Journal of Computational Physics 230 (2011) 7723–7735

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

A hybrid (Monte Carlo/deterministic) approach for multi-dimensional radiation transport

Guillaume Bal [a,1], Anthony B. Davis [b,c,2], Ian Langmore [a,*]

[a] Department of Applied Physics and Applied Mathematics, Columbia University, 200 S.W. Mudd Building, 500 W. 120th Street, New York, NY 10027, USA
[b] Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Mail Stop 169-237, Pasadena, CA 91109, USA
[c] Kavli Institute for Theoretical Physics, Kohn Hall, University of California, Santa Barbara, CA 93106-4030, USA

ARTICLE INFO

Article history:
Received 19 June 2010
Received in revised form 8 May 2011
Accepted 25 June 2011
Available online 5 July 2011

Keywords:
Linear transport
Monte Carlo
Hybrid methods
Importance sampling
Variance reduction
3D rendering
Remote sensing

ABSTRACT

A novel hybrid Monte Carlo transport scheme is demonstrated in a scene with solar illumination, scattering and absorbing 2D atmosphere, a textured reflecting mountain, and a small detector located in the sky (mounted on a satellite or a airplane). It uses a deterministic approximation of an adjoint transport solution to reduce variance, computed quickly by ignoring atmospheric interactions. This allows significant variance and computational cost reductions when the atmospheric scattering and absorption coefficient are small. When combined with an atmospheric photon-redirection scheme, significant variance reduction (equivalently acceleration) is achieved in the presence of atmospheric interactions.

1. Introduction

1.1. Motivation and background

Forward and inverse linear transport models find applications in many areas of science including neutron transport [1–3], medical imaging and optical tomography [4,5], radiative transfer in planetary atmospheres [6–8] and in oceans [9,10], as well as the propagation of seismic waves in the solid Earth [11]. In this paper, we focus on the solution of the forward transport problem by the Monte Carlo (MC) method with, as our main application, remote sensing (an inverse transport problem) of the atmosphere/surface system [12]. In our demonstration, light is emitted from the Sun and propagates in a complex environment involving absorption and scattering in the atmosphere and reflection at the Earth's surface before (a tiny fraction of) it reaches a narrowband detector, typically mounted on a airplane or a satellite.

The integro-differential transport Eq. (1) 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 (in principle) no

* Corresponding author. Tel.: +1 415 272 6321.
E-mail addresses: gb2030@columbia.edu (G. Bal), Anthony.B.Davis@jpl.nasa.gov (A.B. Davis), ianlangmore@gmail.com (I. Langmore).
[1] Tel.: +1 212 854 4731.
[2] Tel.: +1 818 354 0450, +1 805 893 4111.

0021-9991/$ - see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2011.06.029

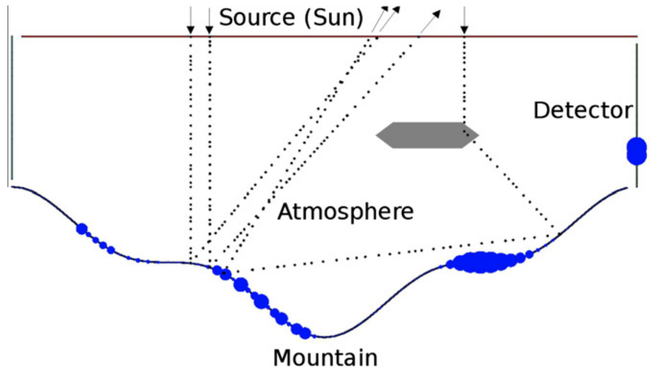

Fig. 1. Mountain ($1-\cos^3 x$ shape), cloud, sky, and detector. Dot size indicates relative adjoint flux strength. Large dots on right-hand-side are the detector (dot size is down-scaled for detector). Dot size on mountain indicates that portions of the mountain are shaded from the detector, and that the surface albedo is varying. See Section 3.1 for specifics, as used in the present study.

discretization error. The drawback is that they can be very slow to converge. MC methods converge at a rate $(variance/N)^{1/2}$ where N is the number of simulations, and the *variance* is that of each photon 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, even in a forward simulation, MC methods must be accelerated.

One approach to speedup MC simulations is to use quasi-Monte Carlo methods, which steepen the convergence rate from $\sim N^{-1/2}$ to a more negative exponent. However, most MC speedup efforts focus on reducing the variance of each photon. See [2,3] or the review of more recent work on neutron transport in [14,15,13] and on 3D atmospheric radiative transfer in [16,17]. See also [18] for a thorough introduction to the MC techniques, including variance reduction, used 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 [2,3].[3] Often, one uses some heuristic (such as proximity to the detector), or some function to measure the "importance" of each region of phase space. In *splitting methods* [2,3], the photon is split into two or more photons upon identifying that a photon is in a region of high importance. The weight of each photon is then decreased proportionately. 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 terminated with a certain probability, i.e., high chance of absorption if the weight is low; in the rarer alternative outcome of the Bernoulli trial, the weight is increased to keep the simulation numerically unbiased. So there is typically a slight cost in variance to improve efficiency (by terminating low-weighted trajectories). 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 and scattering properties at any point can be modified, provided photons are re-weighted correctly.

It has long been recognized that the adjoint transport solution is a natural importance function [19,2,3,20–22,14,23,15]. One can use approximations of the adjoint solution—typically a coarse deterministic solution—to reduce variance. The result is a *hybrid* method (deterministic and MC). The AVATAR method uses an adjoint approximation to determine weight windows [22]. The CADIS scheme in [14] uses an adjoint approximation in both source biasing and weight-window determination. An adaptive technique that successively refines the solution in "important" regions, using the adjoint to designate such regions, is described in [24,25]. In [19,2,3,20], a zero-variance technique is outlined that uses the true adjoint solution to launch photons that all reach the detector with the same weight ...which happens to be the correct answer. This method is of course impractical since determining the exact adjoint solution everywhere is harder than determining some specific integral of that solution, which is usually the goal of a MC simulation. The LIFT method [20,21] therefore 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 Fig. 1 for the type of geometry considered in this paper. The problem we consider has a fixed, partially-reflective, complex-shaped lower 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 (optically thin) atmosphere. Calculation of the approximate adjoint solution used to emulate 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 atmospheric scattering (infinite MFP limit). The computation of the adjoint solution thus becomes a radiosity problem with much reduced dimensionality compared to the full transport problem. This, of course, can only reduce variance in proportion to the number of

[3] Note the somewhat confusing terminology: on the one hand, a method is statistically biased if the expected outcome is not the intended one. On the other, the practice of re-directing photons in favorable directions and/or reducing the number of scattering events is also called biasing. In the latter case the photon has its weight adjusted so that the simulation is unbiased.