volume. While the WKB approximation was obtained using the assumption of strongly forward peaked phase functions, there is no other restriction on the specific phase function

form. In practice, the direct and indirect radiance components are computed separately and if the phase function allows backscattering, it will be computed explicitly in the direct radiance component computation.

5. Algorithm

As discussed in section 3, the radiance L received from direction $\vec{\omega}$ at the observation point \mathbf{x} is composed of three components:

$$L(\mathbf{x}, \vec{\omega}) = L_{un}(\mathbf{x}, \vec{\omega}) + L_{sc}(\mathbf{x}, \vec{\omega}) + L_e(\mathbf{x}, \vec{\omega}).$$

The unscattered component $L_{un}^i(\mathbf{x}, \vec{\omega})$ represents the amount of unscattered light due to the i^{th} light source:

$$L_{un}^i(\mathbf{x}, \vec{\omega}) = L_{light}^i(\mathbf{x}, \vec{\omega}) \exp \left(- \int_0^\infty c(\mathbf{x} - \vec{\omega}s) ds \right), \quad (15)$$

where $L_{light}^i(\mathbf{x}, \vec{\omega})$ is the radiant exitance in direction $\vec{\omega}$ from the i^{th} light source. In practice, the unscattered radiance L_{un} and the emitted radiance L_e are also the source for the scattered radiance¹³:

$$Q_S(\mathbf{x}, \vec{\omega}) = a(\mathbf{x}) L_e(\mathbf{x}, \vec{\omega}) + b(\mathbf{x}) \sum_I^{allNlights} P(\vec{\omega}, \vec{\omega}_I) L_{un}^i(\mathbf{x}, \vec{\omega}_I). \quad (16)$$

To compute the total radiance L in the medium, all external and internal sources of radiation need to be propagated through the volume to the point \mathbf{x} on the volume that the camera is looking at. The evolution operator G from equation 4 propagates all energy to a given observation point and direction. We use the propagator G from equation 14 for the rendering algorithm described in the next subsection. We also use results from Tessendorf and Wasson⁴⁴ to develop a rendering algorithm using the path integration formulation.

Monte Carlo Ray Tracing

Monte Carlo ray tracing is an accurate algorithm for solving the radiative transfer equation in arbitrary media. We use it for comparison to evaluate our approximations. We march through the medium in direction $\vec{\omega}$ sampling points along the ray. The light from previous step is attenuated and the light that is inscattered into the viewing direction $\vec{\omega}$ is gathered. The inscattered light is collected recursively for each inscattered ray:

$$L_{n+1}(\mathbf{x}, \vec{\omega}) = \sum_I^{allNlights} L_{un}(\mathbf{x}, \vec{\omega}_I) P(\mathbf{x}, \vec{\omega}_I, \vec{\omega}) b(\mathbf{x}) \Delta x + \frac{4\pi}{N} \sum_{i=1}^M L_{sc}(\mathbf{x}, \vec{\omega}_i) P(\mathbf{x}, \vec{\omega}_i, \vec{\omega}) b(\mathbf{x}) \Delta x + \exp(-c(\mathbf{x}) \Delta x) L_n(\mathbf{x} + \vec{\omega} \Delta x, \vec{\omega})$$

where M is the number of directional samples taken. While Monte Carlo ray tracing is robust and powerful, it is also slow because of the large number of rays needed.

5.1. Path Integration Approximation

Our algorithm exploits the WKB approximation from previous section. The WKB approximation computes the multiply-scattered light by finding the most probable path and then analytically integrates scattered radiance along this path and some neighborhood around this path including quadratic fluctuations around this path (equation 14).

The approximate radiance is then the sum of the singly scattered (“direct”) and multiply scattered components⁴⁴:

$$L_{ssc}(\mathbf{x}, \vec{\omega}) = \int_0^{\infty} Q_S(\mathbf{x} - s\vec{\omega}, \vec{\omega}) \exp\left(-\int_0^s ds' c(\mathbf{x} - (s-s')\vec{\omega})\right) ds, \quad (17)$$

and

$$L_{msc}(\mathbf{x}, \vec{\omega}) = \sum_{i=1}^{n_w} \int_{4\pi} d\omega'' \int_0^{\infty} ds_i \{Q_S(\mathbf{x}'', \vec{\omega}'') \times w[\mathbf{x}''] \times P_{MS}(\theta, \ell)\}, \quad (18)$$

where

$$w[\mathbf{x}''] = \exp(-c(\mathbf{x}'')\ell/b(\mathbf{x}'')) \times (\exp(\ell) - 1), \quad (19)$$

and \mathbf{x}'' is an intermediate point on the curved path: $\mathbf{x}'' = \mathbf{x} + \mathbf{d}(\ell)$. The first term in equation 19 is the intensity of the ray at a sampling point along the path while the second term is the number of rays that propagate along this path. Note that the number of rays grows exponentially with the distance. Brute force Monte Carlo ray tracing algorithms would have trouble with this due to enormous number of potentially spawned rays. The path integration formulation does not have problems with this exponential growth, because the rays are dealt with via analytic integration.

As we saw earlier, the MPP we construct is not uniquely defined by the boundary conditions. The total length of the path (either arclength S or total spatial variability ℓ) remains a free parameter since it is not possible to find this parameter without preprocessing or involved analysis. To find the true MPP one would need to perform a search among paths with different lengths and choose the one with minimal effective attenuation. This process can be very expensive for inhomogeneous media. Fortunately, in practice it is sufficient to construct a small number n_w of “quasi-MPPs” with different total lengths (spatial variabilities) using the process described above and sum their contributions, as shown by equation 18. Based on the density and optical properties of the volume, spatial variability of the path is estimated. For example, a raymarching can be performed along three axes and diagonals to estimate the total length of the path parameter. These estimates are then used as a free parameter for the MPP generation. Further investigation is needed to estimate the total path length more robustly and systematically. However, we found that naïve approach worked well in practice without spending extra time probing the volume.

The contributions of diffuse (or indirect) outside light

Algorithm 1 Rendering Algorithm Using WKB Approximation

```

// Preprocessing for efficiency purposes
for each light source  $S$ 
    Compute a light attenuation through the media by ray marching
    Store light attenuation in a spatial data structure (e.g. octree, deep shadow map)
end for

// Unscattered (direct) radiance  $L_{ssc}$  computation
Given starting point  $\mathbf{x}$  and viewing direction  $\vec{\omega}$ 
for each sample point along the ray in direction  $\vec{\omega}$ 
    Compute the unscattered radiance contribution from equation 17
end for

// Scattered (indirect) radiance  $L_{msc}$  computation
// using WKB approximation
Choose number of important paths  $n_w$ 
Choose the sampling length  $d\ell$ 
for each important path  $\vec{\beta}_{n_w}$ 
    for each sample point along the path  $\vec{\beta}_{n_w}$ 
        Find  $ds$  and a point  $\mathbf{x}'$  on the path from equation 8
        for each set of sampling directions  $\Omega$ 
            Compute the scattered radiance contribution from equation 18
        end for
    end for
    Add diffuse light source contributions (equation 20)
end for
    
```

sources are included as a radiance source at the boundary:

$$L(\mathbf{x}, \vec{\omega}) = \exp\left\{-\int_0^{S_{boundary}} ds' c(\mathbf{x} - \vec{\omega}(s-s'))\right\} L_d(\vec{\omega}) + \sum_{n_w} \int_{4\pi} d\omega'' w_b[\mathbf{x}'] P_{MS}(\theta, \ell_b) L_d(\vec{\omega}''), \quad (20)$$

where L_d is the diffuse radiance in a given direction. Similarly to the weight w in equation 19 for the collimated beam¹³, the weight for the diffuse light source contribution is:

$$w_b[\mathbf{x}'] = \exp(-c\ell_b/b) \times (\exp(\ell_b) - 1), \quad (21)$$

and ℓ_b is the number of scattering events when the path reached the boundary \mathbf{x}' of the volume.

The rendering algorithm can be implemented using an existing ray marching routine with little difficulty. For every viewing ray $\vec{\omega}$, we first find an intersection with the volume. This is starting point \mathbf{x} . The direct component (equation 17) is computed using standard ray marching. At every sample point, the source term (direct radiance and emission) Q_S is computed and the intensity is then exponentially attenuated. To speed up computation we precompute light attenuation in the volume and store it at all possible depths. For this we use a simplified version of a deep shadow map²³. From each light source we compute attenuation along the ray and store it at every sampling point. This speeds up rendering, because we

only pay the price once and the rest is a table lookup. Therefore, the source term Q_S is just a lookup in the spatial data structure storing the precomputed attenuation for every light source. The indirect component (equation 18) is computed by constructing the most probable path (or some number of them) using the procedure presented in Section 4.3. We march along the curved path in steps such that the number of scattering events is uniform. At every sampling point along the curved path, we compute the contribution from many directions over the entire sphere. Sampling directions are chosen based on viewing and lighting directions (i.e. sample more densely in directions close to the viewing and lighting directions). We have a fixed number of directions we sample at every point. One can use a more sophisticated metric to determine how many directions to sample. In our implementation, the number of sampling directions is a parameter to the rendering algorithm. Although many directions are sampled at every sampling point along the path, this consists just of a few operations since no rays are actually spawned. This is unlike in a traditional Monte Carlo algorithm where additional rays are created at every sampling point which causes exponential growth in number of rays. Similarly to computing the direct component, the source term Q_S can be just looked up in the spatial data structure. Once the boundary \mathbf{x}' of the volume is reached, the diffuse light contributions are explicitly added (equation 20). The rendering algorithm is summarized in Algorithm 1.

6. Results and Discussion

The algorithm from the previous section was implemented in C++ (compiled with g++ compiler) as an extension to the traditional Monte Carlo raytracer. All examples were rendered on a Pentium IV 1.7 Ghz with 512 Mb of memory running Linux. The results of the algorithm for a uniform density volume is shown in Figure 4. These images used the scattering parameters for skim milk from Jensen *et al.*¹⁸ and Henyey-Greenstein phase function with $g = 0.9$. The images are 512 by 384 and used 9 samples per pixel. The runtimes were 25 minutes for the Monte Carlo method and 8 minutes for the PI method. Note that the path integration solution is slightly darker than the brute-force Monte Carlo solution. This is because path integration assumes that paths near the MPP are where all of the light comes from, ignoring contributions that are far from the MPP. However, the degree of darkening is small, showing that the contributions from those paths is small as well.

Figure 5 shows the results for a non-uniform density. The images are 512 by 384 and used 9 samples per pixel. The runtimes were 55 minutes for the PI method and 580 minutes for the Monte Carlo solution. The cloud densities were generated using a Perlin-style noise function, and a Gaussian phase function was used with $\mu = 0.11$. Again, notice some slight darkening.

Figure 6 shows the results for a non-uniform density cloud




Figure 4: Upper: MC raytracing. Lower: PI Approximation.

with different illumination that includes indirect illumination and skylight contribution. The image is 640 by 480 pixels and used 9 samples per pixel. The runtime was 95 minutes.

Most of the theory in this paper is known to other scientific communities. We have described a more intuitive way of constructing path integrals avoiding the heavy mathematics of functional integration that would ordinarily be required. We have shown that this theory can be applied to produce reasonably accurate images. This was accomplished by finding the MPP and explicitly computing contributions along that path. All other paths are dealt with via analytic integration around the MPP. The MPP can be also used in traditional Monte Carlo simulations as a form of variance reduction. Our formulation can be used as a starting point for a sophisticated method such as Metropolis Light Transport, which effectively attempts to find the MPP by random walk. We have not pursued a highly efficient implementation; there is much room for the development of algorithms based on path integration. For example, a hierarchical implementation analogous to¹⁵ would surely be beneficial. Our images should be viewed as a proof-of-concept that path integration is a valuable tool for dealing with volumetric lighting. We have not shown that it has advantages over the dipole approximation¹⁸ to the diffusion approxi-




Figure 5: Upper: MC raytracing. Lower: PI Approximation.



Figure 6: A cloud rendered using the Path Integration approximation. Diffuse light source contributions were added at the boundary of the volume.

mation for dense uniform media, and in fact there is reason to prefer diffusion as density increases because path integration ignores backscattering. However, we have shown that path integration can be applied to sparse and spatially-varying volumes where the diffusion approximation is not appropriate¹⁹. Thus path integration should be viewed as a

complement to diffusion, and should become part of the toolbox of dealing with volume-based scatterers. Many classical rendering problems could be reformulated and solved faster using the path integral formulation. Understanding the parallels with the classical formulation and thinking of problems in this framework opens many directions for future research.

7. Acknowledgements

Thanks to Milan Ikits, Emil Praun, and David Weinstein for reading drafts of the paper. The reviewers provided valuable suggestions for improving this paper. This work was partially supported by NSF grant 99-78099.

References

1. Raphael Aronson. Radiative transfer implies a modified reciprocity relation. *J. Opt. Soc. Am. A*, 14(2):486–490, February 1997.
2. James Arvo. Transfer equations in global illumination. Global Illumination, SIGGRAPH '93 Course Notes, August 1993.
3. G. I. Bell and S. Glasstone. *Nuclear Reactor Theory*. Van Nostrand Reinhold, New York, 1970.
4. Craig F. Bohren. *Clouds in a glass of beer*. John Wiley and Sons, 1987.
5. Craig F. Bohren. Multiple scattering of light and some of its observable consequences. *American Journal of Physics*, 55(6):524–533, June 1987.
6. Masud Chaichian and Andrei Demichev. *Path Integrals in Physics Volume I. Stochastic Processes and Quantum Mechanics*. Institute of Physics Publishing, 2001.
7. C. C. Constantinou and C. Demetrescu. Physical interpretation of virtual rays. *J. Opt. Soc. Am. A*, 14(8):1799–1803, 1997.
8. J. Dorsey, A. Edelman, H. W. Jensen, J. Legakis, and H. Pedersen. Modeling and rendering of weathered stone. In *Proceedings of SIGGRAPH 1999*, pages 225–234, August 1999.
9. Ronald Fedkiw, Jos Stam, and Henrik Wann Jensen. Visual simulation of smoke. In *Proceedings of ACM SIGGRAPH 2001*, pages 15–22, August 2001.
10. R. P. Feynman and A. R. Hibbs. *Quantum mechanics and path integrals*. McGraw-Hill Higher Education, 1965.
11. R. G. Giovanelli. *Diffusion through non-uniform media*. Progress in Optics. North-Holland, Amsterdam, 1963.
12. Eugene P. Gross. Path integral approach to multiple scattering. *J. Math. Phys.*, 24(2):399–405, 1983.
13. Akira Ishimaru. *Wave Propagation and Scattering in Random Media, Volume I: Single scattering and transport theory; Volume II: Multiple scattering, turbulence, rough surfaces and remote sensing*. Academic Press, New York, 1978.
14. Steven L. Jacques and Xuijing Wang. Path integral description of light transport in tissues. In *Proceedings of optical tomography and spectroscopy of tissue: theory, instrumentation, and human studies II*, pages 488–499, 1997.

15. Henrik Wann Jensen and Juan Buhler. A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics*, 21(3):576–581, July 2002.
16. Henrik Wann Jensen and Per H. Christensen. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proceedings of SIGGRAPH 98*, pages 311–320, July 1998.
17. Henrik Wann Jensen, Justin Legakis, and Julie Dorsey. Rendering of wet materials. In *Eurographics Rendering Workshop 1999*, Granada, Spain, June 1999. Springer Wein / Eurographics.
18. Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. In *Proceedings of SIGGRAPH 2001*, pages 511–518, August 2001.
19. Jan J. Koenderink and Andrea J. van Doorn. Shading in the case of translucent objects. In Bernice E. Rogowitz and Thrasyloulos N. Pappas, editors, *Proc. SPIE Vol. 4299 Human Vision and Electronic Imaging VI*, pages 312–320, 2001.
20. Eric Lafortune. *Mathematical Models and Monte Carlo Algorithms*. PhD thesis, Katholieke Universiteit Leuven, 1996.
21. Flor Langouche, Dirk Roekaerts, and Enrique Tirapegui. *Functional Integration and Semiclassical Expansions*. D. Reidel Publishing Company, 1982.
22. Hendrik P. A. Lensch, Michael Goesele, Phillippe Bekaert, Jan Kautz, Marcus A. Magnor, Jochen Lang, and Hans Peter Seidel. Interactive rendering of translucent objects. In *Proceedings of Pacific Graphics*, 2002.
23. Tom Lokovic and Eric Veach. Deep shadow maps. In *Proceedings of ACM SIGGRAPH 2000*, pages 385–392, July 2000.
24. Steven D. Miller. Stochastic construction of a feynman path integral representation for green's functions in radiative transfer. *Journal of Mathematical Physics*, 39(10):5307–5315, 1998.
25. S. G. Narasimhan and S. K. Nayar. Shedding light on the weather. In *Proceedings of CVPR*, June 2003.
26. Duc Quang Nguyen, Ronald P. Fedkiw, and Henrik Wann Jensen. Physically based modeling and animation of fire. *ACM Transactions on Graphics*, 21(3):721–728, July 2002.
27. Mark Pauly, Thomas Kollig, and Alexander Keller. Metropolis light transport for participating media. In *Rendering Techniques 2000: 11th Eurographics Workshop on Rendering*, pages 11–22, June 2000.
28. Lev T. Perelman, Jun Wu, Irving Itzkan, and Michael S. Feld. Photon migration in turbid media using path integrals. *Physical Review Letters*, 72(9):1341–1344, 1994.
29. Lev T. Perelman, Jun Wu, Yang Wang, Irving Itzkan, Ramachandra R. Dasari, and Michael S. Feld. Time-dependent photon migration using path integrals. *Physical Review E*, 6(51):6134–6141, 1995.
30. Frederic Pérez, Xavier Pueyo, and François X. Sillion. Global illumination techniques for the simulation of participating media. In *Eurographics Rendering Workshop 1997*, pages 309–320, June 1997.
31. Matt Pharr and Patrick M. Hanrahan. Monte carlo evaluation of non-linear scattering equations for subsurface reflection. In *Proceedings of SIGGRAPH 2000*, pages 75–84, July 2000.
32. A. Ya. Polishchuk and R. R. Alfano. Fermat photons in turbid media: an exact analytic solution for most favorable paths – a step toward optical tomography. *Optics Letters*, 20(19):1937–1939, 1995.
33. A. Ya. Polishchuk, J. Dolne, F. Liu, and R. R. Alfano. Average and most-probable photon paths in random media. *Optics Letters*, 22(7):430–432, 1997.
34. Rudolph W. Preisendorfer. *Light Transport on Discrete Spaces*. Pergamon Press, 1965.
35. S. Premože. Analytic Light Transport Approximations for Volumetric Materials. In *Proceeding of Pacific Graphics*, 2002.
36. Holly Rushmeier. Rendering participating media: Problems and solutions from application areas. In *Fifth Eurographics Workshop on Rendering*, pages 35–56, Darmstadt, Germany, June 1994.
37. Jos Stam. Multiple scattering as a diffusion process. In *Eurographics Rendering Workshop 1995*. Eurographics, June 1995.
38. V. I. Tatarskii, A. Ishimaru, and V. U. Zavorotny. *Wave Propagation in Random Media*. Society of Photo-optical Instrumentation Engineers, Bellingham, Washington, 1993.
39. Jerry Tessendorf. Radiative transfer as a sum over paths. *Phys. Rev. A*, 35:872–878, 1987.
40. Jerry Tessendorf. Comparison between data and small-angle approximations for the in-water solar radiance distribution. *Journal of Optical Society of America A*, 5:1410–1418, 1988.
41. Jerry Tessendorf. Time-dependent radiative transfer and pulse evolution. *J. Opt. Soc. Am. A*, (6):280–297, 1989.
42. Jerry Tessendorf. The Underwater Solar Light Field: Analytical Model from a WKB Evaluation. In Richard W. Spinrad, editor, *SPIE Underwater Imaging, Photography, and Visibility*, volume 1537, pages 10–20, 1991.
43. Jerry Tessendorf. Measures of temporal pulse stretching. In Gary D. Gilbert, editor, *SPIE Ocean Optics XI*, volume 1750, pages 407–418, 1992.
44. Jerry Tessendorf and David Wasson. Impact of multiple scattering on simulated infrared cloud scene images. In *Proceedings SPIE. Characterization and propagation of sources and backgrounds*, volume 2223, pages 462–473, 1994.
45. Eric Veach and Leonidas J. Guibas. Metropolis light transport. In *SIGGRAPH 97 Conference Proceedings*, pages 65–76, August 1997.
46. Michael J. Wilson and Ruikang K. Wang. A path-integral model of light scattered by turbid media. *J. Phys. B: At. Mol. Opt. Phys.*, 34:1453–1472, 2001.