

Importance Sampling and Adjoint Hybrid Methods in Monte Carlo Transport with Reflecting Boundaries

Guillaume Bal*, and Ian Langmore †

June 19, 2010

Abstract

Adjoint methods form a class of importance sampling methods that are used to accelerate Monte Carlo (MC) simulations of transport equations. Ideally, adjoint methods allow for zero-variance MC estimators provided that the solution to an adjoint transport equation is known. Hybrid methods aim at (i) approximately solving the adjoint transport equation with a deterministic method; and (ii) use the solution to construct an unbiased MC sampling algorithm with low variance. The problem with this approach is that *both* steps can be prohibitively expensive. In this paper, we simplify steps (i) and (ii) by calculating only parts of the adjoint solution. More specifically, in a geometry with limited volume scattering and complicated reflection at the boundary, we consider the situation where the adjoint solution “neglects” volume scattering, whereby significantly reducing the degrees of freedom in steps (i) and (ii). A main application for such a geometry is in atmospheric remote sensing. Volume scattering is then incorporated using an analog sampling algorithm (or more precisely a simple modification of analog sampling called a heuristic sampling algorithm) in order to obtain unbiased estimators. In geometries with weak volume scattering (with a domain of interest of size comparable to the transport mean free path), we demonstrate numerically significant variance reductions and speed-ups (figures of merit).

Keywords: Linear Transport; Monte Carlo; Hybrid Methods; Importance Sampling; Variance Reduction; Atmospheric Remote Sensing

1 Introduction

Forward and inverse linear transport models find applications in many areas of science including medical imaging and optical tomography [1], radiative transfer in the atmosphere and the ocean [4, 12, 14], neutron transport [6, 16], as well as the propagation of seismic waves in the earth crust [15]. In this paper, we focus on the solution of the forward transport problem by the Monte Carlo method with remote sensing (an inverse transport problem) of the atmosphere as our main application. Light is emitted from the sun and propagates in a complex environment involving

*Department of Applied Physics and Applied Mathematics, Columbia University, 200 S.W. Mudd building, 500 W. 120th street, New York NY, 10027; 212-854-4731, gb2030@columbia.edu

†Corresponding author. Department of Applied and Applied Mathematics, Columbia University, 200 S.W. Mudd building, 500 W. 120th street, New York NY, 10027; 415-272-6321, ianlangmore@gmail.com

absorption and scattering in the atmosphere and at the Earth’s surface before (a tiny fraction of) it reaches a detector, typically mounted on a plane or a satellite.

The transport equation may be solved numerically in a variety of ways. Monte Carlo (MC) simulations model the propagation of individual photons along their path and are well adapted to the complicated geometries encountered in remote sensing. Photons scatter and are absorbed with prescribed probability depending on the underlying medium. The output from the simulation, e.g., the fraction of photons that hit a detector, is the expected value of a well-chosen random variable. These simulations are very easy to code, embarrassingly parallel to run, and suffer no discretization error (in principle). The drawback is that they can be very slow. Monte Carlo methods converge at a rate $= (\text{Variance}N^{-1/2})$ where N is the number of simulations, and the variance is that of each shot fired. In remote sensing, the (relative) variance is high in large part because the detector is typically small and thus most photons are not recorded by the detector. In order to be effective, MC methods must be accelerated.

Most efforts to speed MC simulations focus on reducing the variance of each shot. See [16, 13] or the review of more recent work (on neutron transport) in [9]. See also [20] for a thorough introduction to the MC techniques in computer graphics. In problems with a small detector, this is achieved by directing photons toward that detector (and re-weighting to keep calculations unbiased). When *survival biasing* is used, photons have their weight decreased rather than being absorbed [16, 13]. Often, one uses some heuristic (such as proximity to the detector), or some function to measure the “importance” of each region of phase space. Splitting methods [16, 13], upon identifying that a photon is in a region of high importance, split the photon into two or more photons. The weight of each photon is then decreased. Propagating many photons with a low weight is not desirable, therefore splitting is often accompanied by *Russian roulette*. Here, if a photon enters a region of low enough importance, then the photon is killed off with a certain probability (high chance of death if the weight is low). Typically a *weight window* is used to enforce regions of low/high importance. *Source biasing* techniques change the source distribution in order to more effectively reach the detector. More generally, the absorption/scattering properties at any point can be modified, provided shots are re-weighted correctly.

It has long been recognized that the adjoint transport solution is a natural and optimal importance function [16, 13, 17, 18, 9, 19, 7]. One can use well-chosen approximations of the adjoint solution (typically a rough deterministic solution) to reduce variance. The result is a *hybrid* method (deterministic+MC). The AVATAR method uses an adjoint approximation to determine weight windows [19]. The CADIS scheme in [9] uses an adjoint approximation in both source biasing and weight-window determination. An adaptive technique that successively refines the solution in “important” regions (and uses the adjoint to designate such regions) is described in [10, 11]. In [16, 17], a zero-variance technique is outlined that uses the true adjoint solution to fire photons that all reach the detector with the same weight (which happens to be the correct answer). This method is impractical since determining the exact adjoint solution everywhere is harder than determining some integral of that solution. The LIFT method [17, 18] uses an approximation of the adjoint solution to approximate this zero-variance method.

We adapt the zero-variance technique to the particular problem we have at hand; see figure 1 for the type of geometry considered in this paper. The problem we consider has a fixed, reflective, complex boundary, and relatively large mean-

free-path (MFP) in the sense that a large fraction of the photons reaching the detector have not scattered inside the atmosphere. The calculation of the approximate adjoint solution used to approximate zero-variance techniques is difficult and potentially very costly. What we demonstrate in this paper is that partial, “localized” (in an appropriate sense) knowledge of the adjoint solution still offers very significant variance reductions. More specifically, we calculate adjoint solutions that accurately account for the presence of the boundary but do not account for volume scattering (infinite MFP). The calculation of the adjoint solution thus becomes a radiosity problem with much reduced dimensions compared to the full transport problem. This, of course, can only reduce variance in proportion to the number of “ballistic” photons that never interact with the volume. Moreover, an adjoint solution that does not “see” volume scattering cannot be used alone as a variance reduction scheme for otherwise volume scattering would be neglected and the simulation biased, which is not allowed. When combined with simple rules for allowing volume scattering and sending some photons directly from the volume to the detector, our hybrid method yields very significant variance reduction at relatively minimal cost. Furthermore, the methodology studied is applicable whenever any method is available to deterministically pre-calculate flux over any subset of paths. For instance, complex propagation of light in clouds and its importance could be pre-calculated locally and incorporated into the MC simulations in a similar fashion. This modular approach to the description of the adjoint solution is well-adapted to the remote sensing geometry and avoids complicated, global, and hence expensive deterministic calculations of adjoint transport solutions. Our treatment of the reflecting boundary described in detail in this paper is a first step toward modular adjoint transport calculations and their variance reduction capabilities in remote sensing.

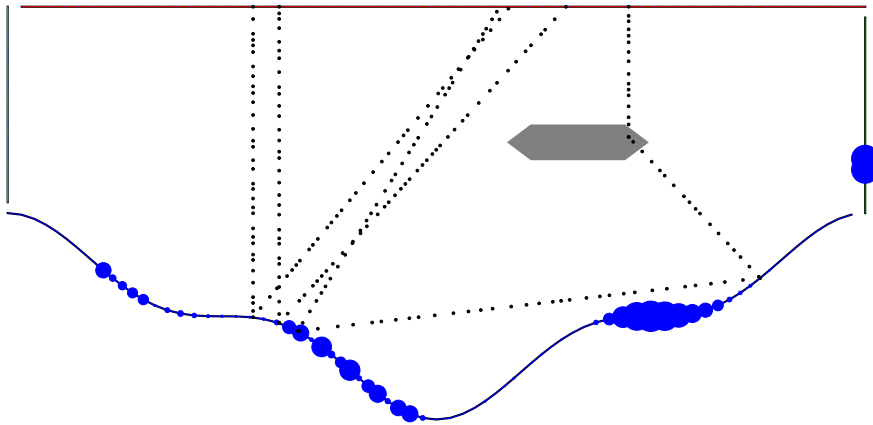


Figure 1: Mountain ($1 - \cos^3 x$ shape), cloud, sky, and detector. Dot size indicates relative adjoint flux strength. Large dots on right-hand-side are the detector (dot size is down-scaled for detector). Dot size on mountain indicates that portions of the mountain are shaded from the detector, and that the scattering albedo is non-constant.

The rest of the paper is structured as follows. Section 2 presents basic information about the transport equation with reflecting boundary. Section 3 presents our main theoretical results on hybrid acceleration of Monte Carlo by deterministic adjoint calculations. We adopt an importance sampling viewpoint [3] that

is common in the statistical literature. This means we view the modifications to absorption/scattering as a change of probability measure and the re-weighting as a Radon-Nikodym derivative (Jacobian). This allows us to fit many methods together under one framework. In particular, source-biasing, the zero-variance scheme, our approximation of it, and our “heuristic” volume-to-detector adjustment are put in this light. This allows us to obtain estimates of variance as a function of scattering/absorption coefficients and the accuracy of the deterministic solver.

Sections 3.1 and 3.2 recall the main ideas behind importance sampling and the use of adjoint transport solutions. We recall how zero-variance chains can be constructed and show how they can be approximated by small-variance chains. In the absence of volume scattering, a small variance chain is constructed in section 3.3. The modularity mentioned earlier in this section is implemented by a regularization methodology introduced in (41) in section 3.4.1. The Surface Adjoint Importance (SAI) method, used to incorporate the adjoint solutions that accurately describe the surface defined in section 3.3 in a scheme that also handles volume scattering, is described in detail in section 3.4. The variance reduction and speedup that can be gained from the proposed methodology are presented in section 4. Several details in the derivation and the proof of the results of section 3 and the numerical implementation of the simulations of section 4 are postponed to Appendix A.

2 Transport with Reflecting Boundaries

Let $X \subset \mathbb{R}^d$ ($d = 3$ in practice and $d = 2$ in our numerical simulations) be an open (spatial) domain with smooth boundary ∂X . Denote $X \cup \partial X$ by \bar{X} . For $x \in X$ photons will have velocities $v \in \mathbb{S}^{d-1}$, the unit sphere, and we call the pair $z = (x, v) \in Z$. When $x \in \partial X$ we separate directions into incoming and outgoing. With ν_x the outward unit normal vector at $x \in \partial X$ we have $\Gamma_{\pm} := \{(x, v) : x \in \partial X, \pm v \cdot \nu_x > 0\}$. Note that z always is interpreted as the pair (x, v) , and for example $z_j = (x_j, v_j)$.

Photons will be cast along rays, and travel until they hit the boundary. We define the forward and backward propagation times as $\tau_{\pm}(z) := \min_{t>0} \{x \pm tv \in \partial X\}$. We also define the forward and backward spatial and phase-space propagations $x_{\pm}(z) := x \pm \tau_{\pm}(z)v$, $z_{\pm}(z) := (x_{\pm}(z), v)$. The rays themselves are denoted by a starting point and direction, $r(z) := \{x + tv : 0 < t < \tau_{+}(z)\}$.

Define an integral over $\bar{Z} := Z \cup \Gamma_{-} \cup \Gamma_{+}$ by

$$\int_{\bar{Z}} f(z) dz := \int_Z f(z) dz + \int_{\mathbb{S}^{d-1}} \int_{\partial X} f(z) d\mu(x) dv,$$

where $d\mu$ the surface measure on ∂X and an inner product by

$$\langle f, g \rangle = \int_{\bar{Z}} f(z)g(z) dz.$$

Some functions are defined only, for example, on Γ_{-} . In that case we extend the function to \bar{Z} by setting it equal to zero off of Γ_{-} .

Light traveling through X encounters an absorption cross section $\sigma_a(x)$, scattering kernel $\theta(x, v \rightarrow v')$, and scattering cross section $\sigma_s(x) := \int_{\mathbb{S}^{d-1}} \theta(x, v \rightarrow v') dv'$, which is assumed independent of v . The total cross section $\sigma := \sigma_a + \sigma_s$. The

exponential of σ is denoted by

$$E_\sigma(x_1, x_2) := e^{-\int_0^{|x_1-x_2|} \sigma(x_1+tv_1) dt},$$

where $v_1 = \widehat{x_2 - x_1} := (x_2 - x_1)|x_2 - x_1|^{-1}$. We define $E_{\sigma_a}, E_{\sigma_s}$ similarly. Once a photon collides with the boundary, it is scattered with probability $\alpha(x)$. In that case, the probability distribution $\Theta(x, v \rightarrow v')$ determines the new direction. This implies

$$\int_{\nu_x \cdot v' > 0} \Theta(x, v \rightarrow v') dv' = 1.$$

We model photon flux density u in our medium with source s by

$$\begin{aligned} v \cdot \nabla_x u(z) + \sigma(x)u(z) &= Ku(z) \\ u|_{\Gamma_-}(z) &= \frac{K(u|_{\Gamma_+})(z)}{|\nu_x \cdot v|} + \frac{s(z)}{|\nu_x \cdot v|}, \end{aligned} \quad (1)$$

where

$$\begin{aligned} Kf(z) &= \int_{\mathbb{S}^{d-1}} \theta(x, v' \rightarrow v) f|_Z(x, v') dv', & z \in Z \\ Kf(z) &= \alpha(x) \int_{\nu_x \cdot v' > 0} \Theta(x, v' \rightarrow v) |\nu_x \cdot v'| f|_{\Gamma_+}(x, v') dv' & z \in \Gamma_-. \end{aligned} \quad (2)$$

Since the transport problem is linear, we normalize s so that

$$\int_{\Gamma_-} s(x, v) d\mu(x) dv = 1. \quad (3)$$

Multiplying the identity $v \cdot \nabla_x u(x - tv, v) + \sigma(x - tv)u(x - tv, v) = Ku(x - tv, v)$ by the integrating factor $E_\sigma(x, x - tv)$ and integrating t from 0 to $\tau_-(z)$ we find that u satisfies the following integral transport equation:

$$u = LKu + Ls, \quad \text{so that} \quad u = \sum_{n=0}^{\infty} (LK)^n Ls = L \sum_{n=0}^{\infty} (KL)^n s, \quad (4)$$

where (with $z \in Z \cup \Gamma_+$)

$$Lf(z) := \int_0^{\tau_-(z)} E_\sigma(x, x - tv) f|_Z(x - tv, v) dt + \frac{E_\sigma(x, x_-(z))}{|\nu_{x_-(z)} \cdot v|} f|_{\Gamma_-}(z_-(z)).$$

Then (4) motivates us to define ψ_o solving

$$\psi_o = KL\psi_o + s, \quad \text{so that} \quad \psi_o = \sum_{n=0}^{\infty} (KL)^n s \quad \text{and} \quad u = L\psi_o. \quad (5)$$

The decompositions (4) and (5) of the transport solution into components corresponding to increasing orders of scattering is standard in forward and inverse transport theory. We refer the reader to e.g. [2, 5, 16] for additional details.

2.1 Coefficient assumptions and measurement setup

The function $g(z) := g(x, v)$ describes the phase-space representation of the detector. We will see that the Monte Carlo detector is defined as $\bar{g}(z) := g(z)|\nu_x \cdot v|^{-1}$. We assume that the source/detector are nonzero only on the incoming/outgoing boundaries: $\text{supp}(s) \subset \Gamma_-$, $\text{supp}(g) \subset \Gamma_+$. Finally, we assume that the detector is non-scattering, $\alpha(x) = 0$ for $(x, v) \in \text{supp}(g)$. We also have $\alpha \equiv 0$ on the sky and left/right sides to model photons that escape our domain. These assumptions are satisfied for source radiation coming from the sun and detectors on high-elevation planes or satellites. The methodology we present could easily be adapted to detectors placed in the volume.

Our measurement is the phase space integral $\langle g, u \rangle$. All numerical methods employed will approximate this integral. When $g(z) = \nu_x \cdot v$ (for x on the support of the detector), the detector is measuring photon flux. This corresponds to counting Monte Carlo photons that pass through the support of \bar{g} .

2.2 Adjoint solutions and operator decomposition

We will see that it is the adjoint operator (and its kernel) that is needed to define the Markov chain transition kernels in MC simulations. We denote adjoint operators by $*$, and adjoint is defined with respect to the inner product $\langle \cdot, \cdot \rangle$. The methods used in this paper rely on a decomposition of the operator $(LK)^*$ into C^* (ray Casting) and S^* (Scattering) operators. We have

$$C^* f(z_1) := \int_0^{\tau_+(z_1)} E_\sigma(x_1, x_1 + tv_1) f|_Z(x_1 + tv_1, v_1) dt + E_\sigma(x_1, x_+(z_1)) f|_{\Gamma_+}(z_+(z_1)),$$

when $z_1 \in Z \cup \Gamma_-$, and

$$S^* f(z_1) := \begin{cases} \int_{\mathbb{S}^{d-1}} \theta(x_1, v_1 \rightarrow v_2) f|_Z(x_1, v_2) dv_2, & z_1 \in Z \\ \alpha(x_1) \int_{\nu_{x_1} \cdot v_2 < 0} \Theta(x_1, v_1 \rightarrow v_2) f|_{\Gamma_-}(x_1, v_2) dv_2, & z_1 \in \Gamma_+. \end{cases}$$

While $C^* \neq L^*$, we still have $C^* S^* = (KL)^*$, which implies of course that $SC = KL$. We also note that $C^* \bar{g} = L^* g$. The notation x_1, v_1, z_1 , and x_2, v_2, z_2 is suggestive of the fact that these variables will later represent photon positions/velocities at the first, second, third, etc. . . position.

Define the adjoint ψ_o^* by

$$\psi_o^* = C^* S^* \psi_o^* + C^* \bar{g}, \quad \text{so that} \quad \psi_o^* = \sum_{n=0}^{\infty} (C^* S^*)^n C^* \bar{g} = C^* \sum_{n=0}^{\infty} (S^* C^*)^n \bar{g}. \quad (6)$$

Then definitions of ψ_o, ψ_o^* imply $\langle s, \psi_o^* \rangle = \langle C^* \bar{g}, \psi_o \rangle$, and therefore

$$\langle s, \psi_o^* \rangle = \langle C^* \bar{g}, \psi_o \rangle = \langle L^* g, \psi_o \rangle = \langle u, g \rangle. \quad (7)$$

In other words, the adjoint solution $\psi_o^*(z)$ is a weight giving the ‘‘importance’’ of a source at point z on our measurement $\langle u, g \rangle$. This is the first fundamental reason for the use of the adjoint solution in Monte Carlo transport; see e.g. [16, 17] and also Theorem 3.1 below. Note that ψ_o^* can be shown to solve

$$\begin{aligned} -v \cdot \nabla_x \psi_o^* + \sigma \psi_o^* &= S^* \psi_o^*, \\ \psi_o^*|_{\Gamma_+} &= S^*(\psi_o^*|_{\Gamma_-}) + \bar{g}. \end{aligned} \quad (8)$$

The relation in (6) motivates us to define ψ_i^* solving

$$\psi_i^* = S^* C^* \psi_i^* + \bar{g}, \quad \text{so that} \quad \psi_i^* = \sum_{n=0}^{\infty} (S^* C^*)^n \bar{g}.$$

We also have the relations

$$\psi_o^* = C^* \psi_i^*, \quad \psi_i^* = S^* \psi_o^* + \bar{g}. \quad (9)$$

Both ψ_i^* and ψ_o^* appear naturally in Monte-Carlo transport. When constructing transition kernels (that determine casting/direction changes), one will be a normalization constant for the other (implicitly or explicitly). We make the distinction explicit due to the following heuristics. We may think of $\psi_i^*(z_1)$ as the *incoming importance* at z_1 . To it we associate an arrow directed into point x_1 with direction v_1 . $\psi_o^*(x_1, v_1)$ is the *outgoing importance* at (x_1, v_1) since it is the integral of incoming importance at all possible points (x_2, v_1) along the ray $x_1 + tv_1$. Likewise, away from the support of g , $\psi_i^* = S^* \psi_o^*$, meaning that the incoming importance at x in direction v is the integral of all importance exiting x .

It is important to note that all chains described here alternate casts with direction changes. Casts move particles from a point (x_1, v_1) to a point $(x_1 + tv_1, v_1)$ on a one-dimensional line segment while direction changes move particles from a point (x_2, v_1) to a point (x_2, v_2) on a $(d-1)$ -dimensional sphere. One could alternatively try devising a scheme that moves $z_j \rightarrow z_{j+1}$ directly. This significantly increases computational cost since, given z_1, z_2 may lie anywhere on the d dimensional manifold $\{x_1 + t_1 v_1 : 0 < t < \tau_+(z_1)\} \times \mathbb{S}^{d-1}$. Thus, sampling z_2 directly would require handling a d dimensional data structure rather than a 1 dimensional and a $(d-1)$ dimensional data structures for alternate casts and direction changes. This is our main motivation for introducing the operators C^* and S^* rather than directly working with $(LK)^*$.

2.3 Transport when $\sigma = 0$

When $\sigma \equiv 0$ (the ‘‘surface’’ regime), we have $C^* = C^s$ (with kernel k_{C^s}), $S^* = S^s$ (with kernel k_{S^s}) where

$$C^s f(z_1) := \int_0^{\tau_+(z_1)} f|_Z(x_1 + tv_1, v_1) dt + f|_{\Gamma_+}(z_+(z_1)),$$

$$S^s f(z_1) := \begin{cases} 0, & z_1 \in Z \\ \alpha(x_1) \int_{\nu_{x_1} \cdot v_2 < 0} \Theta(x_1, v_1 \rightarrow v_2) f|_{\Gamma_-}(x_1, v_2) dv_2, & z_1 \in \Gamma_+ \end{cases}$$

We then define ψ_i^s as the solution to

$$\psi_i^s = S^s C^s \psi_i^s + \bar{g},$$

and let $\psi_o^s := C^s \psi_i^s$. Since $\psi_i^s|_Z \equiv 0$,

$$\psi_o^s(z) = \psi_i^s(z_+(z)), \quad z \in Z \cup \Gamma_-,$$

and for $z \in \Gamma_+$,

$$S^s C^s \psi_i^s(z_1) = \alpha(x_1) \int_{\nu_{x_1} \cdot v_2 < 0} \Theta(x_1, v_1 \rightarrow v_2) \psi_i^s(z_+(x_1, v_2)) dv_2.$$

In other words, we can solve for ψ_i^s by paying attention only to the boundary, and then propagate it to compute ψ_o^s .

3 Monte Carlo with Reflecting Boundaries

Monte Carlo consists of simulating transport one photon at a time. Photons propagate along straight lines until they interact with the underlying medium, where they are either absorbed or scattered into another direction, or reach the detector where they are collected. It can be shown that photon paths terminate (with probability one) after finitely many collisions. Following [16], paths will be written $\omega = (z_0, \dots, z_{\tau-1}, (x_\tau, \mathfrak{d}))$. So the initial point $z_0 = (x_0, v_0) \in \bar{\mathcal{Z}}$, and subsequently we choose x_1 by casting a ray, then v_1 by changing direction, then x_2 , then v_2 and so on until absorption. At this stopping time τ , x_τ is chosen, and then v_τ is set equal to \mathfrak{d} , the “dead velocity.” The chain is now terminated. Let

$$\Omega := \{((x_0, v_0), \dots, (x_{\tau-1}, v_{\tau-1}), (x_\tau, \mathfrak{d})) : v_j = \widehat{x_{j+1} - x_j}\}.$$

All casts and direction changes (including “death”) are made by drawing random variables. We thus introduce a probability measure on the set of paths Ω . We note that $\{\omega \in \Omega : \tau(\omega) = n\} = \{\tau = n\}$ is the set of paths terminating after $n - 1$ scattering events.

A probability measure on Ω is a map P from the (measurable) subsets of Ω into $[0, 1]$. It corresponds to a method of choosing paths. Given a set $A \subset \Omega$ of possible paths, $P[A]$ is the probability that a path lies in A . $P[\tau = n] := P[\{\omega : \tau(\omega) = n\}]$ is the probability that the chain terminates at step n . Let D denote the paths that end up hitting the detector. Then $P[D]$ is the probability of hitting the detector. With the indicator function $\mathbb{1}_D(\omega)$ defined as $\mathbb{1}_D(\omega) = 1$ if $\omega \in D$ and zero otherwise, we have the notation

$$P[D] = \int_D dP(\omega) = \int_\Omega \mathbb{1}_D(\omega) dP(\omega) = \mathbb{E}\{\mathbb{1}_D\}.$$

Here, $\mathbb{E}\{\}$ denotes mathematical expectation (ensemble averaging) w.r.t. P .

3.1 Monte Carlo and Importance Sampling

The *analog* measure P^a closely follows the physics of photon propagation (at least one reasonable physical model for photon propagation). Monte Carlo simulations based on this measure have very large (relative) variance because most of the photons do not reach the detector. Several standard methods exist to modify the measure to steer more photons toward the detector in an unbiased way, i.e., in a way that does not modify the detector reading $\langle u, g \rangle$. We start with a presentation of the analog chain and then present the main ideas of importance sampling to reduce variance in MC simulations. We also present the (standard) survival biasing chain, which forms a basis for comparison and a component in our composite SAI chain.

3.1.1 Analog Sampling

We first define the analog transition kernels $k_{C^*}^a$ and $k_{S^*}^a$, associated to the operators C^* and S^* . The analog ray casting transition kernel is

$$k_{C^*}^a(z_1 \rightarrow x_2) := [\delta_{r(z_1)}(x_2)\sigma(x_2) + \delta(x_2 - x_+(z_1))] E_\sigma(x_1, x_2).$$

Above, $\delta_{r(z_1)}(x_2)$ is the “delta function” in \mathbb{R}^d concentrated along the ray $r(z_1)$. It forces x_2 to be along the path $x_1 + tv_1$, $t > 0$. $\delta(x_2 - x_+(z_1))$ forces x_2 to be

on the boundary at $x_+(z_1)$. Since

$$\frac{d}{dt}(1 - E_\sigma(x, x + tv)) = \sigma(x + tv)E_\sigma(x, x + tv), \quad (10)$$

we have $\int_{\bar{X}} k_{C^*}^a(z_1 \rightarrow x_2) dx_2 = 1$. This means that the probability of termination during an analog casting event,

$$p_{C^*}^a(z_1) := 1 - \int_{\bar{X}} k_{C^*}^a(z_1 \rightarrow x_2) dx_2 = 0.$$

Next, the direction change kernel is given by

$$k_{S^*}^a((x_2, v_1) \rightarrow v_2) := \begin{cases} \theta(x_2, v_1 \rightarrow v_2) \sigma(x_2)^{-1}, & x_2 \in X \\ \alpha(x_2) \Theta(x_2, v_1 \rightarrow v_2), & x_2 \in \partial X. \end{cases}$$

We find that the probability of termination during a direction change is given by

$$p_{S^*}^a(x_2) = \begin{cases} \sigma_a(x_2)/\sigma(x_2), & x_2 \in X \\ 1 - \alpha(x_2), & x_2 \in \partial X. \end{cases}$$

These kernels lead to the standard algorithm 1 [16]: we sample z_0 from the normalized source s (written $z_0 \sim s$), then cast according to $k_{C^*}^a$. Then particle is absorbed with probability $p_{S^*}^a$. If the photon is not absorbed, we change direction using a pdf proportional to $k_{S^*}^a$ ($k_{S^*}^a$ doesn't integrate to one, so it is not a pdf). Then we cast again and so on until we are absorbed. This defines the chain $\omega = ((x_0, v_0), \dots, (x_{\tau-1}, v_{\tau-1}), (x_\tau, \mathfrak{d}))$. At this point, we define the random variable modeling detector reading:

$$\xi_a(\omega) := \frac{\bar{g}(X_\tau, V_{\tau-1})}{p_{S^*}^a(X_\tau)}. \quad (11)$$

Note that our assumptions on α imply $p_{S^*}^a \equiv 1$ on the support of \bar{g} . The simplest

Algorithm 1 Analog

- 1: Draw $z_0 \sim s$, set $j \leftarrow 0$
 - 2: **while** $v_j \neq \mathfrak{d}$ **do**
 - 3: Draw $x_{j+1} \sim k_{C^*}^a(z_j \rightarrow \cdot)$
 - 4: With probability $p_{S^*}^a(x_{j+1}, v_j)$, $v_{j+1} = \mathfrak{d}$
 - 5: **if** $v_{j+1} \neq \mathfrak{d}$ **then**
 - 6: Draw v_{j+1} from a distribution $\propto k_{S^*}^a((x_{j+1}, v_j) \rightarrow \cdot)$
 - 7: **end if**
 - 8: $j \leftarrow j + 1$
 - 9: **end while**
 - 10: Record $\xi_a(\omega) = \bar{g}(x_j, v_{j-1})$
-

example is when the detector measures flux through the surface. In this case $\bar{g}(z) \equiv 1$ on $\text{supp}(g)$ and we simply count MC photons hitting the detector. Chains generated using algorithm 1 induce the analog probability measure

$$\begin{aligned} dP^a(\omega) &= s(z_0) k_{C^*}^a(z_0 \rightarrow x_1) k_{S^*}^a((x_1, v_0) \rightarrow v_1) \\ &\quad \times \cdots \times k_{C^*}^a(z_{\tau-2} \rightarrow x_{\tau-1}) k_{S^*}^a((x_{\tau-1}, v_{\tau-2}) \rightarrow v_{\tau-1}) \\ &\quad \times k_{C^*}^a(z_{\tau-1} \rightarrow x_\tau) p_{S^*}^a(x_\tau, v_{\tau-1}) dz_0 \cdots dz_{\tau-1} dx_\tau. \end{aligned} \quad (12)$$

It is instructive to write this out in the case where photons only interact with the volume, and then reach the detector. In this case (keeping in mind that $p_{\mathcal{G}^*}^a \equiv 1$ on the detector, and ignoring the $dz_0 \cdots dx_\tau$), dP^a becomes

$$s(z_0)\delta_{r(z_0)}(x_1)E_\sigma(x_0, x_1)\theta(x_1, v_0 \rightarrow v_1) \cdots \delta_{r(z_{\tau-1})}(x_\tau)E_\sigma(x_{\tau-1}, x_\tau). \quad (13)$$

We recall the normalization (3) from which we deduce that $\int_\Omega dP^a = 1$. Note first that the above chain is terminated with a cast and use of $k_{\mathcal{C}^*}^a$. Second, note that the measure above is a multiplication of singular measures and must be carefully defined. E.g. recall that we must fix z_1 in order for $k_{\mathcal{C}^*}^a(z_1 \rightarrow x_2)$ to be well defined; see the proof of theorem 3.1 (in the appendix) for details.

The next theorem shows that the chain $\xi_a(\omega)$ is indeed *unbiased*.

Theorem 3.1. *With $\mathbb{E}_a \{ \cdot \}$ denoting expectation under the measure P^a , we have*

$$\mathbb{E}_a \{ \xi_a \} = \langle u, g \rangle.$$

This result is standard in the absence of a boundary; see e.g. [16]. Its proof is sketched in the appendix. Algorithm 1 is a method for producing one draw (shot) $\xi_a(\omega)$ from P^a . As is “always” done with Monte Carlo techniques, we produce N draws $\{ \xi_a(\omega^i) \}_{i=1}^N$ in an identical fashion, then estimate

$$\langle u, g \rangle = \mathbb{E}_a \{ \xi_a \} \approx \frac{1}{N} \sum_{i=1}^N \xi_a(\omega^i).$$

3.1.2 Importance sampling

Here we give a quick introduction to importance sampling and show how it relates to our scheme. Given the analog probability measure dP^a , we can use a different measure $d\tilde{P}$ for sampling. With ξ_a defined as in (11), and $\tilde{\xi} := \xi_a \left| \frac{dP^a}{d\tilde{P}} \right|$,

$$\langle u, g \rangle = \mathbb{E}_a \{ \xi_a \} = \int_\Omega \xi_a dP^a = \int_\Omega \xi_a \left| \frac{dP^a}{d\tilde{P}} \right| d\tilde{P} = \mathbb{E} \{ \tilde{\xi} \}_{\tilde{P}},$$

where the Radon-Nikodym derivative $\left| \frac{dP^a}{d\tilde{P}} \right|$ (the Jacobian) must be defined on $\text{supp}(\xi_a)$. This happens precisely when, for any measurable $A \subset \Omega$ such that $P^a(A) > 0$, we also have $\tilde{P}(A) > 0$. In this case we say that P^a (or dP^a) is *absolutely continuous* with respect to \tilde{P} (or $d\tilde{P}$). Then we can estimate the measurement in one of two ways:

1. $\langle u, g \rangle \approx \frac{1}{N} \sum_{i=1}^N \xi_a(\omega_i)$, where ω_i are sampled according to P^a
2. $\langle u, g \rangle \approx \frac{1}{N} \sum_{i=1}^N \tilde{\xi}(\omega_i)$, where ω_i are sampled according to \tilde{P} .

For uncorrelated samples, the variance in either case ($\xi = \xi_a$ or $\xi = \tilde{\xi}$) is

$$\text{Var} \left\{ \frac{1}{N} \sum_{i=1}^N \xi(\omega_i) \right\} = \frac{\text{Var} \{ \xi \}}{N}.$$

So it will suffice to study $\text{Var} \{ \xi \}$ and the time needed per sample to calculate speedup. Here are a few expressions for variance of a random variable $\xi : \Omega \rightarrow \mathbb{R}$:

$$\text{Var} \{ \xi \} = \int_\Omega (\xi - \mathbb{E} \{ \xi \})^2 dP = \sum_{n=0}^{\infty} \int_{\tau=n} (\xi - \mathbb{E} \{ \xi \})^2 dP = \mathbb{E} \{ \xi^2 \} - \mathbb{E} \{ \xi \}^2.$$

The behavior of g puts some fundamental limits on variance for the *analog* chain. Let $D \subset \Omega$ be the set of paths that reach the detector, and suppose the “real life” detector measures flux through the surface, $g(z) = \nu_x \cdot v$ (on its support). Then $\bar{g} \equiv 1$ on the support of g so that $\xi_a(\omega) = \mathbb{1}_D(\omega)$ and the Monte Carlo detector acts as a photon counter. Then, for $\text{P}^a[D] \ll 1$,

$$\frac{\sqrt{\text{Var}\{\xi_a\}}}{\langle u, g \rangle} = \frac{\sqrt{\text{P}^a[D](1 - \text{P}^a[D])}}{\text{P}^a(D)} \approx \frac{1}{\sqrt{\text{P}^a[D]}}.$$

So for a small detector, the relative variance of analog MC is quite large. Since both methods are unbiased, variance is reduced if and only if

$$0 < \mathbb{E}_a \{\xi_a^2\} - \mathbb{E} \left\{ \tilde{\xi}^2 \right\}_{\tilde{\text{P}}} = \int_{\Omega} \xi_a^2 \left[1 - \left| \frac{d\text{P}^a}{d\tilde{\text{P}}} \right| \right] d\text{P}^a.$$

The goal of importance sampling is thus to make $d\tilde{\text{P}} \gg d\text{P}^a$ on as much of $\text{supp}(\xi_a)$ as possible. However, $d\tilde{\text{P}}$ must integrate to one and we must have $\left| \frac{d\text{P}^a}{d\tilde{\text{P}}} \right|$ defined on D (which we don’t have *a priori* access to).

We now describe some simplified importance sampling situations. They serve to bring intuition to our model. Suppose first that we devise an algorithm whose corresponding chain has measure

$$d\tilde{\text{P}}(\omega) = \begin{cases} G d\text{P}^a(\omega), & \omega \in D \\ \frac{1 - G\text{P}^a[D]}{1 - \text{P}^a[D]} d\text{P}^a(\omega), & \omega \notin D, \end{cases} \quad (14)$$

with $1 \leq G \leq \text{P}^a[D]^{-1}$. Now $\tilde{\xi} = \mathbb{1}_D/G$ and

$$\frac{\text{Var}\{\xi_a\}}{\text{Var}\{\tilde{\xi}\}} = \frac{1 - \text{P}^a[D]}{G^{-1} - \text{P}^a[D]}.$$

Theoretically, we can set $G = \text{P}^a[D]^{-1}$ and achieve infinite variance reduction, i.e., find a zero-variance method which gives the right result with probability 1. Assuming knowledge of $\text{P}^a[D]$ of course means we know the desired integral we are attempting to measure and thus is not practical. Moreover, practically, we cannot know how to increase $d\tilde{\text{P}}$ uniformly (and exclusively) for the *a priori* unknown $\omega \in D$, and thus some error is made. But this simple argument shows the possibility of achieving zero-variance MC. This will be utilized later in this section after we introduce importance sampling based on the adjoint transport calculations.

More practically, we may still devise schemes that increase the draws from some known, controlled, set $B \subset \Omega$, “stealing” them from $\Omega \setminus B$. In the simplified case where we change the measure on B by a multiplicative constant b and on $\Omega \setminus B$ by an appropriate constant so that mass is preserved, we obtain that

$$d\tilde{\text{P}}(\omega) = \begin{cases} b d\text{P}^a(\omega), & \omega \in B \\ \frac{1 - b\text{P}^a[B]}{1 - \text{P}^a[B]} d\text{P}^a(\omega), & \omega \in B^c, \end{cases} \quad \tilde{\xi}(\omega) = \begin{cases} \frac{1}{b}\xi(\omega), & \omega \in B \\ \frac{1 - \text{P}^a[B]}{1 - b\text{P}^a[B]}\xi(\omega), & \omega \in B^c. \end{cases} \quad (15)$$

Here, we have defined $B^c = \Omega \setminus B$. Then, assuming that $\xi = \mathbb{1}_D$ (i.e., that the detector counts photons),

$$\mathbb{E} \left\{ \tilde{\xi}^2 \right\}_{\tilde{\text{P}}} = \frac{1}{b} \text{P}^a[D \cap B] + \frac{1 - \text{P}^a[B]}{1 - b\text{P}^a[B]} \text{P}^a[D \cap B^c]. \quad (16)$$

Let us now optimize the choice of b to maximize variance reduction. Variance is significantly reduced when B is a good approximation of D , i.e., when $\mathbb{P}^a[D \cap B]$ is relatively close to $\mathbb{P}^a[D]$. How good an approximation we need may be quantified as follows. We recall that $\mathbb{P}^a[D \cap B] = \mathbb{P}^a[D]\mathbb{P}^a[B|D]$. We remind the reader that $\mathbb{P}[B|D]$ is the conditional probability of the event B given D . In other words, it is the probability that $\omega \in B$ given that the path ω reaches the detector.

Let us introduce the factors

$$\beta = b\mathbb{P}^a[D], \quad \gamma = \frac{\mathbb{P}^a[B]}{\mathbb{P}^a[D]}, \quad a = \frac{(1 - \gamma\mathbb{P}^a[D])(1 - \mathbb{P}^a[B|D])}{\mathbb{P}^a[B|D]\mathbb{P}^a[D]}. \quad (17)$$

Starting from (16), some algebra shows that

$$\text{Var} \left\{ \xi^2 \right\} = \mathbb{P}^{a^2}[D] \left(\mathbb{P}^a[B|D] \left(\frac{1}{\beta} + \frac{a}{1 - \gamma\beta} \right) - 1 \right).$$

Minimizing the above expression allows us to maximize the variance reduction. We find that for the optimal value of β_{opt} equal to $(\sqrt{\gamma}(\sqrt{\gamma} + \sqrt{a}))^{-1}$, the minimal variance is given by

$$\text{Var} \left\{ \tilde{\xi}^2 \right\}_{\min} = \mathbb{P}^{a^2}[D] \left(\mathbb{P}^a[B|D] (\sqrt{\gamma} + \sqrt{a})^2 - 1 \right).$$

This shows that the maximal variance reduction is given by

$$\frac{\text{Var} \left\{ \xi^2 \right\}}{\text{Var} \left\{ \tilde{\xi}^2 \right\}} \Big|_{\max} = \frac{1 - \mathbb{P}^a[D]}{\mathbb{P}^a[D]} \frac{1}{\mathbb{P}^a[B|D] (\sqrt{\gamma} + \sqrt{a})^2 - 1}. \quad (18)$$

When $B \equiv D$, we find that $\gamma = 1$ and $a = 0$. In that case, we find again that the above value is $+\infty$ and that the chain $\tilde{\xi}$ has zero variance.

In practice however, it is unlikely that a will be small. Since $\mathbb{P}^a[D]$ is small, we find that a is approximated by $\frac{1 - \mathbb{P}^a[B|D]}{\mathbb{P}^a[B|D]\mathbb{P}^a[D]}$. Since $\mathbb{P}^a[D] \ll 1$ for small detectors, a is likely to be large even for reasonable approximations of D by B . It turns out that even in that case, we can still expect good variance reductions. When $a \gg 1$ and γ close to 1, we observe that

$$\frac{\text{Var} \left\{ \xi^2 \right\}}{\text{Var} \left\{ \tilde{\xi}^2 \right\}} \Big|_{\max} \approx \frac{1}{1 - \mathbb{P}^a[B|D]}, \quad \beta_{\text{opt}} \approx \frac{1}{\sqrt{a\gamma}} \approx \left(\frac{\mathbb{P}^a[B|D]\mathbb{P}^a[D]}{\gamma(1 - \mathbb{P}^a[B|D])} \right)^{\frac{1}{2}}. \quad (19)$$

We observe that for a choice of b close to $\mathbb{P}^a[D]^{-1}\beta_{\text{opt}}$, we obtain very reasonable variance reduction when B is chosen so that $1 - \mathbb{P}^a[B|D] \ll 1$ but not necessarily $a \lesssim 1$ which is equivalent to $1 - \mathbb{P}^a[B|D] \lesssim \mathbb{P}^a[D]$ and imposes constraints on B that are not practical. We use the notation $a \lesssim b$ to denote “ $a \leq Cb$ for some $C < \infty$.” In figure 2, we show the variance reduction (18) for several values of $\mathbb{P}^a[B|D]$ (left) and its approximation by (19) (right), which works quite well when a is large and not so well when a is small as expected from theory. In all plots, $\mathbb{P}^a[D] = 0.002$, which is close to our actual simulations in section 4.

Note that (15) may be improved as follows when we know the existence of a set C such that $C \cap D = \emptyset$. Paths in C do not reach the detector and thus we want to give them a vanishing weight. The measure in (15) then needs to be modified as

$$d\tilde{\mathbb{P}}(\omega) = \begin{cases} b d\mathbb{P}^a(\omega), & \omega \in B \\ 0, & \omega \in C \\ \frac{1 - b\mathbb{P}^a[B]}{1 - \mathbb{P}^a[B] - \mathbb{P}^a[C]} d\mathbb{P}^a(\omega), & \omega \in (\Omega \setminus C) \setminus B. \end{cases} \quad (20)$$

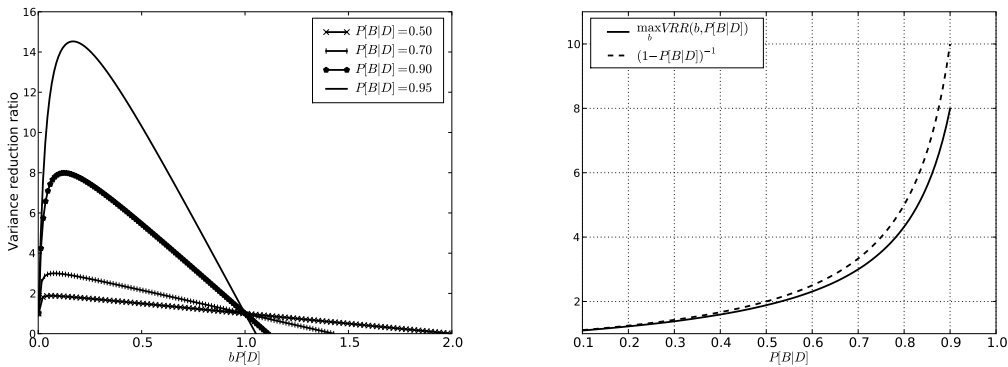


Figure 2: Variance reduction by importance sampling. Left: The ratio $VRR(b, P^a[B|D]) := \text{Var}\{\xi\} / \text{Var}\{\tilde{\xi}\}$ is plotted vs. $bP^a[D]$ for a number of different $P^a[B|D]$. Right: $\max_{b>0} VRR(b, P^a[B|D])$ is plotted versus $P^a[B|D]$. In both cases variance is calculated in the regime (16).

This leads to

$$\mathbb{E}\left\{\tilde{\xi}^2\right\}_{\tilde{P}} = \frac{1}{b}P^a[D \cap B] + \frac{1 - P^a[B] - P^a[C]}{1 - bP^a[B]}P^a[D \cap B^c]. \quad (21)$$

The situation (21) is preferable to (16) when $P^a[C] > 0$. The optimal value for b is obtained as before with $1 - \gamma P^a[D]$ in the definition of a replaced by $1 - P^a[C] - \gamma P^a[D]$.

3.1.3 Modular importance sampling

Finding the “right” set B is a difficult task: photons making it to the detector may undergo complicated interactions with the volume scatterers and the reflecting boundary. Moreover, in most settings of importance sampling, the derivative $\left|\frac{dP^a}{dP}\right|(\omega)$ does not take only two values as in the simplified setting (15). B should be replaced by one or several subsets $B = B_1 \cup B_2 \cup \dots$ where the weight $\left|\frac{dP^a}{dP}\right|(\omega)$ should be allowed to vary.

The *modularity* that we mention in the introduction consists of choosing sets B by appropriate approximations to the adjoint transport solution that are relatively simple to calculate and have a large intersection with D , the set of paths reaching the detector. For instance, a subset B_1 could correspond to particles reaching the detector after interacting with the boundary, B_2 to particles reaching the detector after one scattering event in a cloud, and so forth.

The importance sampling schemes considered in this paper are all based on changes of measure of the form $\left|\frac{dP^a}{dP}\right|(\omega)$ that generalize that seen in (16) or (21). We summarize them here.

The survival biasing method defined in section 3.1.4 below eliminates the volume and surface absorption of photons. Hence it “steals” shots from a subset of photons that were absorbed before reaching the detector, and moves them into some subset of D . We are therefore in the regime (21) (at least approximately as $\left|\frac{dP^a}{dP}\right|(\omega)$ is not constant on B); see also theorem 3.2 below.

The heuristic volume scattering method defined in section 3.1.5 below scatters (with probability < 1) photons in the volume directly toward the detector (rather than using the phase function θ). It has measure $dP_{heu,qv}$ uniformly larger than dP_{sb} on the set of paths that scatter once in the volume then hit the detector. It modifies the measure (often increasing it) on the set of paths that have their last interaction in the volume, then reach the detector. It steals shots from the set that interact last with the boundary, then hit the detector. A very rough approximation would put us in the regime (16) with $B_{heu} = \{\omega \in D : x_{\tau-1} \in X\}$.

The ideal zero-variance chain derived in section 3.2.1 below sends all photons to the detector. It uses an exact calculations of the adjoint solution to sample only from D , and is in the regime (14) with $G = P^a[D]^{-1}$; see theorem 3.3 below. When only approximate expressions for the adjoint solution are available, the zero-variance chain may be modified to yield small variance chains. However, in practice, the calculation of both the adjoint solution (step (i) in the abstract) and of the change of measure $\left| \frac{dP^a}{dP} \right|(\omega)$ (step (ii) in the abstract) is prohibitively expensive.

The SAI method defined in detail in section 3.4 below is our main example of a modular approach to importance sampling. In that method, we devise a subset $B = B_1 \cup B_2$, where B_1 corresponds to particles that do not undergo any volume scattering and where B_2 corresponds to photons that are sent straight to the detector after undergoing volume scattering. We will see that the method involves the calculation of an adjoint solution in the absence of volume scattering and that the calculation of $\left| \frac{dP^a}{dP} \right|(\omega)$ on B_1 , B_2 , and $\Omega \setminus (B_1 \cup B_2)$ is relatively straightforward. Moreover, we will see $B_1 \cup B_2$ is a good approximation of D when volume scattering is not too large although B_1 and B_2 individually are not necessarily good approximations of D . In the simplified calculations in (19) and in Fig.2, we observe that for $P^a[B_1|D] = 0.45$ and $P^a[B_2|D] = 0.45$, we may ideally have $P^a[B_1 \cup B_2|D] = 0.9$, with a potential variance reduction of order 10 whereas the variance reduction from B_1 or from B_2 alone would at best be a factor 2.

3.1.4 Survival Biasing

Here we define a classical chain where no photons are absorbed in X , although some are possibly in ∂X (for use in our application where we have perfectly absorbing boundaries). This will be related to the analog chain via importance sampling. Define

$$k_{C^*}^{sb}(z_1 \rightarrow x_2) := [\delta_{r(z_1)}(x_2)\sigma_s(x_2) + \delta(x_2 - x_+(z_1))] E_{\sigma_s}(x_1, x_2),$$

$$k_{S^*}^{sb}((x_2, v_1) \rightarrow v_2) := \begin{cases} \frac{\theta(x_2, v_1 \rightarrow v_2)}{\sigma_s(x_2)}, & x_2 \in X \\ \alpha^{sb}(x_2)\Theta(x_2, v_1 \rightarrow v_2), & x_2 \in \partial X, \end{cases}$$

with $\alpha^{sb}(x) = 1$ when $\alpha(x) > 0$ and $\alpha^{sb}(x) = 0$ when $\alpha(x) = 0$. We then have:

$$p_{C^*}^{sb}(z_1) = 0 \quad \text{and} \quad p_{S^*}^{sb}(x_2, v_1) = \begin{cases} 0, & x_2 \in X \\ 1 - \alpha^{sb}(x_2), & x_2 \in \partial X. \end{cases}$$

The Radon-Nikodym derivative is obtained by formally dividing dP^a by dP_{sb} , where dP^a is defined in (12), and dP_{sb} is defined analogously. Since, for $(x_\tau, v_{\tau-1}) \in \text{supp}(g)$, $\alpha(x_\tau) = \alpha^{sb}(x_\tau) = 0$, the Radon-Nikodym derivative, restricted to

$\{\omega : x_\tau \in \pi_x \text{supp}(g)\}$ is

$$\begin{aligned} \left| \frac{dP^a}{dP_{sb}} \right| &= E_{\sigma-\sigma_s}(x_0, x_1, \dots, x_\tau) \gamma_{a,sb}(x_1) \cdots \gamma_{a,sb}(x_{\tau-1}), \\ \gamma_{a,sb}(x) &:= \begin{cases} 1, & x \in X, \\ \alpha(x)/\alpha^{sb}(x), & x \in \partial X. \end{cases} \end{aligned} \quad (22)$$

Defining

$$\xi_{sb} := \xi \left| \frac{dP^a}{dP_{sb}} \right|,$$

we have

$$\mathbb{E}_{sb} \{\xi_{sb}\} = \mathbb{E}_a \{\xi_a\} = \langle u, g \rangle.$$

Theorem 3.2 (Variance reduction by eliminating absorption). *We have*

$$\text{Var} \{\xi_{sb}\} \leq \text{Var} \{\xi_a\}$$

with equality if and only if absorption is zero (with probability = 1) on analog paths that reach the detector.

Proof. Since both methods are unbiased, it will suffice to consider the expected value of the random variable squared. Since $E_{\sigma-\sigma_s}(x, y) \leq 1$, (with equality if and only if $\sigma = \sigma_a$ along the path from x to y), and for $j < \tau$, $\gamma_{a,sb}(x_j) \leq 1$ (with equality if and only if $\alpha(x_j) = 1$),

$$\mathbb{E} \{\xi_{sb}^2\}_{P_{sb}} = \mathbb{E} \left\{ \xi_a^2 \left| \frac{dP^a}{dP_{sb}} \right| \right\}_{P^a} \leq \mathbb{E} \{\xi_a^2\}_{P^a},$$

with equality occurring only under the specified conditions. \square

Note that since photons are not absorbed, their path length could be much longer than in standard analog sampling. This could result in a decrease in our figure of merit (see section 4). In nuclear reactor applications, the multiplication of particles with very small weights becomes a real issue and several techniques such as Russian roulette have been developed to address this [17, 9]. In remote sensing applications with a reasonably large mean free path, and the chance of escape into the atmosphere, this is much less of an issue and thus is not considered in this paper.

3.1.5 Heuristic volume scattering adjustment

In this section, we present a very simple (and classical) direction change kernel to be used as part of any modular scheme to handle volume scattering (we use it as part of SAI). We modify the volume scattering kernel in order to direct photons toward the detector. When a large fraction of photons reach the detector with only zero or one volume scattering event (i.e. when σ_s is small), this is a reasonable method. Although better methods do exist, we include this to demonstrate our modular variance reduction paradigm. We introduce a regularization parameter $q_v \in (0, 1]$. We draw from our modified method a fraction of the time approximately proportional to $1 - q_v$.

Let x_{d_0} be the midpoint of the detector (assume one detector). For $q_v \in [0, 1]$, $x_2 \in X$, put

$$q_{heu}(x_2, v_1) := (q_v - 1) \frac{\theta(x_2, v_1 \rightarrow \widehat{x_{d_0} - x_2})}{\|\theta(x_2, v_1 \rightarrow \cdot)\|_{L^\infty}} + 1. \quad (23)$$

For $x_2 \in X$, let $f_V(x_2 \rightarrow v_2)$ be uniform on $\{v \in \mathbb{S}^{d-1} : r(x_2, v) \cap \pi_x \text{supp}(g) \neq \emptyset\}$. We define the *heuristic scattering adjustment* direction change kernel by

$$k_{S^*}^{heu}((x_2, v_1) \rightarrow v_2) := \begin{cases} [1 - q_{heu}(x_2, v_1)]f_V(x_2 \rightarrow v_2) + q_{heu}(x_2, v_1)k_{S^*}^{sb}((x_2, v_1) \rightarrow v_2), & x_2 \in X \\ k_{S^*}^{sb}((x_2, v_1) \rightarrow v_2), & x_2 \in \partial X. \end{cases}$$

So we are aimed toward the detector via f_V with probability $1 - q_{heu}$. The ratio of θ to its L^∞ norm in (23) is proportional to the analog probability of heading toward the detector; certainly we don't want to send particles toward the detector when the analog chain would *never* do that (the Radon-Nikodym derivative would be zero in this case).

For convenience, we calculate here the change of measure associated to the chain that uses survival biasing on the boundary and volume, as well as heuristic direction changes in the volume. This is the *heuristic chain*

$$\begin{aligned} \left| \frac{dP_{heu, q_v}}{dP_{sb}} \right| &= \gamma_{heu, sb}(x_1, z_2) \cdots \gamma_{heu, sb}(x_{\tau-2}, z_{\tau-1}), \\ \gamma_{heu, sb}(x_1, z_2) &= \begin{cases} \frac{(1 - q_{heu})f_V(x_2 \rightarrow v_2)\sigma_s(x_2) + q_{heu}\theta(x_2, v_1 \rightarrow v_2)}{\theta(x_2, v_1 \rightarrow v_2)}, & x_2 \in X \\ 1, & x_2 \in \partial X. \end{cases} \end{aligned} \quad (24)$$

To produce one draw ω from the heuristic chain, we follow algorithm 2. This could then be used to estimate $\langle u, g \rangle$. In this paper, we combine the heuristic chain with an adjoint-based method. See section 3.4.

3.2 Adjoint-based importance sampling

In this section, we first show how knowledge of the exact adjoint transport solution allows us to devise a zero-variance method. This generalizes to the case of transport with boundaries well-known results for volume scattering [16, 17]. When the adjoint solution is approximated, e.g., by a deterministic calculation, we show how a non-analog MC chain may be generated. We show that when the approximation of the adjoint solution is of order h for a ‘‘mesh’’ size $h \ll 1$, then the MC variance is of order h^2 in ideal circumstances (and larger in more complex geometries). We will present in section 3.4 a hybrid method that only calculates *important* parts of the adjoint solution at a minimal computational cost while still offering sizable variance reductions.

3.2.1 The zero-variance chain

Here we describe a chain that uses an exact adjoint solutions (ψ_i^*, ψ_o^*) and has zero variance. We show that draws from the chain can be made in a manner similar to analog, with modified scattering cross-sections. Obtaining ψ_i^*, ψ_o^* is more difficult than solving our original problem (they must be obtained everywhere), hence as we mentioned earlier this method is impractical.

Algorithm 2 Heuristic scattering adjustment

```

1: Draw  $z_0 \sim s$ , set  $j \leftarrow 0$ 
2: while  $v_j \neq \mathfrak{d}$  do
3:   if  $x_2 \in X$  then
4:     Compute  $q_{heu}(x_{j+1}, v_j)$  using (23). With probability  $1 - q_{heu}$  set  $switch \leftarrow \mathbf{true}$ 
5:     if  $switch$  then
6:       Draw  $v_2 \sim f_V(x_2 \rightarrow \cdot)$ 
7:     else
8:       Draw  $v_2 \sim k_{S^*}^{sb}((x_2, v_1) \rightarrow \cdot)$ 
9:     end if
10:  else
11:    With probability  $p_{S^*}^{sb}(x_2) = 1 - \alpha^{sb}(x_2)$ ,  $v_{j+1} \leftarrow \mathfrak{d}$ 
12:  end if
13:  if  $v_{j+1} \neq \mathfrak{d}$  then
14:    Draw  $v_{j+1}$  from a distribution  $\propto k_{S^*}^a((x_{j+1}, v_j) \rightarrow \cdot)$ 
15:  end if
16:   $j \leftarrow j + 1$ 
17: end while
18: Record  $\xi_{heu}(\omega) = \bar{g}(x_j, v_{j-1})$ 

```

Our Markov chain formulation phrases the use of the adjoint in terms of transition kernels. This was done explicitly in [16] and implicitly in [17]. Unlike [16] we explicitly write out the modified ray-casting and direction-change kernels. Unlike either scheme, we explicitly use both the incoming ψ_i^* and outgoing ψ_o^* adjoint solutions. In [16] ψ_i^* was used (implicitly) and in [17] both were used (implicitly). Define

$$\begin{aligned}
k_{C^*}^*(z_1 \rightarrow x_2) &:= [\delta_{r(z_1)}(x_2) + \delta(x_2 - x_+(z_1))] E_\sigma(x_1, x_2) \frac{\psi_i^*(x_2, v_1)}{\psi_o^*(z_1)}, \\
k_{S^*}^*((x_2, v_1) \rightarrow v_2) &:= \begin{cases} \theta(x_2, v_1 \rightarrow v_2) \frac{\psi_o^*(x_2, v_2)}{\psi_i^*(x_2, v_1)}, & x_2 \in X, \\ \alpha(x_2) \Theta(x_2, v_1 \rightarrow v_2) \frac{\psi_o^*(x_2, v_2)}{\psi_i^*(x_2, v_1)}, & x_2 \in \partial X. \end{cases}
\end{aligned}$$

So we modify the casting by the ratio of the importance of the point we will enter to the importance of the point we are exiting. We modify direction changes by the ratio of the importance entering x_2 to that exiting.

Using the equations defining the adjoint solutions, we have

$$\begin{aligned}
p_{C^*}^*(z_1) &:= 1 - \int_{\bar{X}} k_{C^*}^*(z_1 \rightarrow x_2) dx_2 = 1 - \frac{C^* \psi_i^*(z_1)}{\psi_o^*(z_1)} = 0, \\
p_{S^*}^*(x_2, v_1) &:= 1 - \int_{\mathbb{S}^{d-1}} k_{S^*}^*((x_2, v_1) \rightarrow v_2) dv_2 = 1 - \frac{S^* \psi_o^*(x_2, v_1)}{\psi_i^*(x_2, v_1)} \quad (25) \\
&= \frac{\bar{g}(x_2, v_1)}{\psi_i^*(x_2, v_1)} = \begin{cases} 1, & (x_2, v_1) \in \text{supp}(\bar{g}) \\ 0, & \text{otherwise.} \end{cases}
\end{aligned}$$

The last equality used the fact that $\psi_i^* = \bar{g}$ on the support of \bar{g} (since $\alpha = 0$ there). So all photons reaching the detector are collected.

We also define a new (normalized) source

$$s^* := \frac{s\psi_o^*}{\langle s, \psi_o^* \rangle}.$$

In other words, we bias the photons leaving the source so that they leave in directions with high importance.

Note that sampling is done by alternately casting along a line, then changing direction, just as in an analog scheme. Since (off the detector) both $k_{C^*}^*$ and $k_{S^*}^*$ integrate to one, they are probability densities. To sample from $k_{C^*}^*(z_1 \rightarrow x_2)$ it therefore suffices to cast along the ray $r(z_1)$ and integrate $k_{C^*}^*$ as we go. Once the integral is greater than some uniform random number $u \sim \mathcal{U}[0, 1]$, we scatter. The relation (10) along with algorithm 1 show that this same procedure is done in standard analog Monte Carlo. Sampling from $k_{S^*}^*$ may also be done just as in analog Monte Carlo.

We define dP^* in the same manner as dP^a . This yields,

$$dP^*(\omega) = s^*(z_0)k_{C^*}^*(z_0 \rightarrow x_1)k_{S^*}^*((x_1, v_0) \rightarrow v_1) \cdots k_{C^*}^*(z_{\tau-1} \rightarrow x_\tau)p_{S^*}^*(x_\tau, v_{\tau-1}) \\ \times dz_0 \cdots dx_\tau.$$

It is instructive to see that most terms involving ψ^* cancel in the calculation of the Radon-Nikodym derivative $\left| \frac{dP^a}{dP^*} \right|$. Restricting ourselves to paths that do not interact with the boundary and ignoring $dz_0 \cdots dx_\tau$, $dP^*(\omega)$ takes the form:

$$\frac{s(z_0)\psi_o^*(z_0)}{\langle s, \psi_o^* \rangle} \delta_{r(z_0)}(x_1)E_\sigma(x_0, x_1) \frac{\psi_i^*(x_1, v_0)}{\psi_o^*(z_0)} \theta(x_1, v_0 \rightarrow v_1) \\ \times \delta_{r(z_1)}(x_2)E_\sigma(x_1, x_2) \frac{\psi_i^*(x_2, v_1)}{\psi_o^*(z_1)} \cdots \frac{\bar{g}(x_\tau, v_{\tau-1})}{\psi_i^*(x_\tau, v_{\tau-1})} \\ = \langle s, \psi_o^* \rangle^{-1} s(z_0) \delta_{r(z_0)}(x_1)E_\sigma(x_0, x_1) \theta(x_1, v_0 \rightarrow v_1) \delta_{r(z_1)}(x_2)E_\sigma(x_1, x_2) \cdots \bar{g}(x_\tau, v_{\tau-1}).$$

This easily combines with (13) to yield (26) below for the Radon-Nikodym derivative $\left| \frac{dP^a}{dP^*} \right|$ restricted to paths that do not interact with the boundary.

More generally, for all paths, which account for both volume and boundary interactions, we verify (after careful algebra) that the Radon-Nikodym derivative (restricted to the set D) is still given by

$$\left| \frac{dP^a}{dP^*} \right| = \frac{\langle s, \psi_o^* \rangle}{\bar{g}(x_\tau, v_{\tau-1})}. \quad (26)$$

Since (for $(x_\tau, v_{\tau-1}) \in \text{supp}(\bar{g})$), $p_{S^*}^a(x_\tau, v_{\tau-1}) = 1$, the appropriate random variable to measure is

$$\xi^* := \left| \frac{dP^a}{dP^*} \right| \xi_a = \left| \frac{dP^a}{dP^*} \right| \frac{\bar{g}(x_\tau, v_{\tau-1})}{p_{S^*}^a(x_\tau, v_{\tau-1})} = \langle s, \psi_o^* \rangle.$$

In the event of highly scattering media, it would be advantageous to use a scheme that would reduce the number of scattering events seen by a photon. More generally, we would hope that a less expensive route to the detector could be taken. Unfortunately, the next theorem shows that the zero-variance scheme cannot do this. Fortunately, it also shows that the modified chain does not take a more expensive route to the detector. We remind the reader that $P[A|D]$ is the conditional probability of the event A given D . In other words, it is the probability that $\omega \in A$ given that we reach the detector.

Theorem 3.3 (Identical Paths). *Let $A \subset \Omega$ be measurable, and let $D \subset \Omega$ denote the paths that end with $(x_\tau, v_{\tau-1}) \in \text{supp}(g)$, then*

$$\mathbb{P}^*[A] = \frac{\int_A \bar{g}(X_\tau, V_{\tau-1}) d\mathbb{P}^a}{\langle \psi_o^*, s \rangle}.$$

In the special case $\bar{g} \equiv 1$ on $\text{supp}(g)$, then $\mathbb{P}^[A] = \mathbb{P}^*[A | D] = \mathbb{P}^a[A | D]$. In other words, the paths taken to the detector in the modified scheme are the exact same as in the analog scheme.*

The following corollary follows by letting $A = \{\tau = n\}$.

Corollary 3.4 (Constant Collision Ratios).

$$\mathbb{P}^*[\tau = n] = \frac{\int_{\tau=n} \bar{g}(X_\tau, V_{\tau-1}) d\mathbb{P}^a}{\langle \psi_o^*, s \rangle}.$$

In the special case $\bar{g} = 1_{\text{supp}(g)}$, then $\mathbb{P}^[\tau = n] = \mathbb{P}^a[\tau = n | D]$. In other words, the number of collisions photons have before hitting the detector is the same in the analog or zero-variance scheme.*

These results show that the zero-variance chain (when the MC detector counts photons) is precisely in the regime (14) with $G = \mathbb{P}^a[D]^{-1}$. In other words, we increase the measure uniformly (an optimum amount) on the set of paths that reach the detector.

Proof of theorem 3.3.

$$\mathbb{P}^*[A] = \int_A \left| \frac{d\mathbb{P}^*}{d\mathbb{P}^a} \right| d\mathbb{P}^a = \frac{1}{\langle s, \psi_o^* \rangle} \int_A g(X_n, V_{n-1}) d\mathbb{P}^a.$$

This proves the first part. When $\bar{g} \equiv 1$ on $\text{supp}(g)$, the above becomes

$$\mathbb{P}^*[A] = \frac{\mathbb{P}^a[A \cap D]}{\langle s, \psi_o^* \rangle}.$$

This leads to

$$\mathbb{P}^a[D] = \sum_{n=0}^{\infty} \mathbb{P}^a[(\tau = n) \cap D] = \langle s, \psi_o^* \rangle \sum_{n=0}^{\infty} \mathbb{P}^*[\tau = n] = \langle s, \psi_o^* \rangle.$$

The result then follows from the definition of conditional probability. \square

3.2.2 Approximations of the zero variance chain

After seeing the zero-variance chain, one immediately gets the idea of using approximations to ψ_i^* , ψ_o^* in a variance reduction method. Assuming one can generate these approximations (e.g., by using a deterministic solver), it still remains to construct a bona fide chain (a probability density integrating to 1) and to sample from the corresponding chain (this is step (ii) in the abstract). We show here that an arbitrarily coarse adjoint approximation can be used in an approximation of the zero-variance scheme.

The approximation $d\mathbb{P}^h$ of $d\mathbb{P}^*$ can be used in the so-called ‘‘asymptotic regime’’ where $\psi_o^h \approx \psi_o^*$. In this setting, the calculation of the adjoint solutions and the sampling from the measure $d\mathbb{P}^h$ may be prohibitively expensive as

the number of degrees of freedom necessary to adequately represent the adjoint solution is typically very large.

This approximation can also be used to guide photons along paths to the detector. There, it only needs to perform sufficiently well and no longer needs to be very accurate. In this case we draw from dP^h (a “not-necessarily-good” approximation of dP^*) only an optimized fraction of the time, while e.g. using the analog measure to sample from the rest of the time. In many cases good speedup is obtained. The latter methodology is implemented by the SAI chain in section 3.4.

Assume one has $\psi_i^h \approx \psi_i^*$, and $\psi_o^h \approx \psi_o^*$. Then, following the recipe of the zero-variance chain, we could set

$$k_{C^*}^h(z_1 \rightarrow x_2) := [\delta_{r(z_1)}(x_2) + \delta(x_2 - x_+(z_1))] E_\sigma(x_1, x_2) \frac{\psi_i^h(x_2, v_1)}{\psi_o^h(z_1)},$$

$$k_{S^*}^h((x_2, v_1) \rightarrow v_2) := \begin{cases} \theta(x_2, v_1 \rightarrow v_2) \frac{\psi_o^h(x_2, v_2)}{\psi_i^h(x_2, v_1)}, & x_2 \in X, \\ \alpha(x_2) \Theta(x_2, v_1 \rightarrow v_2) \frac{\psi_o^h(x_2, v_2)}{\psi_i^h(x_2, v_1)}, & x_2 \in \partial X. \end{cases} \quad (27)$$

However, many difficulties arise. For example, it is not clear that ψ_i^h, ψ_o^h are nonzero whenever ψ_i^*, ψ_o^* are. In this case, the modified chain will not send photons along all paths that the analog chain does, and the result will be biased. Assuming we take care of this problem, a more insidious issue arises: What are the values of the integrals $\int_X k_{C^*}^h dx, \int k_{S^*}^h dv$? If both integrate to one (away from the detector), then, as in the zero variance chain, we use them as pdfs and sample directly from them (say with an accept-reject method, or by pre-calculating a cdf). If they integrate to less than one, this gives us a probability of absorption, and we need to know this. If they integrate to more than one (very possible), then one can still sample from a pdf proportional to them. However, this proportionality constant must be known when the Radon-Nikodym derivative is calculated.

To formalize this, we propose modified kernels of the form

$$k_{C^h}^a(z_1 \rightarrow x_2) := [\delta_{r(z_1)}(x_2) + \delta(x_2 - x_+(z_1))] E_\sigma^h(x_1, x_2),$$

$$k_{C^h}^h(z_1 \rightarrow x_2) := k_{C^h}^a(z_1 \rightarrow x_2) \frac{\psi_i^h(x_2, v_1)}{\psi_o^h(z_1)},$$

$$k_{S^h}^a((x_2, v_1) \rightarrow v_2) := \begin{cases} \theta^h(x_2, v_1 \rightarrow v_2), & x_2 \in X, \\ \alpha^h(x_2) \Theta^h(x_2, v_1 \rightarrow v_2), & x_2 \in \partial X. \end{cases} \quad (28)$$

$$k_{S^h}^h((x_2, v_1) \rightarrow v_2) := k_{S^h}^a((x_2, v_1) \rightarrow v_2) \frac{\psi_o^h(x_2, v_2)}{\psi_i^h(x_2, v_1)}.$$

The new coefficients ($E_\sigma^h, \theta^h, \alpha^h, \Theta^h$) are chosen such that the kernels integrate to one or less. In most cases, away from the detector, one would choose the integrals to be one (so particles are not absorbed). We also assume that we have approximations of the detector and source, $\bar{g}^h \approx \bar{g}, s^h \approx s$. Note that the calculation of such coefficients may prove to be quite expensive numerically (this is step (ii) introduced in the abstract). In some sense, the coefficients in (28) may be seen as normalized versions of the coefficients introduced in (27). However, finding rules to calculate this normalizing constants efficiently is non-trivial. We will address this issue in the following section in the simplified setting where volume scattering is absent.

Note that such normalizing constants would easily be calculated if $k_{C^h}^a$ and $k_{S^h}^a$ were the kernels of operators C^h and S^h , respectively, and ψ_o^h and ψ_i^h were obtained by solving the equations $\psi_o^h = C^h S^h \psi_o^h + C^h \bar{g}^h$ and $\psi_i^h = S^h C^h \psi_i^h + \bar{g}^h$. Indeed as in (25), we would then obtain that

$$1 - \int_{\bar{X}} k_{C^h}^h(z_1 \rightarrow x_2) dx_2 = 1 - \frac{C^h \psi_i^h(z_1)}{\psi_o^h(z_1)} = 0,$$

$$1 - \int_{\mathbb{S}^{d-1}} k_{S^h}^h((x_2, v_1) \rightarrow v_2) dv_2 = 1 - \frac{S^h \psi_o^h(x_2, v_1)}{\psi_i^h(x_2, v_1)}. \quad (29)$$

However, such operators S^h and C^h would preserve the singularities of the transport equation (primarily propagation along straight lines) and are therefore cannot be discrete. Their kernels in (28) are infinite dimensional and cannot be reduced to (finite dimensional) matrices. Any reduction to a matrix form involves approximations that will modify the structure of the singularities in (28) and render the integrals in (29) more complicated to estimate.

In any case, assuming that our construction (28) defines a bona fide change of measures (i.e. \mathbb{P}^a is absolutely continuous with respect to \mathbb{P}^h) so that the Radon-Nidodym derivative restricted to $\{\omega : z_\tau \in \text{supp}(\bar{g})\}$ is well defined, then the latter is given by:

$$\left| \frac{d\mathbb{P}^a}{d\mathbb{P}^h} \right| = \frac{\langle s^h, \psi_o^h \rangle}{\bar{g}^h(x_\tau, v_{\tau-1}) s^h(z_0)} \beta_{a,h}(x_0, \dots, x_\tau) \gamma_{a,h}(z_1, \dots, z_{\tau-1}),$$

$$\gamma_{a,h}(z_1, z_2) := \begin{cases} \frac{\theta(x_2, v_1 - v_2)}{\theta^h(x_2, v_1 - v_2)}, & x_2 \in X \\ \frac{\alpha(x_2) \Theta(x_2, v_1 - v_2)}{\alpha^h(x_2) \Theta^h(x_2, v_1 - v_2)}, & x_2 \in \partial X, \end{cases} \quad (30)$$

$$\gamma_{a,h}(z_1, \dots, z_n) := \gamma_{a,h}(z_1, \dots, z_{n-1}) \gamma_{a,h}(z_{n-1}, z_n).$$

$$\beta_{a,h}(x_1, x_2) := \frac{E_\sigma(x_1, x_2)}{E_\sigma^h(x_1, x_2)},$$

$$\beta_{a,h}(x_0, \dots, x_n) := \beta_{a,h}(x_0, \dots, x_{n-1}) \beta_{a,h}(x_{n-1}, x_n).$$

Since $(\theta, \Theta, \alpha, E_\sigma) \neq (\theta^h, \Theta^h, \alpha^h, E_\sigma^h)$ a priori, the telescopic cancellations in (26) no longer occur in (30). We then set

$$\xi^h := \xi_a \left| \frac{d\mathbb{P}^a}{d\mathbb{P}^h} \right|, \text{ so that } \mathbb{E}_h \{ \xi^h \} = \langle u, g \rangle.$$

When $\psi_o^h \approx \psi_o^*$ we expect $\text{Var} \{ \xi^h \} \ll 1$. The rate of convergence is studied here in the ideal setting where the following assumptions are satisfied:

Assumptions 3.1. Assume there exist $\rho, C > 0$ such that, for all small enough h ,

- (i) $|\langle s^h, \psi_o^h \rangle / \langle u, g \rangle - 1| \leq Ch$,
- (ii)

$$\left| \frac{\bar{g}}{\bar{g}^h} - 1 \right| + \left| \frac{s}{s^h} - 1 \right| + |\gamma_{a,h}(z_1, z_2) - 1| + |\beta_{a,h}(x_1, x_2) - 1| \leq Ch$$

- (iii) $\mathbb{P}^h[\tau = n] \leq C e^{-\rho n}$

- (iv) $\text{supp}(\psi_o^*) = \text{supp}(\psi_o^h)$, and $\text{supp}(\psi_i^*) = \text{supp}(\psi_i^h)$

Assumptions (i) and (ii) follow if all approximations are $O(h)$ in the uniform norm, all coefficients are bounded below (on their support), and the support of the true and approximate coefficients are the same. The third assumption (iii) is standard in the transport regime with not-too-small mean free path and simply indicates that long-distance, multiple-scattering paths are improbable. Assumptions (ii), (iv) ensure $\left| \frac{dP^a}{dP^h} \right|$ exists everywhere.

Verifying assumptions (i) and (ii) is extremely constraining. However, in this idealized setting, we have the following theorem, whose proof is postponed to section A.2.

Theorem 3.5 (Convergence in the asymptotic regime). *Assume that we meet Assumptions 3.1. Then as $h \rightarrow 0$,*

$$\text{Var} \left\{ \xi^h \right\} \leq \langle u, g \rangle^2 C' h^2,$$

for some $C' > 0$ depending on C and ρ .

This result shows that importance samplings with small variance can be achieved provided that accurate approximations to adjoint transport solutions are available.

The aim of all remaining sections and the introduction of the SAI method is precisely an attempt at using an adjoint approximation that (i) is inexpensive to calculate; and (ii) generates a measure that is both easy to sample from and has small variance.

3.3 Reflecting boundaries without volume scattering

Here we discretize the operator appearing in section 2.3. This operator arises in the limit of zero volume interactions ($\sigma \rightarrow 0$). We first present a discretization of the adjoint solution in section 3.3.1 and then show how the adjoint solution can be used to obtain a non-analog MC algorithm with small variance in section 3.3.2.

3.3.1 Surface-limit adjoint problem

In this limit, we have $\psi_i^* \rightarrow \psi_i^s$. Here, at discretization level h , we approximate $\psi_i^h \approx \psi_i^s$ and $\psi_o^h \approx \psi_o^s$. In this section we assume the boundary is sufficiently smooth (of class C^3).

To simplify computation of our numerical solution we make the assumptions

$$\Theta(x, v \rightarrow v') = \mathbb{1}_{\nu_x \cdot v > 0}(x, v) \kappa(x, v'), \quad g(z) = |\nu_x \cdot v| g_0(x),$$

so that $\bar{g}(z) = g_0(x)$. We recall that $\mathbb{1}_A$ is the “indicator” function of the set A . The result is that ψ_i^s is then a function of position only. This significantly improves the speed of solving the adjoint problem, as well as the memory requirements for using it. Theoretical results in this paper do not need this assumption, which we make here as a matter of convenience.

We will now discretize the coefficients and approximate the operator appearing in section 2.3. For $z_1 \in \Gamma_+$,

$$S^s C^s \psi_i^s(z_1) = \alpha(x_1) \int_{\nu_{x_1} \cdot v_2 < 0} \kappa(x_1, v_2) \psi_i^s(z_+(x_1, v_2)) dv_2.$$

Notice that $S^s C^s f$ is function depending only on x , and in fact only on the boundary values of f . Since \bar{g} depends only on x , $\psi_i^s = \sum_{k=0}^{\infty} (S^s C^s)^k \bar{g}$ will depend only on x . We thus define

$$\varphi(x) := \psi_i^s|_{\Gamma_+}(x, v).$$

We find that $\varphi : \partial X \rightarrow \mathbb{R}$ satisfies the equation

$$\varphi = Q\varphi + g_0, \quad Qf(x_1) := \alpha(x_1) \int_{\nu_{x_1} \cdot v_2 < 0} \kappa(x_1, v_2) f(x_+(x_1, v_2)) dv_2.$$

In discretizing this operator, and integrals over directions in general, we use the change of variables,

$$\begin{aligned} \int_{\nu_x \cdot v < 0} f(z_+(x, v)) dv &= \int_{\partial X} f(x', v) \partial_\nu N(x, x') d\mu(x'), \\ \partial_\nu N(x, x') &:= \frac{\nu_x \cdot (x' - x)}{|x' - x|^d}. \end{aligned} \quad (31)$$

The term $\partial_\nu N$ is normal derivative (at x) of the free-space Green's function for the Laplacian. One can show (see e.g. the section on double-layer potentials in [8]) that for $x, x' \in \partial X$, $\nu_x \cdot (x' - x) \lesssim |x' - x|^2$. Therefore it is in fact an integrable function. When $d = 2$ we have more.

Lemma 3.1. *When $d = 2$, if ∂X is C^{k+2} , then $\partial_\nu N(x, x')$ is $C^k(\partial X \times \partial X)$.*

The proof is postponed to Appendix A.3. We now discretize the operator Q . First split the boundary into non-overlapping segments $\{\partial X_j\}_{j=0}^{N_p-1}$ with ∂X_j centered at x_j , with measure $|\partial X_j| \leq h$. Denote by Rf the (orthogonal) projection of f onto the space of piecewise constant functions (constant on each segment ∂X_j). We also think of Rf as a vector in \mathbb{R}^{N_p} and Rf_j its components. Then, after the change of variables (31) we have (at gridpoint x_i)

$$\begin{aligned} Qf(x_i) &= \alpha(x_i) \int_{\partial X} \kappa(x_i, \widehat{x - x_i}) \partial_\nu N(x_i, x) f(x) d\mu(x) \\ &\approx \alpha(x_i) \sum_{\substack{0 \leq j \leq N_p-1 \\ j \neq i}} |\partial X_j| \kappa(x_i, \widehat{x_j - x_i}) \partial_\nu N(x_i, x_j) f(x_j) \\ &:= \sum_j \mathcal{Q}_{ij}^h Rf_j. \end{aligned} \quad (32)$$

This implicitly defines the matrix \mathcal{Q}^h . So long as ∂X is C^2 , the above integrand L^1 (bounded in two dimensions), hence we are justified in approximating it as such.

We now define our discrete approximation to φ as the piecewise constant function (vector) φ^h solving

$$\varphi^h = \mathcal{Q}^h \varphi^h + R\bar{g}. \quad (33)$$

We then define approximations $\psi_i^h \approx \psi_i^*$, $\psi_o^h \approx \psi_o^*$ by

$$\psi_i^h(x, v) := \varphi^h(x), \quad \psi_o^h(z_-(x, v)) := \psi_i^h(x, v), \quad (x, v) \in \Gamma_+. \quad (34)$$

The next proposition is used to apply convergence theorems to the SAI chain.

Proposition 3.1. *Suppose $d = 2$, $\|\alpha - \alpha^h\|_{L^\infty} \lesssim h$, $\|\kappa - \kappa^h\|_{L^\infty} \lesssim h$. Then as operators $: L^1(\partial X) \rightarrow L^\infty(\partial X)$, we have*

$$\|Q - Q^h\| \lesssim h.$$

Furthermore, assuming $\|Q\| < 1$, $\|Q^h\| < 1$, then we have

$$\|\psi_i^h - \psi_i^s\|_{L^\infty(\Gamma_+)} \lesssim h.$$

Proof. The first inequality follows by a bound on the coefficients of $Q - Q^h$, keeping in mind that the apparent singularity is actually a bounded function in dimension 2. The second follows from $\varphi = \sum_{n=0}^{\infty} Q^n R \bar{g}$, $\varphi^h = \sum_{n=0}^{\infty} (Q^h)^n R \bar{g}$, and repeated application of relations similar to $ab - \tilde{a}\tilde{b} = (a - \tilde{a})b + \tilde{a}(b - \tilde{b})$. \square

In our implementation, we have chosen to represent angular integrals as integrals over the boundary. This works for two reasons. First, as our adjoint solution depends only on position it is convenient to evaluate these sums. Second, if instead a discretization were chosen that was uniform in angle, then (with only finitely many angles) one would often miss the (small) detector in evaluation of the integral.

3.3.2 Surface-adjoint approximations and non-analog chains

The SAI chain makes use of the approximate surface adjoint solutions ψ_i^h, ψ_o^h from (34). They are used almost exactly as in the zero-variance scheme.

We define the transition kernels following the zero-variance recipe (section 3.2.1). Keeping in mind $\psi_o^h(z_-(x, v)) = \psi_i^h(x, v)$, we have

$$k_{C^h}^h(z \rightarrow x) := \delta(x - x_+(z)),$$

so photons are cast from one boundary point to another with no absorption, exactly as they are in the continuous case. No discretization error occurs with casting. To define the scattering kernel $k_{S^h}^*$ we first recall the zero-variance kernel, which, since $\psi_o^s(z_-(x, v)) = \psi_i^s(x, v) = \varphi(x)$, takes the form:

$$k_{S^*}^*((x, v_{in}) \rightarrow v) = \alpha(x)\kappa(x, v) \frac{\varphi(x_+(x, v))}{\varphi(x)}.$$

We now discretize directions on every segment ∂X_j . We recall that x_j is the center of ∂X_j . Let V_{ij} be the set of directions best approximated by $v_{ij} := \widehat{x_j - x_i}$

$$V_{ij} := \{v \in \mathbb{S}^{d-1} : \arg \min_k |v - \widehat{x_k - x_i}| = j\}. \quad (35)$$

With $|V_{ij}|$ the measure of this set, we see that, roughly speaking, $N(x_i, x_j) \approx |V_{ij}|/|\partial X_j|$. With this in mind, we present our method for selecting direction. Suppose we are at $x'_i \in \partial X_i$, with incoming direction v_{in} . First we select a target region ∂X_j using using the discrete pdf

$$j \mapsto \alpha(x_i)\kappa(x_i, v_{ij})\partial_\nu N(x_i, x_j)|\partial X_j| \frac{\varphi^h(x_j)}{\varphi^h(x_i)}. \quad (36)$$

Notice that we use the grid center-point x_i instead of x'_i . So, for fixed $x'_i \in \partial X_i$ we may calculate the pdf by taking the component-wise product of the i^{th} row of

\mathcal{Q} with φ^h , and then divide by $(\varphi^h)_i$. Because of this and (33), the above discrete function does indeed sum (over j) to one, so it is a pdf. In a second step, v is selected from a uniform distribution on V_{ij} . This defines a direction leaving x_i , pointed into the domain, and toward ∂X_j . Note though that a shot leaving x'_i in direction v may not point into the domain since the normal vectors to ∂X at x'_i and x_i are not exactly equal. To correct for this, we introduce a family of rotation operators $\{\mathcal{R}_{x,y} : (x,y) \in \partial X \times \partial X\}$, such that $\mathcal{R}_{x_i,x'_i}(v) = v'$, where v' is a rotation of v that points from x'_i into the domain. In two dimensions, the obvious choice (which we use) of v' will ensure $v' \cdot \nu_{x'_i} = v \cdot \nu_{x_i}$. In three dimensions another reference vector (besides the normal) must be pre-selected at each point. Since both x_i and x'_i belong to ∂X_i , the operator \mathcal{R}_{x_i,x'_i} is close to the identity operator (this generates a “small” rotation).

We may now define our scattering transition kernel at arbitrary points by referring back to the kernel at grid-points. For $x'_i \in \partial X_i$, $v'_j \in V_{ij}$, our scattering transition kernel is

$$k_{S^h}^h((x'_i, v_{in}) \rightarrow v'_j) = k_{S^h}^h((x_i, \mathcal{R}_{x_i,x'_i}^{-1}(v_{in})) \rightarrow \mathcal{R}_{x_i,x'_i}^{-1}(v'_j)),$$

where for $v \in V_{ij}$

(37)

$$k_{S^h}^h((x_i, v_{in}) \rightarrow v) := \alpha(x_i) \kappa(x_i, v_{ij}) \partial_\nu N(x_i, x_j) \frac{|\partial X_j| \varphi^h(x_+(x_i, v_{ij}))}{|V_{ij}| \varphi^h(x_i)}.$$

To put ourselves in the framework (30) we define the discretized coefficients:

$$\alpha^h := R\alpha, \quad \bar{g}^h := Rh, \quad E_\sigma^h = E_\sigma \equiv 1, \quad (38)$$

$$\Theta^h(x_2, v_1 \rightarrow v_2) := k_{S^h}^h((x_2, v_1) \rightarrow v_2) \frac{1}{\alpha^h(x_2)} \frac{\varphi^h(x_2)}{\varphi^h(x_+(x_2, v_2))}. \quad (39)$$

Notice that if we ignore the rotation \mathcal{R} , we have (for $x'_i \in \partial X_i$, $v'_j \in V_{ij}$),

$$\Theta^h(x'_i, w \rightarrow v'_j) := \kappa(x_i, v_{ij}) \partial_\nu N(x_i, x_j) \frac{|\partial X_j|}{|V_{ij}|}.$$

However, in general, the ratio of φ^h will not cancel due to rotation. The Radon-Nikodym derivative $\left| \frac{dP^a}{dP^h} \right|$ is then given by (30).

The next lemma shows that this transition kernel is a pdf.

Lemma 3.2.

$$\int_{\nu_{x'_i} \cdot v' < 0} k_{S^h}^*((x'_i, v) \rightarrow v') dv' = 1 - R\bar{g}(x_i).$$

Proof. For $x'_i \in \partial X_i$,

$$\begin{aligned} \int_{\nu_{x'_i} \cdot v' < 0} k_{S^h}^*((x'_i, v_{in}) \rightarrow v') dv' &= \int_{\nu_{x'_i} \cdot v' < 0} k_{S^h}^*((x_i, \mathcal{R}_{x_i,x'_i}^{-1}(v_{in})) \rightarrow \mathcal{R}_{x_i,x'_i}^{-1}(v')) dv' \\ &= \int_{\nu_{x_i} \cdot v < 0} k_{S^h}^*((x_i, \mathcal{R}_{x_i,x'_i}^{-1}(v_{in})) \rightarrow v) dv. \end{aligned}$$

This follows since rotations preserve measure and $\{v' : \nu_{x'_i} \cdot \mathcal{R}_{x_i,x'_i}^{-1}(v') < 0\} = \{v : \nu_{x_i} \cdot v < 0\}$ by our choice of \mathcal{R} . To integrate the last term, we notice that

$k_{S^h}^*((x_i, w) \rightarrow v)$ is constant for $v \in V_{ij}$. Therefore, using (33),

$$\begin{aligned} \int_{\nu_{x_i} \cdot v < 0} k_{S^h}^*((x_i, w) \rightarrow v) dv &= \sum_{j=0}^{N_p-1} \alpha(x_i) \kappa(x_i, v_{ij}) \partial_\nu N(x_i, x_j) |\partial X_j| \frac{\varphi^h(x_j)}{\varphi^h(x_i)} \\ &= \frac{\sum_{j=0}^{N_p-1} Q_{ij} \varphi_j^h}{\varphi_i^h} = 1 - R\bar{g}, \end{aligned}$$

independent of the incoming direction w . □

In the best of cases, the SAI chain meets the hypothesis of theorem 3.5.

Theorem 3.6. *Assume that for h small enough, there exist $C, \rho > 0$ such that*

- (i) *We have the bounds $|\alpha/\alpha^h - 1| \leq Ch$, $|\Theta/\Theta^h - 1| \leq Ch$*
- (ii) *The boundary ∂X is C^3 and strictly convex*
- (iii) *$P^h[\tau = n] \leq Ce^{-\rho n}$ for some $\rho > 0$*
- (iv) *$\text{supp}(\varphi) \subset \text{supp}(\varphi^h)$*

Then the hypothesis of Theorem 3.5 are met and

$$\text{Var} \left\{ \xi^h \right\} \lesssim h^2. \quad (40)$$

The proof of the theorem can be found in Appendix A.3.

Remark 3.1. Assumptions (i) and (ii) are here to simplify the derivation of the result, which may hold in more general settings. In general, assumptions (i), (iv) (which together imply $\text{supp}(\varphi^h) = \text{supp}(\varphi)$) require our discrete mesh to be chosen to coincide well with the support of the coefficients. Moreover, assumption (i) requires that the rotation caused by \mathcal{R} causes little change in the value of φ^h . Smoothness assumptions on φ would provide this. Assumption (iii) means that multiple scattering is not dominant and holds in our model cases since we have $\alpha \equiv 0$ on the left/right sides and the sky.

Remark 3.2. Often physical coefficients such as the detector support are discontinuous and do not match up exactly with the grid. If $\bar{g}^h \equiv 1$ on its support, φ^h has a very large jump at the boundary of this support. This means that the mismatch in φ^h due to rotations will sometimes be large. Also, a non-convex boundary will cause issues at points where the curvature changes sign. These difficulties all occur at a finite number of points, and lead to an error contribution of $O(h)$ in (40).

Remark 3.3. Note that when $\Theta(x, v \rightarrow v') = 0$ for v or v' a grazing angle (i.e., $|v \cdot \nu_x|$ or $|v' \cdot \nu_x|$ close to 1), we verify that the rotations \mathcal{R} can be set to identity for small h and one can verify that $k_{S^h}^h((x_i, v_{in}) \rightarrow v)$ in (37) also generates a pdf for x'_i close to x_i .

3.4 The Surface Adjoint Importance (SAI) method

The SAI method is a modular method using the surface-adjoint approximation. In the absence of volume interaction, SAI becomes the zero-variance technique when $h \rightarrow 0$ that we saw in Theorem 3.6. In the presence of limited amounts of

volume scattering, we will show that SAI properly modified (a “heuristic module” is added) can be used for significant variance reduction and speedup.

For the rest of the paper, we define dP^h as the measure obtained by approximating the zero-variance chain in the *absence* of volume scattering, i.e. with setting $\sigma = 0$. It is thus defined via (30) with the coefficients given in (38). As we saw in section 3.3.1, solving the radiosity equation for the adjoint solution in the absence of volume interaction is much less costly than solving a full transport equation accounting for volume scattering.

When $\sigma > 0$, neglecting volume scattering as we did in our definition of dP^h causes problems and $\left| \frac{dP^a}{dP^h} \right|$ does not always exist. This occurs due to the fact that the analog chain sends some photons to the detector after experiencing volume interactions, but the chain generated by dP_s does not. Thus dP_s , or its approximation dP^h , cannot be used directly for variance reduction as they provide biased estimates of $\langle g, u \rangle$.

3.4.1 SAI-Heuristic Importance Sampling Scheme

For these reason, we propose the following regularized scheme: For $q = (q_v, q_s)$, with $q_v \in (0, 1]$, $q_s \in [0, 1]$, construct the measure

$$dP_q := (1 - q_s) dP^h + q_s dP_{heu, q_v}. \quad (41)$$

This means we will fire photons using heuristic (replacing the latter by any unbiased scheme would work) with probability q_s , and use $dP^h \approx dP_s$ with probability $1 - q_s$. When $q_s = 0$ we do not account for volume scattering and thus cannot obtain an unbiased estimator. When $q_v = 1$ we are using the SAI approximation combined with survival biasing.

The algorithm based on (41) is our main example of modular calculation of the adjoint solution. Here, volume and boundary scattering are uncoupled. When few particles undergo both volume and boundary scattering, then the above measure can have a very small variance. Here, we have defined dP^h as an approximation to the zero-variance measure dP_s if only surface scattering were present. The volume scattering dP_{heu, q_v} is still handled in a very crude fashion. A more accurate calculation of the adjoint solution accounting for volume scattering would provide larger variance reductions. In the presence of highly scattering clouds for instance, dP_{heu, q_v} would have to be replaced by a more accurate approximation of volume scattering. Yet, the structure of (41) would remain the same.

We then see that three parameters need to be chosen: h , q_s , and q_v . The regularizing parameters q_s , q_v should be chosen as a function of the mean free path. Ideally we could choose $q_s = 0$ (q_v has no effect then) when the MFP is infinite. With finite MFP we must use $q_s > 0$ (in fact close to one, even when $MFP \approx 16$ times the domain diameter). As MFP decreases, both q_s and q_v should decrease to allow for more analog shots that account for complex volume or volume + boundary interactions. The parameter h should then be chosen to maximize the figure of merit: Small values of h generate large computational cost (due to the expense of the deterministic solve) for limited variance reductions since a significant variance comes from shots that interact with the volume. Simulations show that very small values of h yield no measurable improvement in variance.

Note that the asymptotic regime is no longer a good description of the above method. Instead, different subsets of paths to the detector are chosen, and different methods are used to increase the probability of their occurrence. Figure 3 shows both boundary and volume photons being directed toward the detector.

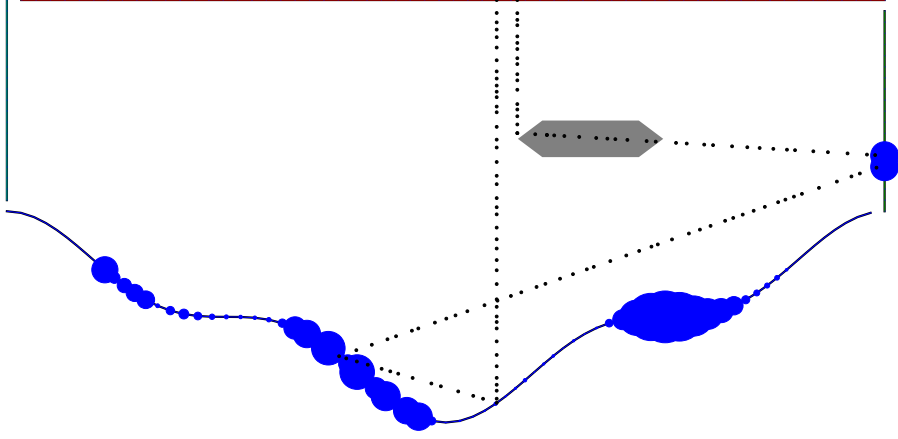


Figure 3: Boundary and volume interactions handled by different modules

The boundary photon was directed using dP^h . The relative size of the adjoint solution on the boundary is indicated by relative dot size. Note that the adjoint solution allows us to account for complex boundary interactions.

The details of the SAI algorithm are as follows. First we produce N draws $\{\omega^i\}_{i=1}^N$ from P_q using algorithm 3, then estimate $\langle u, g \rangle \approx N^{-1} \sum_{i=1}^N \xi^q(\omega^i)$ where $\xi^q = \xi_a \left| \frac{dP^a}{dP_q} \right|$ (expression derived below). When we draw from $k_{S^h}^*$, k_{C^h} in algorithm 3, shots are never absorbed until they reach the support of the discretized detector \bar{g}^h .

Algorithm 3 SAI

- 1: With probability $1 - q_s$, set *switch* \leftarrow **true**
 - 2: **if** *switch* **then**
 - 3: Draw z_0 from a density $\propto s(z)\psi_o^h(z)$
 - 4: **while** $x_j \notin \text{supp}(\bar{g}^h)$ **do**
 - 5: Draw $x_{j+1} \sim k_{C^h}^h(z_j \rightarrow \cdot)$ (In this case we simply set $x_{j+1} \leftarrow x_+(z_j)$)
 - 6: Draw $v_{j+1} \sim k_{S^h}^*(x_{j+1}, v_j \rightarrow \cdot) \propto \Theta^h(x_{j+1}, v_j \rightarrow \cdot)\psi_o^h(x_{j+1}, \cdot)$
 - 7: Set $j \leftarrow j + 1$
 - 8: **end while**
 - 9: Set $v_j \leftarrow \bar{\mathfrak{d}}$
 - 10: **else**
 - 11: Draw $\omega \sim dP_{heu, q_v}$ using algorithm 2
 - 12: **end if**
-

We now derive an expression for $\left| \frac{dP^a}{dP_q} \right|$. Whenever $dP^h = 0$ (say the photon had a volume interaction), $\left| \frac{dP^a}{dP_q} \right|$ (restricted to $(x_\tau, v_{\tau-1}) \in \text{supp}(\bar{g})$) is given by

$$\left| \frac{dP^a}{dP_q} \right| = \frac{1}{q_s} \left| \frac{dP^a}{dP_{heu, q_v}} \right| = \frac{1}{q_s} \frac{\left| \frac{dP^a}{dP_{sb}} \right|}{\left| \frac{dP_{heu, q_v}}{dP_{sb}} \right|},$$

where $\left| \frac{dP^a}{dP_{sb}} \right|$ is given by (22), and $\left| \frac{dP_{heu, q_v}}{dP_{sb}} \right|$ is given by (24). When $dP^h \neq 0$,

we have

$$\left| \frac{dP^a}{dP^q} \right| = \frac{\left| \frac{dP^a}{dP^h} \right|}{(1 - q_s) + q_s \left| \frac{dP^{heu, q_v}}{dP^h} \right|} = \frac{\left| \frac{dP^a}{dP^h} \right|}{(1 - q_s) + q_s \left| \frac{dP^{heu, q_v}}{dP^{sb}} \right| \left| \frac{dP^{sb}}{dP^h} \right|}.$$

So we need expressions for $\left| \frac{dP^a}{dP^h} \right|$ and $\left| \frac{dP^{sb}}{dP^h} \right| = \left| \frac{dP^{sb}}{dP^{sb}} \right| \left| \frac{dP^a}{dP^h} \right|$. Note that $\left| \frac{dP^a}{dP^h} \right|$ is given by (30), but simplifies since $E^h \equiv 1$, and the fact that when $dP^h \neq 0$ we have necessarily taken a path such that all $x_i \in \partial X$, which makes the expression for $\gamma_{a,h}$ (a term in $\left| \frac{dP^a}{dP^h} \right|$) “simple”. Also note that $\left| \frac{dP^{sb}}{dP^{sb}} \right|$ is given by (22) and that $\gamma_{a,sb}$ is “simple” for the same reason $\gamma_{a,h}$ was. We therefore have

$$\begin{aligned} \left| \frac{dP^a}{dP^h} \right| &= \frac{\langle s^h, \psi_o^h \rangle}{\bar{g}^h(x_\tau, v_\tau)} \frac{s(z_0)}{s^h(z_0)} E_{\sigma}(x_0, \dots, x_\tau) \\ &\quad \times \frac{\alpha(x_1)\Theta(x_1, v_0 \rightarrow v_1) \cdots \alpha(x_{\tau-1})\Theta(x_{\tau-1}, v_{\tau-2} \rightarrow v_{\tau-1})}{\alpha^h(x_1)\Theta^h(x_1, v_0 \rightarrow v_1) \cdots \alpha^h(x_{\tau-1})\Theta^h(x_{\tau-1}, v_{\tau-2} \rightarrow v_{\tau-1})}, \\ \left| \frac{dP^{sb}}{dP^h} \right| &= \frac{\langle s^h, \psi_o^h \rangle}{\bar{g}^h(x_\tau, v_\tau)} \frac{s(z_0)}{s^h(z_0)} E_{\sigma_s}(x_0, \dots, x_\tau) \\ &\quad \times \frac{\alpha^{sb}(x_1)\Theta(x_1, v_0 \rightarrow v_1) \cdots \alpha^{sb}(x_{\tau-1})\Theta(x_{\tau-1}, v_{\tau-2} \rightarrow v_{\tau-1})}{\alpha^h(x_1)\Theta^h(x_1, v_0 \rightarrow v_1) \cdots \alpha^h(x_{\tau-1})\Theta^h(x_{\tau-1}, v_{\tau-2} \rightarrow v_{\tau-1})}. \end{aligned}$$

Even though these expressions are complicated to write explicitly, we emphasize that their computational cost is rather minimal compared to the overall cost of solving a transport equation by Monte Carlo.

3.4.2 Optimal parameter selection

Here we outline two procedures to pick values of q_s close to optimal. To simplify, we assume that q_v is fixed in the heuristic module with dP^{heu, q_v} . As before, we assume that $\bar{g} \equiv 1$ on its support so that $\xi = \mathbb{1}_D$ although general \bar{g} could be handled with additional hypotheses.

Note that

$$\left| \frac{dP^a}{dP^q} \right| = \left[(1 - q_s) \left| \frac{dP^h}{dP^a} \right| + q_s \left| \frac{dP^{heu, q_v}}{dP^a} \right| \right]^{-1},$$

and that the quantity to minimize with respect to q_s is thus

$$\mathbb{E}_q \{ (\xi^q)^2 \} = \mathbb{E}_q \left\{ \mathbb{1}_D^2 \left| \frac{dP^a}{dP^q} \right|^2 \right\} = \int_{\Omega} \mathbb{1}_D \left| \frac{dP^a}{dP^q} \right| dP^a.$$

We split the above into integrals over B and $D \setminus B$, where B is the set of paths of particles that reach the detector without undergoing volume scattering (but can have many interactions with the boundary). We assume for simplicity that $\text{supp}(P^h) = B$, i.e., that discretization effects do not significantly modify the support of P_s (otherwise, B should be thought as the support of P^h).

On B , we find that for any subset $B' \subset B$, we have $P^h(B') \gg P^{heu, q_v}(B') = P^a(B')$ (at least when neglecting discretization effects). The reason is that the paths reaching the detector after interacting with the boundary have very high probability density dP^h (this is exactly the role of dP^h : sending particles interacting with the boundary toward the detector). However, for such paths,

$P_{heu,q_v}(B') = P^a(B')$ since heuristic sampling only modifies those paths that undergo volume scattering. For $\omega \in B$, we thus find that

$$\left| \frac{dP^a}{dP_q} \right|(\omega) = \frac{1}{1 - q_s} \left| \frac{dP^a}{dP^h} \right|(\omega) - \varepsilon(\omega),$$

with $0 \leq \varepsilon(\omega) \ll 1$. On $D \setminus B$, $dP^h = 0$ since paths with volume scattering have vanishing weight under the boundary measure dP^h . As a consequence, we have that

$$\mathbb{E}_q \{ (\xi^q)^2 \} \approx \frac{1}{1 - q_s} \int_{B \cap D} \left| \frac{dP^a}{dP^h} \right| dP^a + \frac{1}{q_s} \int_{D \setminus B} \left| \frac{dP^a}{dP_{heu,q_v}} \right| dP^a. \quad (42)$$

The above can be optimized over q_s once the two integrals are known. Since they are both expectations (with respect to dP^a), we can estimate them with an analog simulation, or send particles using dP_q , and use importance sampling. In this way our optimal choice of q_s can be refined as more particles are sent. We find

$$(q_s)_{\text{opt},1} \approx \frac{\sqrt{\alpha}}{1 + \sqrt{\alpha}}, \quad \alpha := \left(\int_{D \setminus B} \left| \frac{dP^a}{dP_{heu,q_v}} \right| dP^a \right) \left(\int_{B \cap D} \left| \frac{dP^a}{dP^h} \right| dP^a \right)^{-1}.$$

Alternatively, and because calculating the integrals in the definition of α is still difficult, we may obtain another guess using only *a priori* estimates of $P^a[D]$ and $P^a[B | D]$. This can be done in the simplified importance sampling framework of section 3.1.2, specifically the regime (16).

First we define $b \in \mathbb{R}$ as the constant that makes $P_q[B] = bP^a[B]$. Second, we recall that $P^a[B] = P_{heu,q_v}[B]$. Thus, using (41), we find

$$bP^a[B] = P_q[B] = (1 - q_s)P^h[B] + q_sP^a[B].$$

Using the approximation $P^h[B] \approx 1$ (neglecting discretization effects), we have

$$b \approx \frac{1 - q_s}{P^a[B]} + q_s,$$

and thus

$$q_s \approx \frac{1 - bP^a[B]}{1 - P^a[B]} = \frac{1 - (bP^a[D])P^a[B | D]}{1 - P^a[D]P^a[B | D]}. \quad (43)$$

Assuming that $dP_q(\omega) = b dP^a(\omega)$ on B , we are approximately in the regime (15). We thus choose $b_{\text{opt}} = P^a[D]^{-1} \beta_{\text{opt}}$ minimizing (16) with $\beta_{\text{opt}} = (\sqrt{\gamma}(\sqrt{\gamma} + \sqrt{a}))^{-1}$, where a and γ are defined in (17). Then, we find $(q_s)_{\text{opt},2}$ in terms of b_{opt} using (43). See section 4.3 for an implementation of this algorithm.

4 Numerical Results

In this section, we implement the scheme described in the preceding section and sample chains numerically based on the measure dP_q for several values of q and h . We compare the variance of the method with the survival biasing measure dP_{sb} . Several details of the implementation are described in section A.4.

The rotation \mathcal{R} described in section 3.3.2 were found to have an extremely limited effect on the calculated solutions. Even with coarse grids, neglecting the

rotations (setting them to the identity matrix) led to less than 0.1% bias (the bias was so small that it could have been error due to not firing enough shots). See also remark 3.3. So the presented results are obtained with the rotations set to identity.

We first introduce the notion of speedup (a.k.a. figure of merit) in section 4.1. We consider two speedups depending on whether the cost of the deterministic adjoint solution is included or not. Then in section 4.2, we show the influence of the discretization parameter h on the convergence of the variance to 0 (and the speedup to infinity) in the absence of volume scattering ($\sigma \equiv 0$) and compare the numerical results with theoretical predictions. Finally, in section 4.3, we include volume scattering and obtain significant variance reductions by appropriate choice of the regularization parameters (q_s, q_v) . Moreover, we show that large speedups are obtained for a relatively large band of values of (q_s, q_v) , whose optimal values very much depends on geometry/scattering/absorption and has to be obtained fairly empirically.

4.1 Speedup (figure of merit)

For all of these methods, define the approximation after N random draws

$$I_N(\xi) := \frac{1}{N} \sum_{n=1}^N \xi(\omega_n).$$

The RMS estimation error ε is given by

$$\varepsilon(\xi) := \sqrt{\mathbb{E} \{|I_N(\xi) - \langle u, g \rangle|^2\}} = \sqrt{\frac{\text{Var} \{\xi\}}{N}}.$$

For a given error level ε , the required number of MC draws is then $N(\varepsilon, \xi) := \text{Var} \{\xi\} / \varepsilon^2$. The required simulation time $T(\varepsilon, \xi)$ for one estimation of $\mathbb{E} \{\xi\}$ is given by

$$T(\varepsilon, \xi) := T_0(\xi) + \tau(\xi)N = T_0(\xi) + \frac{\tau(\xi)\text{Var} \{\xi\}}{\varepsilon^2},$$

where $T_0(\xi)$ is the time needed to compute the deterministic adjoint solution (e.g. at level h when $\xi = \xi^h$), and $\tau(\xi)$ is the expected time for one draw using the appropriate measure for the random variable ξ . We foresee the use of SAI in situations where the boundary remains fixed, but the volume changes (due to e.g. moving clouds over a fixed surface). We therefore consider the time for m simulations using one boundary,

$$T(\varepsilon, \xi, m) := T_0(\xi) + m\tau(\xi)N = T_0(\xi) + m\frac{\tau(\xi)\text{Var} \{\xi\}}{\varepsilon^2},$$

Then we compare schemes through the ‘‘Speedup.’’

$$\text{Speedup}(\xi_1, \xi_2, \varepsilon, m) := \frac{T(\varepsilon, \xi_1, m)}{T(\varepsilon, \xi_2, m)} = \frac{\varepsilon^2 T_0(\xi_1) + m\tau(\xi_1)\text{Var} \{\xi_1\}}{\varepsilon^2 T_0(\xi_2) + m\tau(\xi_2)\text{Var} \{\xi_2\}}.$$

For a deterministic approximation of ψ_i^s , we expect $T_0(\xi) \approx C(\xi)h^{-2(d-1)}$. We in fact measure (with $d = 2$) $T_0(\xi^h) \approx 0.017h^{-2}$. Our ‘‘benchmark’’ scheme is survival biasing. Since ξ_{sb} requires no deterministic solution, the relevant ratio is

$$\text{Speedup}(\xi_{sb}, \xi^q, \varepsilon, m) = \frac{m\tau(\xi_{sb})\text{Var} \{\xi_{sb}\}}{\left(\frac{\varepsilon}{h}\right)^2 C + m\tau(\xi^q)\text{Var} \{\xi^q\}}.$$

We measured speedup when either $m = 10$ or $m = \infty$ (‘‘Ignoring deterministic solve’’).

4.2 Variance reduction without volume interactions

When the volume mean-free-path is infinite, we can show that the variance approaches zero as $h \rightarrow 0$.

First consider the case of a flat boundary. Photons leave the sky, hit the boundary, then either reach the detector or are “absorbed” by the sky or sides. Neglecting edge/detector overlap we are in the regime of assumptions 3.1. Therefore, combining lemma 3.1 with theorem 3.5 we expect approximately $O(h^2)$ convergence. In practice we observed $O(h^{1.6})$ convergence. See figure 4.

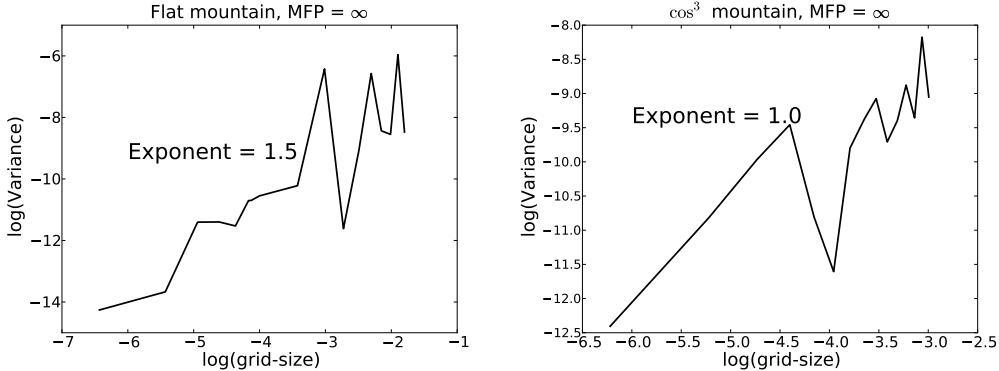


Figure 4: $O(h^\alpha)$ variance behavior. Complicated \cos^3 mountain results in slower convergence.

Second, when the more complex \cos^3 boundary of figure 3 is used, we require $q_s > 0$ for the following reason: Suppose a photon finds itself at the point $x = (-0.4, 1 + \cos^3 0.4)$. On a fine boundary, there is a point nearby that has a direct line to the detector. Therefore, when $\Theta^h(x, v' \rightarrow v)$ will allow for shots directly to the detector. One can see this by noticing that in figure 4.2 the point $(-0.4, 1 + \cos^3 0.4)$ is shaded darkly in the fine boundary (right), indicating that it sees direct illumination from the detector. On a coarse boundary (left) this is not the case. In practice we observed a major contribution to variance due to these effects, and $O(h)$ convergence overall. See figure 4 and also Remark 3.2. Also, (not pictured) we observe that with a flat boundary and fine discretization is used, the optimal regularization parameter is $q_s = 0$. When a coarse discretization or \cos^3 boundary is used, the optimal $q_s \neq 0$.

4.3 Variance reduction with volume interactions

To analyze the variance of the SAI chain in the presence of volume interactions, we adopt the modularity viewpoint explained in section 3.1.2. Note that even when the error $\langle u, g \rangle - \langle \psi_o^h, s^h \rangle$ is high, we still get good variance reduction. See figure 6. This emphasizes the point that the quality of the deterministic solve is not so important in a modular scheme.

Our implementation swept both q_s and q_v . As expected, we see decreasing speedup with increasing volume scattering strength σ . See figure 7.

It is important to note that use of adjoint-enhanced surface scattering, and heuristic volume scattering ($q_s < 1, q_v < 1$) together is especially helpful. In fact, even with a small $\text{MFP} = 1.3 \times \text{Diameter}$, we realize good speedup when $q_s = 0.9$,

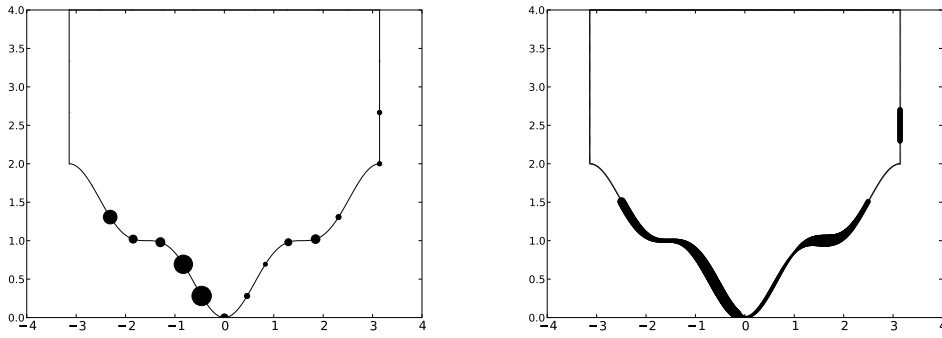


Figure 5: Boundaries discretized on coarse (left) and fine (right) scales. Dots indicate adjoint flux at mesh points. Size is relative to flux strength.

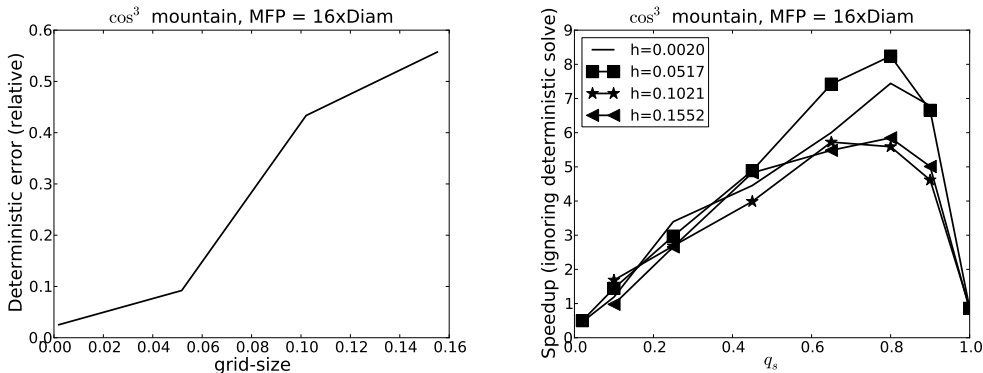


Figure 6: $|\langle u, g \rangle - \langle s^h, \psi_o^h \rangle| / \langle u, g \rangle$ is generally lower for smaller h . However, speedup is still very good even for large h .

$q_v < 1$. Note that if either $q_s = 1$ or $q_v = 1$ (so no use of dP^h or heuristic scattering adjustment), speedup almost disappears. Following the setting described in section 3.1.2 and the discussion at the end of section 3.1.3, this may be explained as follows. Assuming that D is well approximated by $B_1 \cup B_2$ where B_1 and B_2 have approximately the same size and $B_1 \cap B_2 \approx \emptyset$. Then the maximal variance obtained by choosing only B_1 or only B_2 in the importance sampling is roughly a factor 2 whereas the maximal variance obtained by choosing *both* of them is very large. When the computational cost of the deterministic solve is taken into account, we obtain the results in figure 7, which show that both boundary and volume scattering need to be accelerated in order to obtain significant speedups.

4.4 Optimization of the parameter q_s

We now compare the *a priori* estimates of an optimal q_s (computed using the methodology in section 3.4.2) with the observed optimal values (from figure 7). To compute the *a priori* estimates we need estimates for $P^a[D]$, $P^a[B | D]$. First we assume $P^a[B | D] \approx 1 - P^a[V]$, where the set V are the photons that had a volume interaction before dying. Since $P^a[V]$ is rather large, it is easy to estimate with a very short analog simulation. We found that $MFP = 16 \times \text{Diameter}$ corresponded

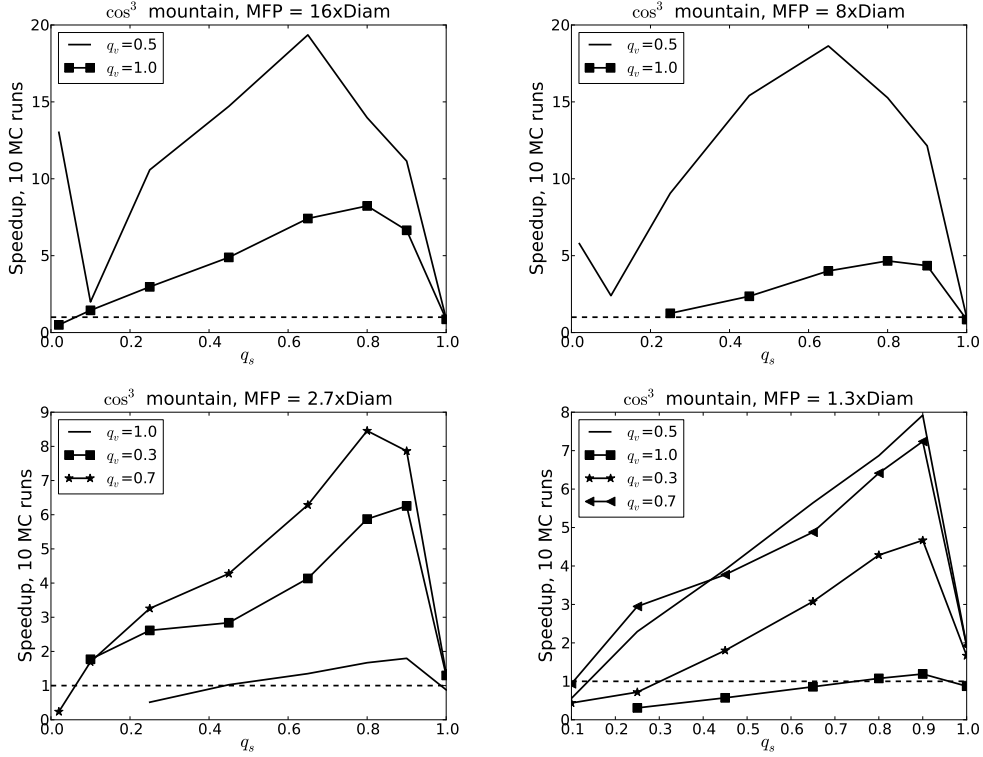


Figure 7: Speedup when using both surface adjoint approximation ψ_o^h (with parameter q_s) and heuristic volume scattering (with parameter q_v)

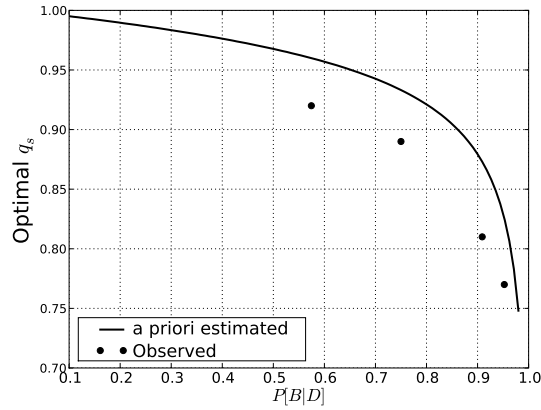


Figure 8: *A priori* Optimal q_s (*a priori* estimate and observed) for q_s is plotted for a variety of $P^a[B|D]$.

to $P^a[V] = 1/21$, and $MFP = (8, 2.7, 1.3) \times \text{Diameter}$ corresponded to $P^a[V] = (1/11, 1/4, 1/2.35)$ respectively. This gives us estimates of $P^a[B|D]$. Now we need an estimate of $P^a[D]$. The true values lie in the range $[0.002325, 0.002484]$. This can be estimated fairly quickly (to within 10% RMS error) using 25,000 survival biased shots. In figure 8 we plot the *a priori* optimal q_s versus $MFP/\text{Diameter}$ using the above estimate for $P^a[B|D]$ and setting $P^a[D] = 0.0024$ (10% errors in $P^a[D]$ make very little difference).

A Appendix

This section collects details left out in the preceding sections.

A.1 Proof of theorem 3.1

Proof of theorem 3.1. For a proof of the theorem in the absence of a boundary, we refer the reader to [16]. The analog probability density is defined in (12). If $Y : \Omega \rightarrow \mathbb{R}$ is a random variable, then

$$\mathbb{E}_a \{Y\} = \sum_{n=1}^{\infty} \int_{\tau=n} Y \, dP^a = \sum_{n=1}^{\infty} \mathbb{E}_a \{Y \mathbb{1}_{\tau=n}\}, \quad (44)$$

where, following the structure in algorithm 1, we have

$$\begin{aligned} \mathbb{E}_a \{Y \mathbb{1}_{\tau=n}\} &= \mathbb{E}_a \{\mathbb{E}_a \{Y \mathbb{1}_{\tau=n} \mid Z_{n-1}\}\} \\ &= \mathbb{E}_a \{\mathbb{E}_a \{\mathbb{E}_a \{Y \mathbb{1}_{\tau=n} \mid Z_{n-1}, Z_{n-2}\} \mid Z_{n-2}\}\} \\ &= \mathbb{E}_a \{\cdots \mathbb{E}_a \{\mathbb{E}_a \{Y \mathbb{1}_{\tau=n} \mid Z_{n-1}, \dots, Z_0\} \mid Z_{n-2}, \dots, Z_0\} \cdots \mid Z_0\}. \end{aligned}$$

Using (5), (7), and (44), it will suffice to show

$$\mathbb{E}_a \{\xi_a \mathbb{1}_{\tau=n}\} = \langle C^* \bar{g}, (KL)^{n-1} s \rangle, \quad n = 1, 2, \dots$$

First note that (since g is a boundary source extended to be zero off of Γ_+)

$$\begin{aligned} \mathbb{E}_a \{\xi_a \mathbb{1}_{\tau=n} \mid Z_{n-1}\} &= \int_{\bar{X}} \frac{\bar{g}(x_n, V_{n-1})}{p_{S^*}^a(x_n)} k_{C^*}^a(Z_{n-1} \rightarrow x_n) p_{S^*}^a(x_n) \, dx_n \\ &= \bar{g}(x_+(Z_{n-1}), V_{n-1}) E_\sigma(X_{n-1}, x_+(Z_{n-1})) = C^* \bar{g}(Z_{n-1}). \end{aligned}$$

So when $n = 1$, we have

$$\begin{aligned} \mathbb{E}_a \{\xi_a \mathbb{1}_{\tau=1}\} &= \mathbb{E}_a \{\mathbb{E}_a \{\xi_a \mathbb{1}_{\tau=1} \mid Z_0\}\} \\ &= \mathbb{E}_a \{C^* \bar{g}(Z_0)\} = \int_{\bar{Z}} s(z_0) C^* \bar{g}(z_0) \, dz_0 = \langle C^* \bar{g}, s \rangle_{\Gamma_-}. \end{aligned}$$

Next note that for $m < \tau$,

$$\mathbb{E}_a \{f(Z_m) \mid Z_{m-1}\} = \int_{\bar{Z}} f(z) k_{T^*}^a(Z_{m-1} \rightarrow z_m) \, dz_m = C^* S^* f(Z_{m-1}).$$

So when $n > 1$, we have

$$\begin{aligned} \mathbb{E}_a \{\xi_a \mathbb{1}_{\tau=n}\} &= \mathbb{E}_a \{\mathbb{E}_a \{\xi_a \mathbb{1}_{\tau=n} \mid Z_{\tau-1}\}\} = \mathbb{E}_a \{C^* \bar{g}(Z_{\tau-1})\} \\ &= \mathbb{E}_a \{\mathbb{E}_a \{C^* \bar{g}(Z_{\tau-1}) \mid Z_{\tau-2}\}\} = \mathbb{E}_a \{(C^* S^*) C^* \bar{g}(Z_{\tau-2})\} \\ &\quad \vdots \\ &= \mathbb{E}_a \{(C^* S^*)^{n-1} C^* \bar{g}(Z_0)\} = \int_{\bar{Z}} s(z_0) (C^* S^*)^{n-1} C^* \bar{g}(z_0) \, dz_0 \\ &= \langle s, (C^* S^*)^{n-1} C^* \bar{g} \rangle = \langle (KL)^{n-1} s, C^* \bar{g} \rangle. \end{aligned}$$

This proves the theorem. \square

A.2 Proof of theorem 3.5

We note that

$$\xi^h = \bar{g}(x_\tau, v_{\tau-1}) \left| \frac{d\mathbb{P}^a}{d\mathbb{P}^h} \right| = \langle u, g \rangle (1 + \varepsilon_h),$$

$$1 + \varepsilon_h(\omega) := \frac{\langle s^h, \psi_o^h \rangle}{\langle u, g \rangle} \frac{\bar{g}(x_\tau, v_{\tau-1})}{\bar{g}^h(x_\tau, v_{\tau-1})} \frac{s(z_0)}{s^h(z_0)} \beta_{a,h}(x_0, \dots, x_\tau) \gamma_{a,h}(z_1, \dots, z_{\tau-1}).$$

We now bound the coefficient error ε_h . First, assumptions 3.1 (i), (ii) give us

$$(1 - Ch)^{\tau+1} \leq 1 + \varepsilon_h \leq (1 + Ch)^{\tau+1}.$$

Using the binomial theorem, we have (for $x > 0$, $m \in \mathbb{N}$)

$$1 - mx e^{mx} \leq (1 - x)^m \leq (1 + x)^m \leq e^{mx}.$$

Therefore

$$(1 - Ch)^{\tau+1} \leq 1 + \varepsilon_h \leq (1 + Ch)^{\tau+1},$$

and thus

$$|\varepsilon_h| \leq hC[\tau + 1]e^{hC[\tau+1]}.$$

Since we assume $\mathbb{P}^h[\tau = n] \leq Ce^{-\rho n}$, we have

$$\begin{aligned} \text{Var} \left\{ \xi^h \right\} &= \sum_{n=0}^{\infty} \int_{\tau=n}^{\infty} (\xi^h - \langle u, g \rangle)^2 d\mathbb{P}^h = \langle u, g \rangle^2 \sum_{n=0}^{\infty} \int_{\tau=n}^{\infty} |\varepsilon_h|^2 d\mathbb{P}^h \\ &\leq \langle u, g \rangle^2 h^2 C^2 \sum_{n=0}^{\infty} [\tau + 1]^2 e^{-(\rho\tau - 2hC[\tau+1])}. \end{aligned}$$

So that, for $h < \rho/(2C)$ the above series converges and the result is proved.

A.3 Proof of some technical results

Proof of Lemma 3.1. Clearly $\partial_\nu N$ is C^{k+2} when $|x' - x| > 0$, so we may restrict our attention to $|x' - x| < \varepsilon$. We prove the lemma then for x', x both in an ε neighborhood of some point. After possibly shrinking ε , we may assume that in this neighborhood ∂X is the graph of a C^{k+2} function f . In other words, with $x = (x_1, x_2)$, and after a rotation and/or translation, this neighborhood is the set $\{(x_1, f(x_1)) : -\varepsilon < x_1 < \varepsilon\}$.

We then have

$$\nu_x = \frac{(-f'(x_1), 1)}{\sqrt{1 + (f'(x_1))^2}}, \quad x - x' = (x_1 - x'_1, f(x_1) - f(x'_1)),$$

and also

$$\begin{aligned} f(x'_1) &= f(x_1) + f'(x_1)(x'_1 - x_1) + R(x_1, x'_1), \\ R(x_1, x'_1) &:= \int_{x_1}^{x'_1} \int_{x_1}^s f''(t) dt ds = (x'_1 - x_1)^2 \int_0^1 \int_0^s f''(t(x'_1 - x_1)) dt ds. \end{aligned}$$

We notice that

$$R(x_1, x_1')(x_1' - x_1)^{-j} \in C^k(\partial X \times \partial X), \quad j = 0, 1, 2.$$

This is all we need since

$$\begin{aligned} \nu_x \cdot (x' - x) &= \frac{R(x_1, x_1')}{\sqrt{1 + (f'(x_1))^2}}, \\ |x - x'|^2 &= (x_1 - x_1')^2 \left[1 + \left(f'(x_1) + \frac{R(x_1, x_1')}{x_1' - x_1} \right)^2 \right], \\ \partial_\nu N(x, x') &:= \frac{\nu_x \cdot (x - x')}{|x - x'|^2}. \end{aligned}$$

□

Proof of theorem 3.6. Our setup so far puts us in the regime of section 3.3.1 with

$$\begin{aligned} E_\sigma^h &= E_\sigma \equiv 1, \quad \bar{g}^h = R\bar{g}, \quad \alpha^h = R\alpha, \\ \Theta^h(x_i', v \rightarrow v_{ij}') &= \kappa(x_i, v_{ij}) \partial_\nu N(x_i, x_j) \frac{|\partial X_j|}{|V_{ij}|}, \quad \text{for } x_i' \in \partial X_i, v_{ij}' \in V_{ij}. \end{aligned}$$

Assumptions (iii), (iv) are the same above and in assumptions 3.1. Using assumption (i) above along with proposition 3.1, we have assumption 3.1 (i). It remains to prove that assumptions 3.1 (ii) is met. Due to assumption (i) above, it will suffice to show

$$1 - C'h \leq \frac{|V_{ij}|}{|\partial X_j|} \frac{1}{\partial_\nu N(x_i, x_j)} \leq 1 + C'h. \quad (45)$$

Due to strict convexity of ∂X , $\partial_\nu N$ is bounded below. Now the differentiability of $\partial_\nu N$ (lemma 3.1) implies that there exists $C' > 0$ such that when $x' \in \partial X_j$,

$$1 - C'h \leq \frac{\partial_\nu N(x_i, x')}{\partial_\nu N(x_i, x_j)} \leq 1 + C'h. \quad (46)$$

Therefore, since

$$|V_{ij}| = \int_{\partial X_j} \partial_\nu N(x_i, x') \, d\mu(x') = \partial_\nu N(x_i, x_j) \int_{\partial X_j} \frac{\partial_\nu N(x_i, x')}{\partial_\nu N(x_i, x_j)} \, d\mu(x'),$$

(46) now implies (45) and the proposition is proved. □

A.4 Parameter choices in numerical simulations

In the simulations performed with $\sigma = 0$ (no volume interactions), we used both a flat surface (so that our domain was $[-\pi, \pi] \times [2, 4]$) and a \cos^3 surface (figure 3). We swept h , with $0.002 < h < 0.2$. We did not use any heuristic scattering adjustment ($q_v = 1.0$). In all cases

$$\Theta(x, y, v \rightarrow v') = \begin{cases} (\nu_x \cdot v')/2, & \nu_x \cdot v' < 0, |x| < 2.5 \\ 0, & \text{otherwise.} \end{cases}$$

The source was mono-directional $v = -\pi/2$ and given by

$$s(x, -\pi/2) = \begin{cases} 1 & |x| < 2.5, \\ 0 & |x| \geq 2.5. \end{cases}$$

In the simulations involving volume interactions ($\sigma > 0$), we used a \cos^3 type surface. We computed speedup in a variety of cases. The mean-free-path $\text{MFP} = \sigma^{-1}$ was varied as well as q_s , h , and q_v . We swept $0.002 < h < 0.15$. In all cases the volume scattering coefficients were constant with $\sigma_s = 2\sigma_a$. The volume scattering was given by

$$\theta(x, v \rightarrow v') = 1 + (v \cdot v')^2.$$

The other coefficients were chosen to have features (in this case oscillations) on a scale coarser than the fine values of h , and finer than the coarse values. The surface scattering coefficient was given (on the mountain) by

$$\Theta((x, y), v \rightarrow v') = (\nu_x \cdot v') \begin{cases} 0 & x > 2.5, \\ 0.75 + 0.25 \sin(2\pi x/0.05) & 1 < x < 2.5, \\ 0.35 + 0.25 \sin(2\pi x/0.05) & -2.5 < x < 1, \end{cases}$$

when $\nu_x \cdot v' > 0$, and 0 when $\nu_x \cdot v' \leq 0$. Off the mountain there was no scattering (perfectly absorbing). The source was mono-directional $v = -\pi/2$ and given by

$$s(x, -\pi/2) = \begin{cases} 1 + 0.25 \sin(2\pi x/0.07) & |x| < 2.5, \\ 0 & |x| \geq 2.5. \end{cases}$$

Acknowledgment

The authors would like to thank Anthony Davis for many useful discussions. This work was supported in part by DOE grant DE-FG52-08NA28779 and NSF grant DMS-0804696.

References

- [1] S. R. Arridge. Optical tomography in medical imaging. *Inverse Problems*, 15:R41–R93, 1999.
- [2] G. Bal. Inverse transport theory and applications. *Inverse Problems*, 25:053001, 2009.
- [3] R. E. Caffisch. Monte carlo and quasi-monte carlo methods. *Acta Numerica*, pages 1–49, 1998.
- [4] S. Chandrasekhar. *Radiative Transfer*. Dover Publications, New York, 1960.
- [5] R. Dautray and J.-L. Lions. *Mathematical Analysis and Numerical Methods for Science and Technology. Vol.6*. Springer Verlag, Berlin, 1993.
- [6] B. Davison and J. B. Sykes. *Neutron Transport Theory*. Oxford University Press, 1957.
- [7] J. D. Densmore and E. W. Larsen. Variational variance reduction for particle transport eigenvalue calculations using monte carlo adjoint simulation. *J. of Comp. Physics*, 192:387–405, 2003.
- [8] G. Folland. *Introduction to partial differential equations*. Princeton University Press, Princeton New Jersey, 1995.
- [9] A. Haghighat and J. C. Wagner. Monte carlo variance reduction with deterministic importance functions. *Prog. in Nuclear Energy*, 42 (1):25–53, 2003.

- [10] Ambrose M. Kong, R. and J. Spanier. Efficient, automated monte carlo methods for radiation transport. *Journal of Computational Physics*, 227:9643–9476, 2008.
- [11] R. Kong and J. Spanier. A new proof of geometric convergence for general transport problems based on sequential correlated sampling methods. *Journal of Computational Physics*, 227:9762–9777, 2008.
- [12] K. N. Liou. *An introduction to atmospheric radiation*. Academic Press, San Diego, CA, 2002.
- [13] I. Lux and L. Koblinger. *Monte Carlo Particle Transport Methods: Neutron and Photon Calculations*. CRC Press, Boca Raton, 1991.
- [14] A. Marshak and A. B. Davis. *3D Radiative Transfer in Cloudy Atmospheres*. Springer, New-York, 2005.
- [15] H. Sato and M. C. Fehler. *Seismic wave propagation and scattering in the heterogeneous earth*. AIP series in modern acoustics and signal processing. AIP Press, Springer, New York, 1998.
- [16] J. Spanier and E. M. Gelbard. *Monte Carlo principles and neutron transport problems*. Addison-Wesley, Reading, Mass., 1969.
- [17] S. A. Turner and E. W Larsen. Automatic variance reduction for three-dimensional monte carlo simulations by the local importance function transform–i: Analysis. *Nuclear science and engineering*, 127:22–35, 1997.
- [18] S. A. Turner and E. W Larsen. Automatic variance reduction for three-dimensional monte carlo simulations by the local importance function transform–ii: Numerical results. *Nuclear science and engineering*, 127:36–53, 1997.
- [19] K. A. Van Riper et al. Avatar - automatic variance reduction in monte carlo calculations. 1997.
- [20] E. Veach. Robust monte carlo methods for light transport calculations. *PhD dissertation, Stanford*, 1997.