

A Hybrid Monte Carlo Method for Accurate and Efficient Subsurface Scattering

Hongsong Li[†] Fabio Pellacini[†] Kenneth Torrance[†]

Program of Computer Graphics, Rhodes Hall, Cornell University, Ithaca, NY 14853, U.S.A.

Abstract

Subsurface scattering is a fundamental aspect of surface appearance responsible for the characteristic look of many materials. Monte Carlo path tracing techniques can be employed with high accuracy to simulate the scattering of light inside a translucent object, albeit at the cost of long computation times. In a seminal work, Jensen et al. [JMLH01] presented a more efficient technique to simulate subsurface scattering that, while producing accurate results for translucent, optically-thick, materials, exhibits artifacts for semi-transparent, optically-thin, ones, especially in regions of high-curvature.

This paper presents a hybrid Monte Carlo technique capable of simulating a wide range of materials exhibiting subsurface scattering, from translucent to semi-transparent ones, with an accuracy comparable to Monte Carlo techniques but at a much lower computational cost. Our approach utilizes a Monte Carlo path tracing approach for the first several scattering events, in order to estimate the directional-diffuse component of subsurface scattering, and switches to a dipole diffusion approximation only when the path penetrates deeply enough into the surface. By combining the accuracy of Monte Carlo integration with the efficiency of the dipole diffusion approximation, our hybrid method produces results as accurate as full Monte Carlo simulations at a speed comparable to the Jensen et al. approximation, thus extending its usefulness to a much wider range of materials.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism I.3.3 [Computer Graphics]: Picture/Image generation

1. Introduction

The simulation of subsurface scattering is of fundamental importance for capturing the appearance of many materials such as marble, cloth, milk and skin. The volumetric effect of subsurface scattering gives such materials a unique look, varying from translucent to semi-transparent, as shown in Figure 1. The degree of transparency can be described by the *optical thickness* $\tau = s/L_{mfp}$, that, for a homogeneous material, is the ratio of the object size s versus the mean free path of scattering L_{mfp} . Optically-thick materials, $\tau \gg 1$, diffusely scatter the incident light showing a translucent appearance; an example of this is Figure 1.a. Optically-thin materials, $\tau \ll 1$, allow the light to travel through the volume without being diffusely scattered and have a semi-


transparent appearance shown for example in Figures 1.b-d. Note that an object often shows regions where the material is optically-thick, for example the statue's torso, as well as regions, especially in areas of high curvature, where the material is optically-thin, for example the statue's hand.

While Monte Carlo path tracing techniques can be used to accurately simulate subsurface effects, the computational cost of such algorithms make them impractical for most applications. In a seminal work, Jensen et al. [JMLH01] showed that an approximate subsurface scattering model can produce very convincing results for translucent, optically-thick, materials at a much lower cost than Monte Carlo simulations. Unfortunately, the approximations introduced in the model make the algorithm inaccurate for a large class of semi-transparent, optically-thin, materials, especially in regions of high curvature, as shown in Figure 2. Our goal is to extend the accuracy of the model presented in [JMLH01]

[†] (hongsong,fabio,ket)@graphics.cornell.edu



Figure 1: Optically-thick to optically-thin materials showing changes in appearance from translucent to semi-transparent. The ratio of the mean free paths is 1:3:10:30 from left to right.



a. Monte Carlo (246 min)

b. Hybrid method (33 min)

c. Jensen et al. (10 min)

Figure 2: Images rendered using a) pure Monte Carlo simulates, b) our hybrid method and c) Jensen et al. approximation.

to a wider range of translucent materials while keeping its low computational cost compared to full Monte Carlo simulations. In developing our technique, we put emphasis on the high accuracy of the simulation which is required for many applications where artifacts are undesirable. In order to achieve our goal, we utilize a Monte Carlo path tracing algorithm to correctly capture the effect of the *first few* scattering events, while switching to the dipole diffusion approximation *only if* a given path penetrates deeply enough into the surface, thus accelerating convergence. This combination, which is the main contribution of our work, provides results as accurate as full Monte Carlo simulations not only for translucent, optically-thick, materials but also for semi-

transparent, optically-thin, ones while maintaining the cost low.

2. Related work

Light transport in a volume can be computed accurately by solving the full radiative transport equation [Cha60]. Given the high computational cost, only a few approaches were proposed in the graphics literature. In particular, Dorsey et al. [DEL*99] simulated the appearance of weathered stone using photon mapping, while Pharr and Hanrahan [PH00] used nonlinear integral scattering functions. Albeit expensive, these techniques provide very accurate simulations valid for most materials.

In a work, Jensen et al. [JMLH01] proposed an efficient method to simulate subsurface scattering based on a BSSRDF approximation that combines the first order approximation, used also in [HK93], with the diffusion method, also found in [Sta95], using a dipole formulation. The combination of these two methods provides a very efficient rendering of subsurface scattering, whose speed can be further improved by using a hierarchical sampling scheme [JB02]. Unfortunately, the approximations introduced in the method limit its applicability to only translucent, optically-thick, materials. Our work tries to extend the accuracy of this method to semi-transparent, optically-thin, materials, while maintaining the computational efficiency.

Shell texture functions [CTW*04], a technique developed concurrently with our own, provide an efficient rendering approach for non-homogeneous translucent materials with surface mesostructure by combining a photon-map-based simulation of scattering in a thin non-homogeneous layer with the diffuse approximation in a homogeneous inner core. While this technique improves over previous methods, its accuracy is compromised by several approximations that limit its applicability to only optically-thick materials. These approximations include the uniformly-diffuse backlighting from the inner core and the limited angular and spatial resolution. For semi-transparent objects, the backlight can be quite directional and the angular variation of subsurface scattering is so dramatic that a large number of light directions need to be sampled for a correct simulation. Finally, the pre-computation times required make it prohibitive in the case of high-quality renderings of deformable objects.

3. Algorithm

3.1. Introduction

Subsurface scattering can be described with the bidirectional surface scattering distribution function (BSSRDF), S , that relates the outgoing radiance $L_o(x_o, \vec{\omega}_o)$ at a point x_o in direction $\vec{\omega}_o$ to the incident flux Φ_i at a point x_i from direction $\vec{\omega}_i$

$$dL_o(x_o, \vec{\omega}_o) = S(x_i, \vec{\omega}_i; x_o, \vec{\omega}_i) d\Phi_i(x_i, \vec{\omega}_i)$$

The BSSRDF captures the behavior of light entering the material and scattering multiple times before exiting. If the BSSRDF is known, the outgoing radiance can be computed by integrating the incident radiance over direction and area. It is important to note that the values of the BSSRDF depend not only on the material properties of the medium, but also on the geometry of the object being rendered, especially in the case of optically-thin materials.

3.2. A complete BSSRDF model

Inspired by the previous BSSRDF and also by the BRDF literature, a comprehensive BSSRDF model can be con-

structed from the sum of a directional-diffuse component and a uniform-diffuse one.

The directional-diffuse component comes from the *first few* scattering events where the angular distribution of the scattered light still has a dependence on the incident direction. Given the small number of scattering events, their total number as well as the spatial and angular distribution of the scattered light depend strongly on the geometry of the medium; in the case of optically-thin, highly-curved objects, the angular and spatial variations of this component can become dramatic.

The uniform-diffuse component originates from the *subsequent many* scattering events where the angular dependence of the light scattered within the material vanishes. This component does not always exist, especially for optically-thin materials, since the light can escape the medium before becoming fully diffused. Therefore, a scale factor for its value is implied. This component tends to have a mild spatial variation that is well captured by methods such as the dipole diffusion approximation.

The complete BSSRDF model can be written as:

$$S(x_i, \vec{\omega}_i; x_o, \vec{\omega}_i) = S_{dd}(x_i, \vec{\omega}_i; x_o, \vec{\omega}_i) + k_{ud}(x_i, \vec{\omega}_i; x_o, \vec{\omega}_i) \cdot S_{ud}(x_i, \vec{\omega}_i; x_o, \vec{\omega}_i) \quad (1)$$

where S is the complete BSSRDF, S_{dd} is the directional-diffuse component, S_{ud} is the uniform-diffuse component and k_{ud} is its scale factor. Note that k_{ud} has a directional dependence that comes from the directional transmission across the material interfaces, while the light transport inside the volumes is directionally independent. The directional-diffuse component in this equation can be expressed more explicitly by:

$$S_{dd}(x_i, \vec{\omega}_i; x_o, \vec{\omega}_i) = \sum_{i=1}^N S^i(x_i, \vec{\omega}_i; x_o, \vec{\omega}_i) \quad (2)$$

where S^i is the contribution of the i -th scattering event and N is the number of scattering events that is large enough to guarantee isotropic scattering.

3.3. Comparison with previous methods

For the case of optically-thick materials, Jensen et al. presented a practical and efficient approximation to evaluate the BSSRDF defined by the sum of two terms, a *single-scatter* term and a *diffuse-dipole approximation* term (for a detailed description the reader should consult [JMLH01]):

$$S(x_i, \vec{\omega}_i; x_o, \vec{\omega}_i) = S^1(x_i, \vec{\omega}_i; x_o, \vec{\omega}_i) + S_d(x_i, \vec{\omega}_i; x_o, \vec{\omega}_i)$$

In this equation, S^1 captures the directional effect of the single scattering event, while S_d describes the isotropic diffuse behavior of the remaining multiple scatters estimated using the dipole approximation.

With respect to the complete BSSRDF, there are two main

approximations taken by this method: the directional-diffuse component is estimated by using only the first scattering event (first-order approximation) and the uniform-diffuse contribution is always present. This is equivalent to setting $k_{ud} = 1$ and $N = 1$ in equations (1) and (2). These approximations assume that only one scattering event is necessary for the light behavior to become isotropic and that the light would not escape the medium before then. While these approximations work well for optically-thick materials, they do not apply to optically-thin materials where, even after several scattering events, the light behavior can remain directional and often the light exits the medium before becoming diffuse. This behavior becomes particularly noticeable in regions of high curvature, showing strong artifacts from the energy imbalance implicit in the approximation, as in Figure 2.

While STF [CTW*04] improves on the original BSSRDF, it cannot deal properly with the multiple scattering as well as the dependance on the object's geometry since all the subsurface scattering events that follow the first one are treated as isotropic and since STF intentionally neglects the boundary effects that can have a great influence on the appearance of a semi-transparent object.

3.4. A hybrid Monte Carlo method

Inspired by [Wan98] we propose a hybrid Monte Carlo method to evaluate the complete BSSRDF, where the directional-diffuse and the uniform-diffuse contributions are evaluated respectively using a path tracing approach and the dipole diffusion approximation. The key to combine these methods is to identify an *isotropic core region*, illustrated in Figure 3.a, inside a participating medium defined as the region beneath the surface where any point has a minimum distance of a mean free path L_{mfp} from any point on the surface. While a similar idea is also being used in the concurrently-developed STF [CTW*04], the methods remain very different in the way they utilize this principle; while in our method the isotropic core region is used to divide the directional-diffuse and uniform-diffuse components, STF divides the two components by the first scattering events.

The fundamental observation is that a photon will contribute to the uniform-diffuse component of the BSSRDF only when it reaches the isotropic core region; otherwise its contribution should be counted in the directional-diffuse component. This observation is justified by the fact that statistically a scattered photon loses its dependence on the incident direction after being scattered over a distance that equals the mean free path and that it will also lose its dependence on the outgoing direction since it will take the same average number of scattering events to reach the surface; thus, it is safe to assume that such a photon contributes to the uniform-diffuse component.

3.4.1. Evaluating the directional-diffuse component

Our algorithm traces paths from the eye and propagates them into the medium using a standard path tracing approach similar to [Wan98]. This process continues until either the path exits the medium or the positions of the *last two scattering events*, indicated by p_m and p_{m+1} , lie inside the core region; for the latter case, the uniform-diffuse component is then computed. This process, illustrated in Figure 3.b, is responsible for the evaluation of S_{dd} and k_{ud} in equation (1).

3.4.2. Evaluating the uniform-diffuse component

The uniform-diffuse component is evaluated using the dipole diffusion approximation for the path at position p_m , where the position of the source $p_{iso} = p_m + L_{mfp} \cdot \vec{s}$ is determined based on the expected scattering direction $\vec{s} = p_{m+1} - p_m$ and the mean free path L_{mfp} while its weight is given by

$$w_{iso} = w_{in} \left(\frac{\sigma_s}{\sigma_s + \sigma_a} \right)^m \cdot \frac{\sigma_s(1-g)}{\sigma_s(1-g) + \sigma_a}$$

where w_{in} is the weight of the photon at incidence, m is the length of the path and σ_s , σ_a , g are the scattering coefficient, the absorption coefficient and the anisotropy factor respectively.

In order to estimate the uniform-diffuse contribution to surface reflectance, an approximate dipole method is used as in Jensen et al. [JMLH01]. In order to support arbitrary geometry, the radially resolved diffuse reflectance R_{da} is estimated from the original formulation R by replacing the distances z_r and z_v of the original method with z_{ra} , the distance between the incidence point and the isotropic source, and $z_{va} = z_{ra} + 4AD$, as shown in Figure 3.c, where (A , D are defined in the original method). Note that the isotropic source may not always be right below the incidence point, which is a reasonable approximation for arbitrary 3D geometry, while we still place the virtual source of the dipole above the incidence point. Under these conditions, the uniform-diffuse component becomes

$$S_{ud}(x_i, \vec{\omega}_i; x_o, \vec{\omega}_o) = \frac{1}{\pi} F_t(\eta, \vec{\omega}_i) R_{da}(\|x_i - x_o\|) F_t(\eta, \vec{\omega}_o)$$

where η is the index of refraction, F_t is the Fresnel transmittance, and R_{da} is the corrected radially-resolved diffuse reflectance. The sampling is accelerated by the importance sampling scheme introduced in [MKB*03].

3.4.3. Estimating surface distance

To test whether a scattering event is inside the isotropic core region, we need to evaluate the minimum distance between the current position and the object surface. To do so, we store the vertices of the model in a KD tree data structure and, at each scattering event, use this data structure to find the closest surface point used to calculate the distance [JC98]. For this method to provide a good approximation for our algorithm, the distance between sample points on the surface has




Figure 3: Hybrid method. a) Paths can either contribute to the directional-diffuse component or, if they enter a core region, to the uniform-diffuse one. b) The first few scattering events are simulated using Monte Carlo techniques, while the dipole diffusion approximation is used for paths that enter and remain in the core region. c) Geometry of the dipole diffusion approximation.

to be smaller than the mean free path. In our implementation, the time for constructing the KD tree is usually negligible compared to the total computation time, while its lookup time is less than the cost of simulating a scattering event; more efficient proximity query algorithms could further improve the performance of our approach.

4. Results

Figure 4 shows a comparison between a pure Monte Carlo technique, our hybrid method, and the Jensen et al. approximation for the case of a complex model of 270,000 polygons lit by a point light source shadowed by a grid-shaped object. The material in the dragon is isotropic and highly scattering with $g = 0$ and $\sigma_s/\sigma_a = 100$. The two series of pictures were generated by scaling σ_s and σ_a to obtain optically-thin and optically-thick materials, whose ratio of mean free paths is 10. A Phong lobe is added to account for the first-surface scattering, producing a glossy appearance. The scene was rendered with 100 samples per pixel at a resolution of 1024x768. Rendering times for a 1Ghz Pentium 4 are reported in Figure 4, where the time needed for constructing the KD tree for the hybrid method is about 10 seconds.

For the optically-thick material, the images rendered with the three methods agree well. While the hybrid method tends to slightly blur the material surface and Jensen's approximation tends to slightly underestimate the subsurface scattering, all three rendered images are of high quality and free of artifacts. The light transport inside the volume can be considered as a local behavior since the mean free path is much smaller than the size of the highly-curved regions of the object's surface; further, the subsurface behavior produces some subtle effects, such as blurred shadow boundaries. The hybrid method is about 10 times faster than the pure Monte Carlo simulation, which is particularly inefficient for this material.

For the optically-thin material, the pure Monte Carlo

method generates a smooth, semi-transparent appearance, which is evidently different from the translucent appearance of the optically-thick case. For this material, the mean free path is comparable to the size of the highly-curved regions of the object surface. Jensen's BSSRDF approximation cannot account for such cases, generating images that are visibly different from the other two methods, with visible artifacts in highly curved regions such as the eyelid and the corner of the mouth. In these regions the subsurface scattering was overestimated, mainly due to the unscaled dipole diffusion approximation. At the same time, Jensen's BSSRDF model gives plausible results for the mildly-curved regions of the object surface. The hybrid method generates rendered images comparable to the pure Monte Carlo method, but with a speedup of about 7 times. Similar relative behavior can be observed in Figure 2. Additional, quantitative comparisons of the three methods, including accuracy and efficiency, are reported in the Appendix.

5. Conclusion and Future Work

In this paper, we present a hybrid algorithm that combines the accuracy of the pure Monte Carlo method with the efficiency of the dipole diffusion approximation. Our approach is able to handle a broad range of participating media, from optically-thick to optically-thin, with much better efficiency than the pure Monte Carlo method and without compromising its accuracy. Covering such a wide range of materials is important since the appearance changes considerably from translucent to semi-transparent as can be clearly seen in Figure 1.

In the future, we plan to extend our method to incorporate non-homogeneous materials, such as layered materials and subsurface structures, as in [CTW*04]. Also we would like to investigate the possibility of further improving the efficiency of our method by removing the distance queries; it might be possible, for example, to scatter a fixed number of times in the medium before switching to the diffusion

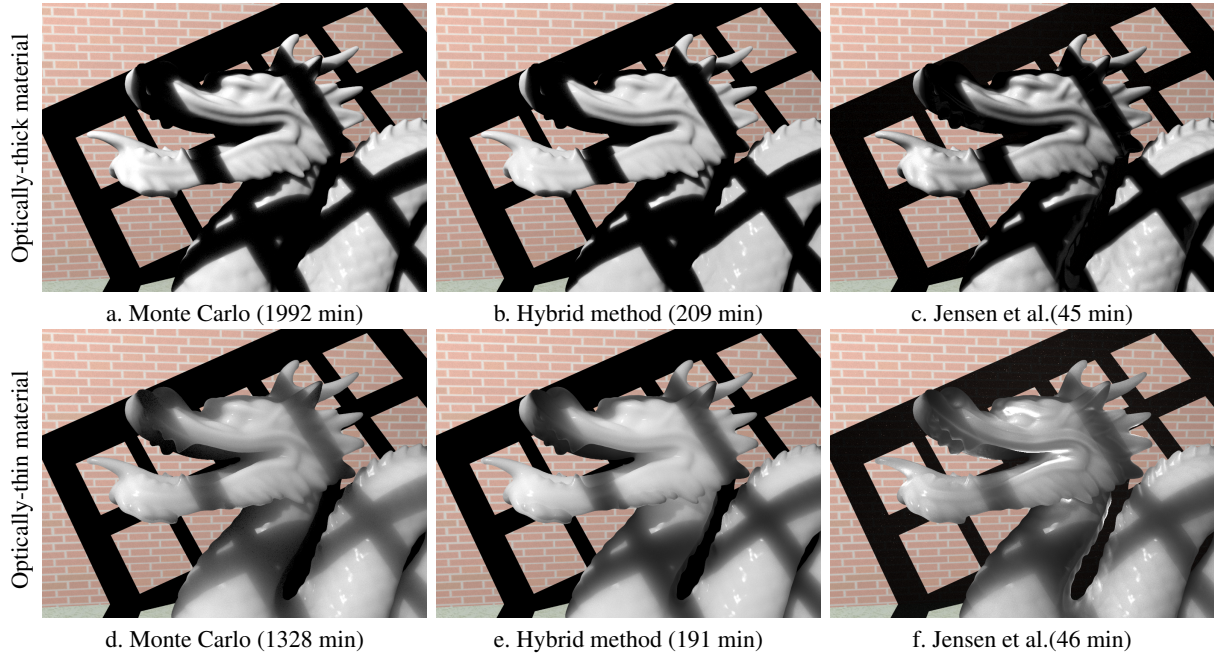


Figure 4: Images rendered using a,d) pure Monte Carlo, c,e) our hybrid method and c,f) Jensen et al. approximation.

approximation. While this will introduce artifacts in many cases, it is possible that these are minor for a wide-enough class of materials.

6. Acknowledgments

We would like to thank Dhruva Karle for help with the implementation, the staff of the Program of Computer Graphics at Cornell and the anonymous reviewers of this paper for many helpful comments, and the financial support of National Science Foundation Grant ACI-0113851.

References

- [Cha60] CHANDRASEKHAR S.: *Radiative Transfer*. Oxford University Press, 1960.
- [CTW*04] CHEN Y., TONG X., WANG J., LIN S., GUO B., SHUM H.-Y.: Shell texture functions. *ACM Transactions on Graphics* 23, 3 (Aug. 2004), 343–353.
- [DEL*99] DORSEY J., EDELMAN A., LEGAKIS J., JENSEN H. W., PEDERSEN H. K.: Modeling and rendering of weathered stone. In *Proceedings of SIGGRAPH 99* (Aug. 1999), Computer Graphics Proceedings, Annual Conference Series, pp. 225–234.
- [HK93] HANRAHAN P., KRUEGER W.: Reflection from layered surfaces due to subsurface scattering. In *Proceedings of SIGGRAPH 93* (Aug. 1993), Computer Graphics Proceedings, Annual Conference Series, pp. 165–174.
- [JB02] JENSEN H. W., BUHLER J.: A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics* 21, 3 (July 2002), 576–581.
- [JC98] JOHNSON D., COHEN E.: A framework for efficient minimum distance computation. In *IEEE Conf. On Robotics and Animation* (1998), pp. 3678–3683.
- [JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Proceedings of ACM SIGGRAPH 2001* (Aug. 2001), Computer Graphics Proceedings, Annual Conference Series, pp. 511–518.
- [MKB*03] MERTENS T., KAUTZ J., BEKAERT P., SEIDEL H.-P., REETH F. V.: Efficient rendering of local subsurface scattering. In *11th Pacific Conference on Computer Graphics and Applications* (2003), p. 51.
- [PH00] PHARR M., HANRAHAN P. M.: Monte carlo evaluation of non-linear scattering equations for subsurface reflection. In *Proceedings of ACM SIGGRAPH 2000* (July 2000), Computer Graphics Proceedings, Annual Conference Series, pp. 75–84.
- [Sta95] STAM J.: Multiple scattering as a diffusion process. In *Eurographics Rendering Workshop 1995* (June 1995), pp. 41–50.
- [Wan98] WANG L. H.: Rapid modeling of diffuse reflectance of light in turbid slabs. *J. Opt. Soc. Am. A15*, 4 (1998), 936–944.




Figure 5: Accuracy comparison for a) semi-infinite slab, b) shadow edge and c) wedge.

Appendix. Validation of accuracy and efficiency

A set of numerical experiments was devised to assess the accuracy and efficiency of the newly proposed hybrid Monte Carlo scheme, a full Monte Carlo simulation, and the Jensen et al. approximation. The new hybrid method produces results whose *numerical* accuracy is comparable to Monte Carlo simulations at a much lower cost. The experiments reveal the source of the inaccuracy that the Jensen et al. approximation displays when rendering optically-thin materials.

6.1. 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 that is 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 5.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 re-

quired to reach convergence. A number of paths were shot into a flat, semi-infinite slab at normal incidence and scattered until either absorbed or transmitted out. Results for varying σ_s/σ_a and g are shown in Figure 6.a. The number of scattering events for the pure Monte Carlo method increases linearly (in log-log scales) 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 the hybrid method significantly accelerates the simulation without any appreciable loss of accuracy.

6.2. 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 edge. 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 7. The camera views the shadow edge with a lateral field of view that is five times the mean free path. Figure 5.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

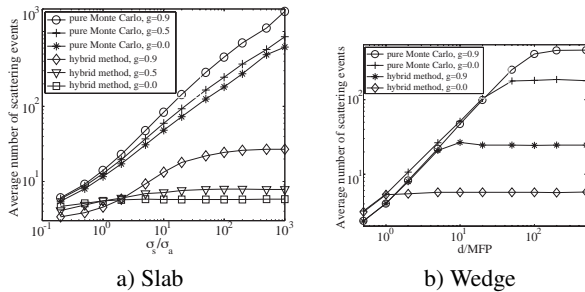


Figure 6: Efficiency comparison for a) slab and b) wedge.




Figure 7: Geometry for the shadow-edge test case.




Figure 8: Comparison of the solution for the wedge example using a) Monte Carlo, b) our hybrid method and c) Jensen et al. approximation.

using a pure Monte Carlo scheme, our hybrid method and the Jensen et al. 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.

6.3. Wedge-like geometry

To show the dependance of the geometry on the subsurface behaviour, we now consider an infinite wedge of a participating medium ($\sigma_s/\sigma_a = 100$, $g = 0$) with a light source/camera arranged as shown in Figure 8.a. With the increase of the wedge thickness, the average number of scattering events increases linearly, resulting in a variation of the scale factor for the uniform-diffuse component k_{ud} from 0 to 1. Figure 5.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 256×256 were rendered at 100 samples per pixel using a pure Monte Carlo scheme, our hybrid method and the Jensen et al. 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, the Jensen et al. model produces results that are quite different with small errors originating from the single scattering term and a larger inconsistency coming from the diffusion 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 these 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 Fig-

ure 8.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, which, in regions 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 8.b.

This experiment not only shows that a Monte-Carlo-simulated directional-diffuse component is essential for a universal solution for 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 number of scattering events for the pure Monte Carlo and the hybrid method, plotted against the normalized thickness of the wedge in Figure 6.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 reaching 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 as 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 time the efficiency advantage of the Jensen et al. approach is less obvious.