

Interactive Light Scattering with Principal-Ordinate Propagation

Oskar Elek^{1,2*} Tobias Ritschel^{1,2} Carsten Dachsbacher³ Hans-Peter Seidel¹
¹ MPI Informatik ² MMCI / Saarland University ³ Karlsruhe Institute of Technology



Figure 1: Dense smoke exhibiting strong multiple anisotropic scattering produced by a steam locomotive under complex environment illumination. Our approach renders it interactively without any precomputations at 10 Hz (NVIDIA GeForce GTX 485 Mobile).

ABSTRACT

Efficient light transport simulation in participating media is challenging in general, but especially if the medium is heterogeneous and exhibits significant multiple anisotropic scattering. We present a novel finite-element method that achieves interactive rendering speeds on modern GPUs without imposing any significant restrictions on the rendered participated medium. We achieve this by dynamically decomposing all illumination into directional and point light sources, and propagating the light from these virtual sources in independent discrete propagation volumes. These are individually aligned with approximate principal directions of light propagation from the respective light sources. Such decomposition allows us to use a very simple and computationally efficient unimodal basis for representing the propagated radiance, instead of using a general basis such as Spherical Harmonics. The presented approach is biased but physically plausible, and largely reduces rendering artifacts inherent to standard finite-element methods while allowing for virtually arbitrary scattering anisotropy and other properties of the simulated medium, without requiring any precomputation.

Index Terms: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Radiosity; I.6.8 [Simulation and Modeling]: Types of Simulation—Parallel

1 INTRODUCTION

Scattering, or translucency, greatly contributes to the appearance of many natural substances and objects in our surrounding. Albeit the problem can be easily formulated as the radiance transfer equation [3, 21], computing a solution can be very costly. Consequently, many existing approaches simplify the problem, e. g., by assuming isotropic scattering or homogeneity of the material, to achieve interactive performance.

In this work we propose a novel interactive algorithm for plausible rendering of heterogeneous participating media with arbitrary

anisotropy. The core of our approach is to propagate light in propagation volumes oriented along the *principal ordinates* of the source illumination. For this we typically use multiple rectilinear grids to propagate environmental (distant) lighting, and spherical grids to account for point light sources. In both cases, one dimension of the grids is aligned with the prominent directional part of the source radiance for which the grid has been created. In contrast to previous methods (e. g., [1, 13]), discretizing the illumination into directional and point light sources enables us to approximately describe the anisotropy (directionality) of light transport by a single scalar value per grid cell. Specifically, this anisotropy value corresponds to a unimodal function implicitly aligned with the respective principal ordinate. In addition to exploiting data locality and the parallelism of GPUs, the benefit of these decisions is a significant reduction of the *false scattering* and *ray effect* artifacts arising in many finite-element methods as a consequence of representing the propagated radiance by, e. g., spherical harmonics or piecewise-constant functions.

Our main contributions can be summarized as follows:

- We introduce a novel approach to finite-element light propagation using implicitly aligned unimodal distributions for regular and spherical grids. This helps reducing propagation artifacts and helps to preserve the directionality of light during the propagation.
- A simplification of lighting by decomposing both environmental and surface illumination (via virtual point lights) into separate principal ordinate grids.
- An observer-centric importance-based selection of principal ordinates and prefiltering for environment lighting, helping to hide its discretized character used in the propagation.

2 PREVIOUS WORK

Offline methods A range of different approaches has been presented to compute solutions to the radiance transport equation for participating environments [3, 21]. However, none of the classic techniques provides a satisfying combination of generality, robustness, and, most importantly in our context, speed. Unbiased Monte-Carlo methods, such as bidirectional path tracing [18] and Metropolis light transport [26] usually require a large number of paths to be traced; in particular in dense media with high scattering anisotropy and albedo (like clouds or milk) the computation time increases tremendously. Caching is often used to speed up the computation, e. g., radiance caching [10], photon mapping [11, 12] or virtual point lights [6].

*e-mail: {oelek, ritschel, hps}@mpi-inf.mpg.de, dachsbacher@kit.edu

However, these methods typically do not handle highly anisotropic scattering very well, even with recent improvements [24, 25], and their performance is often far from interactive.

Finite-Element methods Finite-element methods, including volume radiosity [29], the discrete ordinates method (DOM) [3], light diffusion [32], and lattice-Boltzmann transport (LB) [8] handle highly multiple scattering well. However, in practice they allow only isotropic or moderately anisotropic scattering, and usually suffer from false scattering (smoothing of sharp light beams) and ray effects (selective exaggeration of scattered light due to discretized directions). Light propagation maps [7] significantly reduce the artifacts, but are still limited to rather moderate anisotropy. It can therefore be seen that strong scattering anisotropy is one of the main limiting factors for existing methods. This is unfortunate, as most real-world media exhibit relatively high anisotropy (Henyey-Greenstein [9] coefficient $g \approx 0.9$ or more [23]). Although isotropic approximations are acceptable in some cases, this is generally not a valid assumption and one of the primary motivations for our work.

Interactive rendering Numerous works focus on individual optical phenomena to achieve interactive or real-time performance. These phenomena include light shafts [5, 28], volume caustics [17, 19], shadows [20, 30], and clouds [2]. Various approaches can also be found in classic visualization literature, e. g., half-angle slicing [15] which empirically computes forward scattering for volume visualization. Sometimes precomputation is used to speed up the rendering of heterogeneous translucent objects [31, 33] or smoke using compensated ray marching [34]. In contrast, we target general multiple scattering in participating media without any precomputation or focus on a particular phenomenon.

We build on concepts of DOMs and light propagation volumes [1, 13]. These approaches are attractive for interactive applications as their grid-based local propagation schemes allow for easy parallel implementation on contemporary GPUs. Virtually all existing variants of DOM use a single scene-aligned grid, where every grid cell stores a representation of the directional radiance function using spherical harmonics (SH) or piecewise-constant functions. This representation is then used to iteratively calculate energy transfer between nearby cells, typically within a local 18- or 26-neighbourhood. However, this representation is only suited for moderately anisotropic scattering at best; especially for anisotropic media under complex (high-frequency) illumination such approach causes prominent ray effect and false scattering artifacts (see [7]). We take a different approach and propose to identify the most important light propagation directions (principal ordinates) in the scene and then use *multiple propagation grids* aligned with these directions, instead of a single volume. This enables using a unimodal representation of the angular energy distribution around the principal direction in each grid cell.

3 PRINCIPAL ORDINATE PROPAGATION

The core idea of our method is to reduce the main drawbacks of previous grid-based iterative methods, false scattering and ray effects, by using propagation volumes where the propagation domain is explicitly aligned with approximate principal directions of light transport. Furthermore, we use only a single scalar value per grid cell to describe the local anisotropy of the directional light distribution. In our scheme, we use the well-known Henyey-Greenstein (HG) [9] distribution; the aforementioned value, called the *anisotropy coefficient*, is used to parametrize this distribution. Using principal directions implies that for more complex lighting scenarios we have to use multiple grids that sufficiently well approximate their directionality; for local light sources we propose to use spherical grids centred around them.




Figure 2: For distant (parallel) light we use rectilinear grids aligned with its principal direction, and spherical grids for point light sources. Every grid cell stores only radiance magnitude and anisotropy. The propagation scheme is almost identical for both cases.

\mathbf{d}	(principal) direction
g	scattering anisotropy coefficient
σ_s, σ_a	scattering / absorption coefficient
\mathbf{x}_i	location of grid cell i
L_i, a_i	(per-cell) radiance magnitude and anisotropy
f_{hg}, F_{hg}	HG function and its cumulative distribution
μ	scattering angle cosine
$\mathbf{L}, \mathbf{L}_{acc}$	propagation and accumulation grid
M, m	number of iterations / iteration index
$L_{in}(\mathbf{d})$	incident radiance from direction \mathbf{d}
$\Delta L_{src \rightarrow dst}$	src to dst radiance contribution
$T_i, T_{src \rightarrow dst}$	transmittance to cell i and between cells
Ω_i, Ω_n	solid angle subtended by cell i or ordinate n
N, n	number of principal ordinates / ordinate index

Table 1: Table of symbols (in the order of appearance).

These choices assume that the principal directions can be derived from the initial radiance distribution and do not change strongly when light travels through the medium. However, such variation might occur if the density of the simulated medium changes abruptly. Still, as we discuss in Sec. 3.1.5, violating this assumption does not cause our algorithm to fail, but only leads to decreasing its accuracy.

In the following we first detail our concept of principal ordinate propagation for a single directional source (Sec. 3.1). Then we describe how to extend this scheme to environment illumination (Sec. 3.2) and local light sources (Sec. 3.3) by using multiple importance-sampled rectilinear and spherical propagation volumes, respectively. The propagation scheme is explained using radiance as the radiometric quantity; we assume all other quantities (such as irradiance from environment maps or intensity from point lights) to be converted accordingly. All frequently-used notation is summarized in Table 1.

3.1 Regular grids for directional light

The concept as well as the propagation scheme can be best explained for parallel (distant) light travelling along a direction \mathbf{d} through a region in space (Fig. 2, top). For this case we discretize the space into a uniform rectilinear grid similar to DOM; however, we make sure that one of its dimensions is aligned with \mathbf{d} . For every grid cell i , we store the directional distribution of light and its magnitude L_i (all computations are performed independently per-wavelength, which is omitted here for brevity). The main difference to DOM is that we represent both the *directional distribution* of light and the *phase function* using the HG distribution implicitly aligned with \mathbf{d} . To distinguish radiance anisotropy (directional distributions) from phase functions, we denote the HG parameter for the former as $a_i \in [-1, 1]$, and $g \in [0, 1]$ for the latter (we do not consider negative values of g because of physical implausibility of dominantly-backscattering




Figure 3: The propagation grid aligned with the direction of incidence is initialized with the attenuated radiance and an anisotropy parameter $a_i = 1$. During the propagation both radiance magnitude and anisotropy change towards lower anisotropy.

media). That is, the directional radiance of a grid cell centred at \mathbf{x}_i is $L(\mathbf{x}_i, \omega) = L_i \cdot f_{\text{hg}}(\mu, a_i)$, where f_{hg} is the HG function and $\mu = \omega \cdot \mathbf{d}$ is the cosine of the angle between a direction ω and the principal light direction \mathbf{d} . We assume that the medium is further characterized by its (spatially-varying) scattering coefficient σ_s and absorption coefficient σ_a ; these two quantities as well as the spatially-varying anisotropy of the phase function defined by the HG parameter g are wavelength-dependent and stored for every cell of the medium volume (which exists *independently* of the propagation volumes).

Conceptually, two grids are required in the propagation procedure. The first, *propagation grid*, stores the unpropagated (residual) energy; we will denote it as \mathbf{L} and its state at the iteration $m \in \{1..M\}$, where M is the total number of propagation iterations, as \mathbf{L}^m . The second, *accumulation grid* \mathbf{L}_{acc} , is needed to accumulate the energy transported through the medium over the course of the computation. Two options are available for implementing \mathbf{L}_{acc} : we could either store the overall radiance distribution that has passed through each cell during the propagation, or alternatively store only the observer-dependent out-scattered radiance at each iteration. We opted for the second approach, because storing the entire directional radiance distribution at each cell is much more expensive than just accumulating the outgoing radiance (which is essentially a single scalar value). Although this of course requires recomputing the solution on every observer position change, it is in agreement with our premise of a fully dynamic algorithm without relying on precomputations.

3.1.1 Grid initialization

At the beginning each propagation grid—which is scaled to span the entire medium (Fig. 2, top)—needs to be initialized by the incident radiance at each cell. As no scattering has been accounted for yet, the anisotropy is set to an HG coefficient of $a_i = 1$, an equivalent to the Dirac function in the direction \mathbf{d} (Fig. 3). The radiance magnitude L_i is set to the incident radiance $L_{\text{in}}(\mathbf{d})$ at \mathbf{x}_i , attenuated by absorption and out-scattering. That is, for every cell, we compute the transmittance T_i (from the point where light enters the medium, travelling along \mathbf{d} to \mathbf{x}_i) set to $L_i = L_{\text{in}}(\mathbf{d}) \cdot T_i$. Note that this can be efficiently computed using ray marching: as our grid is aligned with \mathbf{d} we can compute the transmittance incrementally along individual ‘slices’ of the grid along \mathbf{d} in a single pass.

3.1.2 Light energy propagation

In this section, we describe how to iteratively update the grid to simulate the propagation of light. We use a propagation stencil where the radiance of each grid cell is propagated to its 6 direct neighbours in every iteration. More specifically, we perform a gathering-type computation of how much radiance flows *into* each grid cell from its neighbours based on their radiance distributions




Figure 4: Left: Our polar parametrization of the solid sphere. The coloured patches correspond to the approximate solid angles subtended by the cells next to (green), in front (purple) and behind (orange) src . Middle: The HG cumulative function F_{hg} is used to integrate the radiance from the source cell flowing towards the destination cells (depicted as coloured patches of f_{hg} , for $g = 0.5$). Right: On the way the light undergoes scattering and is possibly reduced by absorption.

and then combine these contributions to yield the new distribution at that cell (Fig. 4, right). In the following we denote the neighbouring source cell with index src , and the target destination cell with dst .

Radiance magnitude contribution We first need to determine the amount of radiant energy that flows from cell src towards dst according to the radiance distribution in src . To this end, we efficiently compute the integral of $L(\mathbf{x}_{src}, \omega)$ over the solid angle subtended by dst (denoted as $\Omega_{src \rightarrow dst}$ below) using the closed form of the cumulative HG function $F_{\text{hg}}(\mu, g) = \int_{-1}^{\mu} f_{\text{hg}}(\mu', g) d\mu'$:

$$F_{\text{hg}}(\mu, g) = \frac{1-g^2}{4\pi g} \cdot \left(\frac{1}{(1+g^2-2g\mu)^{1/2}} - \frac{1}{1+g} \right). \quad (1)$$

By this we compute the radiance from src travelling towards to dst using the transmittance $T_{src \rightarrow dst}$ as

$$\Delta L_{src \rightarrow dst} = L_{src} \cdot T_{src \rightarrow dst} \cdot |\phi_1 - \phi_2| \cdot (F_{\text{hg}}(\cos \theta_1, a_{src}) - F_{\text{hg}}(\cos \theta_2, a_{src})) \quad (2)$$

using the following approximate parametrization for the subtended solid angle $\Omega_{src \rightarrow dst}$ (depending on mutual positions of src and dst):

$$(\theta_1, \theta_2, |\phi_1 - \phi_2|) = \begin{cases} (0, \frac{\pi}{4}, 2\pi) & dst \text{ in front of } src \\ (\frac{\pi}{4}, \frac{3\pi}{4}, \frac{\pi}{2}) & dst \text{ next to } src \\ (\frac{3\pi}{4}, \pi, 2\pi) & dst \text{ behind } src \end{cases} \quad (3)$$

(see Fig. 4, left for a sample illustration of the second case of Eq. 3). Since the HG distribution is rotationally-symmetric (Fig. 4, middle) only the absolute value of the difference of the azimuthal angles $|\phi_1 - \phi_2|$ is required. Note that here the transmittance $T_{src \rightarrow dst}$ accounts *just for absorption* that affects the radiance propagation on its way from src to dst . This is because our scheme treats scattering as a decrease of anisotropy and not as an extinction process, as we show below. In practice, we take the averaged absorption coefficients σ_a at the source and destination cells and the distance between their centres t , and apply the Beer-Lambert-Bouguer law; however, ray-marching with a small number of steps might potentially be required to integrate the absorption coefficient if the resolution of the propagation volume is much smaller than the medium grid.

Radiance anisotropy contribution Similarly to absorption attenuating the radiant energy flowing between neighbouring cells, the anisotropy of the energy propagated from src to dst will decrease due to scattering. In agreement with the radiance transfer equation, in our case this can be easily computed exploiting the self-convolution property of the HG distribution [22]: in a medium with scattering anisotropy of g the radiance anisotropy reduces to $a' = a \cdot g^{\sigma_s \cdot t}$ after

travelling a distance t (assuming a constant σ_s along this path). We obtain σ_s and t the same way as for computing $T_{src \rightarrow dst}$ above. The change of radiance anisotropy from src to dst is therefore

$$\Delta a_{src \rightarrow dst} = a_{src} \cdot g^{\sigma_s \cdot t}. \quad (4)$$

We can easily see that this formula cannot lead to an increase of anisotropy, since $g \in [0, 1]$. Additionally, in non-scattering media ($\sigma_s = 0$) the anisotropy will be preserved perfectly.

Combining contributions from neighbours Updating the radiance distribution at the cell dst entails accumulating the contributions from its six neighbours (indexed by src) as

$$L_{dst} = \sum_{src} \Delta L_{src \rightarrow dst}, \quad (5)$$

$$a_{dst} = \frac{\sum_{src} \Delta L_{src \rightarrow dst} \cdot \Delta a_{src \rightarrow dst}}{\sum_{src} \Delta L_{src \rightarrow dst}}. \quad (6)$$

While the radiant energy contributions simply need to be added up, the anisotropy is a weighted average of its neighbours, since the update has to yield an anisotropy value a_{dst} within the valid range. We discuss implications of Eq. 6 in Sec. 3.1.5.

3.1.3 Iterating the solution

The update procedure defined by Eqs. 5 and 6 is performed for every cell of \mathbf{L}^m to yield \mathbf{L}^{m+1} for every iteration m . Implementation-wise, this requires maintaining a second grid identical to the propagation grid and swapping these at each iteration.

Additionally, the results of every propagation iteration need to be accumulated in \mathbf{L}_{acc} by evaluating the updated distributions in \mathbf{L}^{m+1} :

$$L_{acc,i}^{m+1} = L_{acc,i}^m + L^{m+1}(\mathbf{x}_i, \mathbf{c} - \mathbf{x}_i) \quad (7)$$

$$= L_{acc,i}^m + L_i^{m+1} \cdot f_{hg}(\mu, a_i^{m+1}) \quad (8)$$

for every cell i . Here \mathbf{c} is the observer position and μ is therefore the dot product of \mathbf{d} and the view direction.

3.1.4 Upsampling and rendering

When the solution has converged after a sufficient number of iterations, using it for rendering is relatively straightforward. We employ ray-marching to integrate the incoming radiance for every camera ray using the common front-to-back emission-absorption model [21]. In this case the emission term corresponds to the scattered radiance accumulated in \mathbf{L}_{acc} .

As we discuss in Sec. 4, the typical resolutions used for the propagation grids need to be rather small (in most of our examples 20^3 or less) for performance reasons. In order to improve the rendering quality with such low grid resolutions it is desired to upsample them prior to their visualization. We use a 3D version of the joint bilateral upsampling [16] where the density field of the medium (i. e., the spatially varying scattering coefficient) is used as a guidance signal. Typically, the density field is significantly more detailed than the propagation volumes; this detail is “transferred” to the solution by the upsampling. According to our experiments, low-resolution propagation grids are usually sufficient for plausible results.

3.1.5 Discussion of the propagation scheme

Using the unimodal HG function with a single parameter to represent the directional distributions in light transport obviously means that there are distributions in a cell that cannot be represented well. On the other hand, we compensate for this by using multiple grids




Figure 5: Three examples of the local propagation behaviour. Left: all source cells exhibit strong forward scattering which is well-preserved by our propagation scheme. Centre: radiance anisotropy is reduced due to in-scattering from Source 2 which has isotropic radiance distribution. Right: light from Source 1 to destination is almost entirely absorbed. Light from Source 2 should then be deviated “upwards”, which our scheme cannot represent.

(see Sec. 3.2), which in turn can handle anisotropic phase functions significantly better than previous work thanks to the proposed propagation scheme. In comparison, an exceedingly large number of SH coefficients is required to represent highly anisotropic distributions, and this still does not prevent false scattering issues if a local propagation scheme is employed.

In this view the most heuristic step of our scheme is the recombination of reduced anisotropies from the neighbouring cells in Eq. 6. The logic behind this formulation is that the radiance distribution at dst will result from superposing the neighbouring distributions according to how much energy they contribute to dst . The main limitation of this approach lies in the fact that combining multiple HG distributions with different anisotropy values cannot generally be represented by any single HG distribution. Although we have experimented with fitting the resulting HG distribution to the combination of its neighbours in terms of least square error, we found that the simple weighted arithmetic average produces comparable results while keeping the computational cost of this core operation minimal. In addition, Eq. 6 very well preserves the anisotropy of light transported along the principal direction, thus greatly reducing false scattering effects.

Note that there are cases of very heterogeneous media where our approach might locally become too inaccurate (see Fig. 5). If light along the principal direction undergoes strong absorption, while light from other directions does not, the resulting light distribution should possibly become skewed, which cannot be represented within our framework. Although this is obviously a failure case of our representation, occurrences of such strong absorption fluctuations are comparatively rare, and more importantly the resulting radiance magnitude in these cases is typically very small (therefore having little impact on the resulting image). Also note that with multiple propagation volumes we can actually reproduce complex multimodal radiance distributions, despite each grid being composed of unimodal HG distributions.

3.2 Environment lighting: Multiple propagation grids

In the previous section we have described our approach for a single directional light source. In order to account for environmental lighting (typically modelled by an environment map), we need to use multiple grids oriented in different principal directions. In the following we discuss how to choose these directions and, as every grid accounts for light from a finite solid angle, how to prefilter the respective incident radiance to avoid singularity artifacts (see Fig. 6).

Prefiltering A straightforward approach is importance-sampling the environment map to obtain N directions, \mathbf{d}_n , each carrying an




Figure 7: Importance propagation improves overall radiance distribution across the medium and visibility of bright regions behind. This especially holds for high-albedo media with strong scattering anisotropy (here $g = 0.98$) and when using a low number of ordinates (27 here).




Figure 6: The effect of prefiltered initialization on a thin, strongly-scattering medium with increasing anisotropy (left to right). Without prefiltering (top) the individual ordinates become apparent. Using prefiltering (bottom) the resulting images become much smoother and yield the expected appearance (more anisotropic slabs appear more transparent). Note that our technique is energy-conserving (as opposed to, e.g., singularity clamping in instant radiosity).

energy corresponding to its associated portion of the directional domain Ω_n . We can account for the shape of Ω_n when determining the initial directional radiance distributions (parameter a_i in Sec. 3.1.1). Recall that the anisotropy parameter of f_{hg} represents the average cosine of the distribution. We can therefore approximate the initial $a_{n,i} = \int_{\Omega_n} -\mathbf{d}_n \cdot \boldsymbol{\omega} d\omega / \|\Omega_n\|$, the average cosine between \mathbf{d}_n and the directions in Ω_n and use this value for the grid initialization. In practice, $a_{n,i}$ can be approximated without the integration over Ω_n for each ordinate or without even knowing the shape of Ω_n . As we importance-sample the environment map, the importance of the ordinate n is proportional and (up to a factor) very similar to the actual solid angle of Ω_n . Therefore, we use a heuristic that maps the importance $w_n \in (0, 1)$ to anisotropy as $a_{n,i} = (1 - w_n/N)^\beta$: important ordinates are denser in the directional domain and will have small solid angle and high anisotropy, less important ordinates are more sparse, will have larger solid angles and low anisotropy. The scalar factor $\beta > 0$ defines the proportionality and currently needs to be tuned empirically once for each environment map; from our experiments this is a simple and quick task.

Importance propagation The described sampling scheme can be further improved by considering how much illumination from different directions actually contributes to the image. To this end, we introduce an additional *importance propagation* step before sampling




Figure 8: Comparison of our radial propagation to a Monte-Carlo reference for a uniform spherical medium (radius 2.5 m, $\sigma_s = \{0.8, 1, 1.3\} \text{ m}^{-1}$ and unit albedo). The resolution of the radial propagation grid was 32^3 . Our solution differs from the reference mainly due to low (but for this propagation type still present) false scattering, in particular with low anisotropy values. We found that this can be reduced by artificially increasing g , if a specific appearance is desired.

the environment map: we use a regular grid (perspective-warped into the camera frustum and oriented along the view direction) and propagate importance from the camera through the medium. Thanks to the duality of light transport this is equivalent to the radiance propagation as described before. The result of this propagation is a directional importance distribution stored in the grid cells. By ray-marching this grid we project the importance into the directional domain and create a directional importance ‘map’ that aligns with the environment map. We then sample the environment map according to its product with the importance map. We show that in certain situations this step improves the sampling result, especially when a low number of propagation grids is used (see Sec. 4 and Fig. 7). It is also quite cost-effective, since the directional importance function is typically very smooth and therefore only low resolutions for the propagation grid and the directional map are required (all our examples use the resolutions of 16^3 and 32×16 respectively).

3.3 Radial grids for local light sources

In order to extend our method to local light sources, we use spherical grids with two angular coordinates and a radial coordinate which is again aligned with the initial principal directions of the point source (Fig. 2). To obtain more isotropic cell shapes, the spacing of shells along the radial coordinate grows exponentially (in proportion to the radial segment length at a given radius). For parametrizing the spherical domain we use the octahedron parametrization [27] mainly as it is simple, provides reasonably uniform sampling, and above all, it discretizes the domain into a 2D square where every cell has four natural neighbours (plus two along the radial axis), similar to rectilinear grids. The resulting grid is thus topologically equivalent to rectilinear grid and albeit not being uniform, it allows




Figure 9: Workflow of the presented algorithm for a single directional light. For distant environment illumination the volumetric part of the pipeline is very similar, with the exception of rectilinear grids being used to propagate illumination from distant ordinates instead of the combination of VPLs and spherical grids.



Figure 10: In media like clouds the scattering anisotropy plays a significant role in their appearance, thus the common assumption of isotropic scattering prevents a believable rendition of such media. The clouds are rendered by the described method at 12 Hz using 64 ordinates and 20^3 grid resolution for each of them, with 15 propagation iterations. The scattering anisotropy was set to $g = 0.96$.

us to approximately treat the space as locally Euclidean and obtain plausible results again using virtually the same propagation scheme as before. The main difference in the propagation is that we have to account for the quadratic fall-off: although we base our propagation on radiance, we have to explicitly compensate for the varying grid cell sizes resulting from the non-uniform shell spacing. To this end, we scale the radiance when propagating along the principal direction in proportion to the radial coordinate spacing. A simple demonstration of this propagation type for a point light in a simple homogeneous spherical medium is shown in Fig. 8.

Instant radiosity Given the ability to use local point lights, we can use instant radiosity [14] methods, which represent complex illumination as a collection of point lights, to simulate surface-to-volume light transport. Normally these VPLs are obtained from random walks through the scene. In our interactive setting, we generate VPLs using a reflective shadow map (RSM) [4] for every primary light. We importance-sample these RSMs according to surface albedo and (attenuated) irradiance, aiming at keeping the total number of VPLs low. The reflected radiance is then used to initialize the radial propagation grids. Prefiltering can be done in the same way as for environment maps: VPLs with a large importance have a high initial anisotropy and vice versa. Similar to surface lighting, we can use clamping to reduce any remaining singularities [6]. Fig. 9 depicts the pipeline of the algorithm when propagating scattering from one directional light and VPLs generated from its RSM.

4 RESULTS

All results were computed on a laptop PC with a 2.0 GHz Intel Core i7 CPU, 16 GB of RAM and an NVidia GeForce GTX 485 Mobile card with 2 GB of VRAM. In all our measurements we use the framebuffer resolution of 800×600 in order to let the computation time be dominated by the propagation rather than ray-marching. Resolutions of the medium density datasets are typically in the order of tens in each dimension (but effectively enhanced by the procedural noise). Although the number of propagation iterations needs to be chosen empirically at the moment, in general we found that amounts similar to the propagation grid resolution along the propagation dimension is sufficient (around 10–20 in our examples). Other specific scene details are provided in the caption of each discussed figure.

We first tested our method for cloudy media with high scattering anisotropy in comparison to their isotropic versions (Fig. 10). It can be seen that our propagation scheme handles both cases well. Interestingly, grid resolutions as well as computation times required to render plausible participating media are rather insensitive to its anisotropy, i. e., anisotropic media render as fast as isotropic media. Although a larger number of ordinates is required to reproduce high-anisotropy effects, this additional effort is usually compensated by a decreased complexity of the spatial radiance distribution, which enables using coarser propagation grids.




Figure 11: Comparison of our principal ordinates propagation to SHDOM and a Monte-Carlo reference, for a smoke plume 10 m across with $\sigma_s = \{2.9, 3.6, 4.2\} \text{ m}^{-1}$, $\sigma_a = \{3.4, 3.35, 3.4\} \text{ m}^{-1}$ and $g = 0.9$ using the “Uffizi” environment map as illumination. For our technique we used 64 and 125 principal ordinates, grid resolutions of 20^3 and 50^3 , 10 and 30 propagation iterations, respectively. For SHDOM we have used 5 and 10 bands to represent the directional radiance distribution in each cell and the same grid resolutions. SHDOM required a strong prefiltering to avoid ringing and due to false scattering fails to reproduce the high scattering anisotropy. Our method compares well to the reference solution, and even with low-quality settings it matches the overall appearance.

Next, we compare our approach to an unbiased Monte-Carlo reference, as well as SHDOM, in Fig. 11. It is apparent that the described artifacts prevent SHDOM from handling anisotropic media correctly, despite being theoretically capable to do so.

The effect of using different numbers of principal ordinates is shown in Fig. 12. It can be seen that the discretization becomes apparent only with very few ordinates. The importance propagation usually helps to alleviate this by sampling those directions which will influence the solution most significantly. As Fig. 7 demonstrates, this is most likely the opposite side of the medium, suggesting that a simpler empirical heuristic could potentially work in certain cases.

One of the main shortcomings of the importance propagation is its potential temporal incoherency, mostly manifested by temporal flickering. For this reason we filter the importance map both spatially and temporarily, which, however, mainly distributes the incoherency over time. One of our main targets for future work is therefore improving this by explicitly enforcing temporal coherency when the sampled light sources relocate due to camera movement.

Prefiltering helps to improve the rendering quality in most scenarios and we used it to generate all results throughout the paper. It is particularly indispensable for media with an optical thickness insufficient to blur the sampled illumination, e. g., as in Fig. 6, where singularity-like artifacts would appear otherwise. Prefiltering removes these artifacts but still allows to perceive features of the background illumination, thanks to its adaptivity.

Finally in general, Fig. 1 and Fig. 13 show our propagation scheme for both regular and radial grids used to render multiple scattering effects in the volume stemming from direct illumination (light shafts), environment lighting, and indirect surface illumination (using virtual local light sources) under fully-dynamic conditions.

5 CONCLUSION

We propose a novel discrete ordinates method capable of computing light transport in heterogeneous participating media exhibiting light scattering of virtually arbitrary anisotropy. The method does not require any precomputations, which makes it suitable even for simulating dynamic and evolving media. Our representation also adapts to and prefilters the incident lighting. Radiance is represented by the Henyey-Greenstein distribution, and propagated by our novel scheme in volumes oriented along the principal light directions.

In general the steps of the proposed method are physically-plausible (please refer to the supplementary materials for further details). The employed empirical heuristics introduce a certain bias but allow us to make design decisions that result in a near-realtime performance on contemporary graphics hardware.




Figure 12: The smoke dataset with an increasing number of ordinates using the “kitchen” environment map ($g = 0.9$, 20^3 grid resolution, 10 propagation iterations). Accounting for importance improves the results, mainly if low numbers of principal ordinates are used. The typical setting we use is shown on the bottom-centre and takes 10 ms for importance propagation, 4 ms for determining the ordinates, 7 ms for grid initialization, 50 ms for propagation, 10 ms for residuum propagation, 3 ms for upsampling and 11 ms for ray-marching.

The decomposition into a finite number of directions for distant light can only be successful if the variation of the initial light distribution is not too high; this however holds for the HDR environment maps we used in our examples. In addition our prefiltered initialization can be used to avoid discretization artifacts in favour of a smooth approximation. Another limitation that we share with most variants of DOM is the handling of (surface) boundaries. In volumes with high density gradients (close to opaque surfaces) the light distribution might not be faithfully reproduced by the HG basis aligned with the initial light direction. Also the resolution of every principal grid is limited and therefore the general limitations of discrete sampling apply: for finer details more resolution is required. Fortunately, the upsampling and prefiltering help to defer this problem and for typical volume data sets moderate propagation grid resolutions of 8^3 – 20^3 have shown to be sufficient to handle a wide range of illumination conditions and medium properties.

As future work, we would like to extend our propagation to work on hierarchical or nested grids to handle higher details in media as well as illumination. In general, we believe that the effect of complex lighting on dynamic participating media is an exciting visual phenomenon that deserves more dedicated research, e. g., to better understand human perception of volumetric light or the artistic practice applied to depict it.




Figure 13: A scene with both animated medium and illumination, combining scattering from directional and local virtual light sources (running at 9Hz including the generation of the 125 VPLs used to render indirect illumination from surfaces; the medium has a size of 20^3 m with $\sigma_i = \{3.2, 3.3, 3.4\} \text{ m}^{-1}$, $\sigma_a = \{1.15, 1.2, 1.3\} \text{ m}^{-1}$ and $g = 0.7$). The grid size for the directional light is $128^2 \times 16$, with the 16-cell axis oriented along the light shafts (i. e., along the principal ordinate). The radial grids have a resolution of 8^3 each. We use these settings for local light sources in all our examples; note how even this small resolution proves to be sufficient for plausible results.

ACKNOWLEDGEMENTS

We would like to thank Chuong Nguyen, Peter Vangorp and Yulia Gryaditskaya for proofreading, Anton Kaplanyan for discussions, and the anonymous reviewers for their feedback. This work has been supported by the Max Planck Center for Visual Computing and Communication.

REFERENCES

- [1] M. Billeter, E. Sintorn, and U. Assarsson. Real-time multiple scattering using light propagation volumes. In *Proc. I3D*, 2012.
- [2] A. Bouthors, F. Neyret, N. Max, E. Bruneton, and C. Crassin. Interactive multiple anisotropic scattering in clouds. In *Proc. I3D*, 2008.
- [3] S. Chandrasekhar. *Radiative transfer*. Dover Publications, 1960.
- [4] C. Dachsbacher and M. Stamminger. Reflective Shadow Maps. In *Proc. I3D*, pages 203–213, 2005.
- [5] T. Engelhardt and C. Dachsbacher. Epipolar sampling for shadows and crepuscular rays in participating media with single scattering. In *Proc. I3D*, 2010.
- [6] T. Engelhardt, J. Novák, T.-W. Schmidt, and C. Dachsbacher. Approximate bias compensation for rendering scenes with heterogeneous participating media. *Comp. Graph. Forum*, 31(7):2145–2154, 2012.
- [7] R. Fattal. Participating media illumination using light propagation maps. *ACM Trans. Graph.*, 28:7:1–7:11, 2009.
- [8] R. Geist, K. Rasche, J. Westall, and R. Schalkoff. Lattice-Boltzmann lighting. In *Proc. EGSR*, pages 355–362, 2004.
- [9] L. G. Henyey and J. L. Greenstein. Diffuse radiation in the Galaxy. *Astrophysical Journal*, 93:70–83, 1941.
- [10] W. Jarosz, C. Donner, M. Zwicker, and H. W. Jensen. Radiance caching for participating media. *ACM Trans. Graph.*, 27:7:1–7:11, 2008.
- [11] W. Jarosz, D. Nowrouzezahrai, I. Sadeghi, and H. W. Jensen. A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Trans. Graph.*, 30:5:1–5:19, 2011.
- [12] H. W. Jensen and P. H. Christensen. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proc. SIGGRAPH*, pages 311–320, 1998.
- [13] A. Kaplanyan and C. Dachsbacher. Cascaded light propagation volumes for real-time indirect illumination. In *Proc. I3D*, 2010.
- [14] A. Keller. Instant radiosity. In *Proc. SIGGRAPH*, pages 49–56, 1997.
- [15] J. Kniss, S. Premože, C. Hansen, and D. Ebert. Interactive translucent volume rendering and procedural modeling. In *Proc. Visualization*, pages 109–116, 2002.
- [16] J. Kopf, M. F. Cohen, D. Lischinski, and M. Uyttendaele. Joint bilateral upsampling. *ACM Trans. Graph.*, 26(3):96, 2007.
- [17] J. Krüger, K. Bürger, and R. Westermann. Interactive screen-space accurate photon tracing on GPUs. In *Proc. EGSR*, 2006.
- [18] E. P. Lafortune and Y. D. Willems. Rendering participating media with bidirectional path tracing. In *Proc. EGWR*, pages 91–100, 1996.
- [19] G. Liktor and C. Dachsbacher. Real-time volume caustics with adaptive beam tracing. In *Proc. I3D*, 2011.
- [20] T. Lokovic and E. Veach. Deep shadow maps. In *Proc. SIGGRAPH*, pages 385–392, 2000.
- [21] N. Max. Optical models for direct volume rendering. *IEEE Trans. Vis. and Computer Graphics*, 1:99–108, 1995.
- [22] N. Max, G. Schussman, R. Miyazaki, K. Iwasaki, and T. Nishita. Diffusion and multiple anisotropic scattering for global illumination in clouds. *J. WSCG*, 1–3:277–284, 2004.
- [23] S. G. Narasimhan, M. Gupta, C. Donner, R. Ramamoorthi, S. K. Nayar, and H. W. Jensen. Acquiring scattering properties of participating media by dilution. *ACM Trans. Graph.*, 25:1003–1012, 2006.
- [24] J. Novák, D. Nowrouzezahrai, C. Dachsbacher, and W. Jarosz. Progressive virtual beam lights. *Comp. Graph. Forum*, 31:1407–1413, 2012.
- [25] J. Novák, D. Nowrouzezahrai, C. Dachsbacher, and W. Jarosz. Virtual ray lights for rendering scenes with participating media. *ACM Trans. Graph.*, 29(4):60:1–60:11, 2012.
- [26] M. Pauly, T. Kollig, and A. Keller. Metropolis light transport for participating media. In *Proc. EGWR*, pages 11–22, 2000.
- [27] E. Praun and H. Hoppe. Spherical parametrization and remeshing. *ACM Trans. Graph.*, 22(3):340–49, 2003.
- [28] Z. Ren, K. Zhou, S. Lin, and B. Guo. Gradient-based interpolation and sampling for real-time rendering of inhomogeneous, single-scattering media. Technical Report MSR-TR-2008-51, Microsoft Research, 2008.
- [29] H. E. Rushmeier and K. E. Torrance. The zonal method for calculating light intensities in the presence of a participating medium. *Computer Graphics (Proc. SIGGRAPH)*, 21:293–302, 1987.
- [30] M. Salvi, K. Vidimče, A. Lauritzen, and A. Lefohn. Adaptive volumetric shadow maps. In *Proc. EGSR*, pages 1289–1296, June 2010.
- [31] P. Sloan, J. Kautz, and J. Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *ACM Trans. Graph.*, 2002.
- [32] J. Stam. Multiple scattering as a diffusion process. In *Proc. EGWR*, pages 41–50, 1995.
- [33] Y. Wang, J. Wang, N. Holzschuch, K. Subr, J.-H. Yong, and B. Guo. Real-time rendering of heterogeneous translucent objects with arbitrary shapes. *Comp. Graph. Forum*, 29:497–506, 2010.
- [34] K. Zhou, Z. Ren, S. Lin, H. Bao, B. Guo, and H.-Y. Shum. Real-time smoke rendering using compensated ray marching. In *ACM Trans. Graph.*, pages 36:1–36:12, 2008.