Validation of the Numerical Accuracy and Efficiency of the Hybrid Method

Hongsong Li, Fabio Pellacini, Kenneth E. Torrance, Dhruva Karle

TR-PCG-05-04

Tech report (2005)
Unspecified for anonymity

Validation of the Numerical Accuracy and Efficiency of the Hybrid Method

Supplemental material for Online ID: xxx

Abstract

To characterize the accuracy and efficiency of our hybrid Monte Carlo scheme, we devised a set of numerical experiments that compare the newly proposed method with full Monte Carlo simulations as well as the Jensen et al.'s approximation. The results of these tests show that our method produces results whose numerical *accuracy is comparable to Monte Carlo simulations at a much lower computational cost. These experiments also show the source of the inaccuracy that Jensen et al.approximation displays when rendering optically-thin materials.*

1. Semi-infinite slab

To test the convergence of our method, we first calculate the radially-resolved diffuse reflectance of an infinitely narrow beam of light normally incident upon the top surface of a flat, semi-infinite slab. A comparison of the results obtained with the Hybrid method, a full Monte Carlo simulation and the diffuse approximation is found in Figure 1.a where the radius r is normalized with the mean free path. The graph shows that the reflectance predicted by the hybrid method agrees well with the pure Monte Carlo method when r is small and slightly underestimates the diffuse reflectance when r gets large; this loss of energy is due to the approximation used to evaluate the uniform-diffuse component. Nonetheless this error remains small enough to produce very accurate results.

To compare the computational cost of the hybrid method versus a pure Monte Carlo simulation, we run a numerical experiment that compares the average number of scattering events required to reach convergence. To do so a number of path were shot into a flat, semi-infinite slab at normal incidence and scattered until either absorbed of transmitted out. Results for this experiment for varying σ_s/σ_a and g are shown in Figure 2.a. This test shows that the number of scattering events for the pure Monte Carlo method increase linearly with σ_s/σ_a and exponentially with the anisotropy factor g. In a highly scattering ($\sigma_s \gg \sigma_a$), anisotropic ($g \approx 1$) medium, a traced path experiences many scattering events before it gets absorbed or reflected. For the hybrid method, the number of scattering events increases to a plateau, which is the average number of scattering events for a path to reach the isotropic core region. For strongly anisotropic scattering ($g \approx 1$), a path needs to experience more scattering events to smooth out its dependence on the incident direction, leading to a higher upper bound.

This experiment shows that hybrid method significantly accelerates the simulation without any appreciable loss of accuracy.

2. Shadow on a participating medium

A more complex test is the evaluation of convergence for a shadow on a participating medium ($\sigma_s/\sigma_a = 100$, $g = 0$), showing the diffusion of light across a shadow boundary. For our experiment, we illuminate a semi-infinite slab with a point light source blocked by a long shade whose orientation is shown in Figure 3. The camera views the shadow edge with a lateral field of view that is five times the mean free path.

Figure 1.b shows the average of the pixel values y, computed in a direction parallel to the shadow boundary x, plotted against the direction orthogonal to the shadow boundary. Images with resolution of 256x256 were rendered at 100 samples per pixel using a pure Monte Carlo scheme, our hybrid method and Jensen et al.'s approximation and took respectively 96, 6 and 2 minutes.

This experiment shows that the three methods agree well, but Jensen et al.'s approximation is slightly lower than the other two curves, due to its energy loss, and shows a discontinuity that comes from the single-scattering term of the model. This latter inaccuracy generates a sharp edge artifact at the shadow boundary that does not appear with the other two methods.

3. Wedge-like geometry

To show the dependance of the geometry on the subsurface behaviour, we now consider an infinite wedge of participating medium ($\sigma_s/\sigma_a = 100$, $g = 0$) whose geometry is shown in Figure 4.a. With the increase of the wedge thickness, the average number of the scattering events increases linearly, resulting in a variation of the scale factor for the uniform-diffuse component k_{ud} from 0 to 1.

Figure 1.c shows the average of the pixel values, computed in a direction parallel to the wedge y, plotted against the direction orthogonal to the wedge x. Images with resolution of 256x256 were rendered at 100 samples per pixel using a pure Monte Carlo scheme, our hybrid method and Jensen et al.'s approximation and took respectively 12, 11 and 2 minutes. Results are also included for the single-scattering term and the diffusion term of the Jensen et al.model.

In this experiment, the hybrid method agrees very well with the pure Monte Carlo simulation, only slightly underestimating the transmitted light at the thick end of the wedge, an error that comes from the approximation used to account for the uniform-diffuse component. On the other hand, Jensen et al.model produce results that are quite different with small errors originating from the single scattering term and larger inconsistency coming from the diffuse term.

In particular, the diffusion term of the Jensen et al.model predicts an unreasonably large result at the thin edge of the wedge and underestimates the light scattering at the thick end. The sources of this imprecisions are several. First, the distance of the real source of the dipole to the point of illumination z_r should always be no less than one mean free path, as suggested in the original approach, but, by enforcing this limit, the positions of real sources become implausible at the thin edge of the wedge, as shown in Figure 4.c. Thus, the lower bound of the mean free path leads to an overestimation of the diffusion term at the thin end. According to our analysis, we believe the main problem is not how to position the real sources for the dipole method, but how to determine the contribution of the uniform-diffuse component of the subsurface scattering k_{ud} for curved optically-thin participating medium, where, in region of high curvature, should be less than 1. The second source of error is the fact that the original dipole diffusion approximation was derived for the case where the light and camera are on the same side of the material, an incorrect assumption that leads to an underestimation of the reflection at the thick end of the wedge. The accuracy of our hybrid approach comes from the use of the isotropic core region to determine the contribution of the uniform-diffuse term k_{ud}, as illustrated in Figure 4.b.

This experiment not only shows that a Monte-Carlo-simulated directional-diffuse component is essential for a universal solution for the participating media with arbitrary geometry and optical properties, but also helps us explain the blooming artifacts introduced by the Jensen et al.approximation.

A comparison of the average numbers of scattering events for the pure Monte Carlo and the hybrid method, plotted against the normalized thickness of wedge in Figure 2.b gives us an evaluation of the efficiency of our method. The averages for the Monte Carlo method increase linearly with the thickness of the wedge, until reaches an upper bound, representing the furthest distance a photon can travel before getting absorbed completely. The hybrid method shows a similar trend but with lower upper bounds due to the isotropic core region. The speed-up becomes significant when the thickness increases. For this scene, most of the field of view is at the thin end of the wedge. Therefore, the computational times for the pure Monte Carlo and hybrid methods are much closer and at the same the efficiency advantage of the Jensen et al.approach is less obvious.

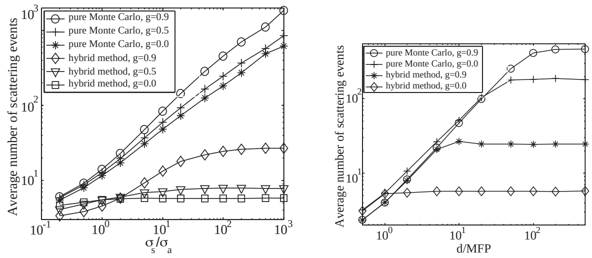

a) Semi-infinite slab b) Wedge

Figure 2: *Accuracy comparison for a) semi-infinite slab, b) shadow edge and c) wedge.*

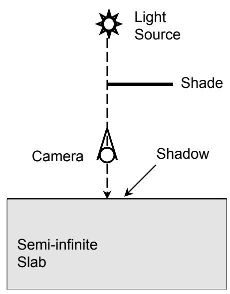

Figure 3: *Geometry for the shadow edge test case.*

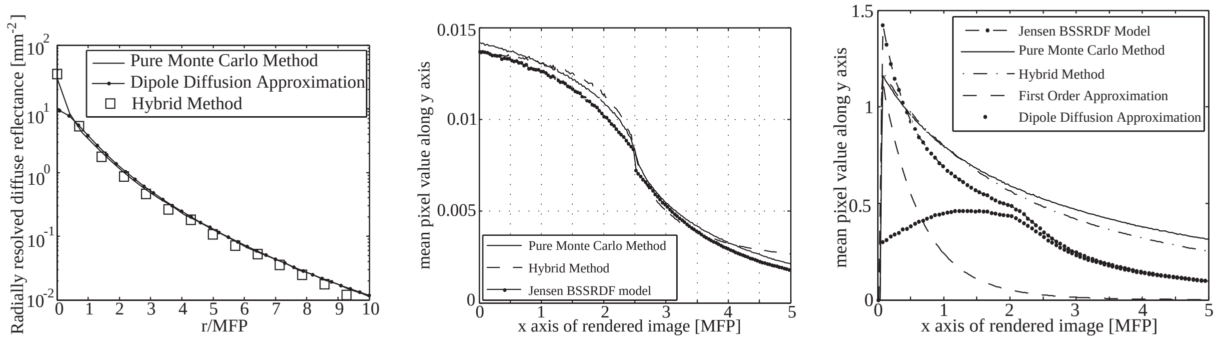

a) Semi-infinite slab b) Shadow edge c) Wedge

Figure 1: *Accuracy comparison for a) semi-infinite slab, b) shadow edge and c) wedge.*

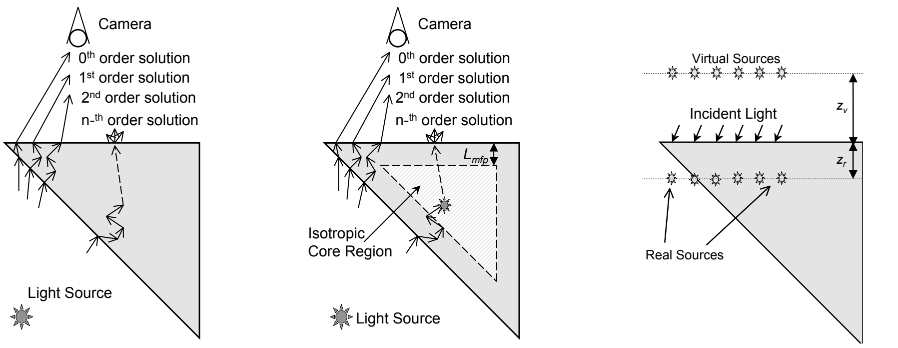

a) Monte Carlo b) Hybrid method c) Jensen et al. dipole approx.

Figure 4: *Comparison of the solution for the wedge example using a) Monte Carlo, b) our hybrid method and c) Jensen et al.dipole diffusion approximation.*