

# Principal-Ordinates Propagation for Real-Time Rendering of Participating Media

Oskar Elek<sup>a,b</sup>, Tobias Ritschel<sup>a,b</sup>, Carsten Dachsbacher<sup>c</sup>, Hans-Peter Seidel<sup>a,b</sup>

<sup>a</sup>Max Planck Institut Informatik

<sup>b</sup>MMCI Cluster of Excellence, Saarland University

<sup>c</sup>Karlsruhe Institute of Technology

---

## 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 Principal-Ordinates Propagation, a novel finite-element method that achieves real-time 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 domains. 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 resulting approach is biased but physically plausible, and largely reduces the rendering artifacts inherent to existing finite-element methods. At the same time it allows for virtually arbitrary scattering anisotropy, albedo, and other properties of the simulated medium, without requiring any precomputation.

*Keywords:* participating media; light scattering; natural phenomena; real-time rendering; physically-based rendering; finite elements

---

## 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, 23], 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 algorithm for plausible real-time 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 rectangular 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. [15, 1]), 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* (numerical dissipation) and *ray effect* (misalignment errors) 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 propose the concept of Principal-Ordinates Propagation (POP), a deterministic finite-element scheme suitable for real-time simulation of anisotropic light transport in heterogeneous participating media (Sec. 3).
- The theory of iterative light propagation in a uniform Euclidean grid using a minimal unimodal propagation basis and explicit alignment with the illumination direction to minimize propagation artifacts and maintain light directionality (Sec. 4).
- An extension of the propagation scheme to handle environmental illumination by decomposing it in a set of discrete directions. This includes several new steps, namely specialized prefiltering, importance propagation, and a separate propagation of isotropic residual energy (Sec. 5).
- An extension to local light sources via spherical grids, enabling the integration of instant radiosity to simulate light interaction between solid objects and the medium (Sec. 6).
- Finally we analyse our approach in a number of diverse scenarios, demonstrating its versatility (Sec. 7).

## 2. Related work

*Offline methods.* A range of different approaches has been presented to compute solutions to the radiance transport equation for participating environments [3, 23]. 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 trac-



Figure 1: Dense smoke exhibiting strong multiple anisotropic scattering produced by a steam locomotive under complex environment illumination. Our approach renders it dynamically without any precomputations at 25 Hz (NVIDIA GeForce GTX 770).

ing [20] and Metropolis light transport [29] 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 [12], photon mapping [14, 13] or virtual point lights [8]. However, these methods typically do not handle highly anisotropic scattering very well, even with recent improvements [27, 28], and their performance is often far from interactive.

*Finite-element methods.* Finite-element methods, including volume radiosity [33], the discrete ordinates method (DOM) [3], light diffusion [36], and lattice-Boltzmann transport [10] 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 [9] significantly reduce the artifacts, but are still limited to rather low scattering 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 [11] coefficient  $g \approx 0.9$  or more [26]). 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 [32, 7], volume caustics [19, 21], shadows [22, 34], and clouds [2]. Various approaches can also be found in visualization literature, e.g. half-angle slicing [17] which empirically computes forward scattering for volume visualization. Sometimes precomputation is used to speed up the rendering of heterogeneous translucent objects [35, 37] or smoke using compensated ray marching [39]. In contrast, we target general multiple scattering in participating media without any precomputation or focus on a particular phenomenon.

We extend the work of Elek et al. [5], building primarily on the concept of DOM [3] and the more recent light propagation

volumes [15, 1]. These approaches are attractive for interactive applications as their grid-based local propagation schemes allow for easy parallel implementation on contemporary GPUs. Our work also shares similarities with the finite-difference time domain method [25], however we only consider the radiance amplitude and in general concentrate on efficiency.

Virtually all existing variants or extensions of DOM use a single scene-aligned propagation grid, where every 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 effects and false scattering artifacts (see [9]). 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 one. This enables using a unimodal representation of the angular energy distribution around the principal direction in each grid cell.

### 3. Principal-Ordinates Propagation

The core idea of our method is to reduce the main drawbacks of previous grid-based iterative methods, namely false scattering and ray effects. These problems stem from the fact that the propagation domain is generally not aligned with the prominent light transport directions. We propose to remedy these issues 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) [11] 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




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.

134 well approximate their directionality; for local light sources we  
135 propose to use spherical grids centred around them.

136 These choices inherently assume that the principal directions  
137 can be derived from the initial radiance distribution and do not  
138 change strongly when light travels through the medium. How-  
139 ever, such variation might occur if the density of the simulated  
140 medium changes abruptly. Still we deem this to be a necessary  
141 compromise if speed is the priority, and as we discuss in Sec. 4.5,  
142 violating this assumption does not cause our algorithm to fail,  
143 but only leads to a gradual decrease of accuracy.

144 In the following sections we first detail our concept of Principal-  
145 Ordinates Propagation for a single directional source (Sec. 4).  
146 Then we describe how to extend this scheme to environment  
147 illumination (Sec. 5) and local light sources (Sec. 6) by using  
148 multiple, importance-sampled, rectilinear and spherical prop-  
149 agation volumes respectively. The propagation scheme is ex-  
150 plained using radiance as the radiometric quantity; we assume  
151 all other quantities (such as irradiance from environment maps  
152 or intensity from point lights) to be converted accordingly. All  
153 frequently-used notation is summarized in Table 1.

#### 154 4. Rectilinear grids for directional light

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




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

171  $\mathbf{x}_i$  is  $L(\mathbf{x}_i, \omega) = L_i \cdot f_{\text{hg}}(\mu, a_i)$ , where  $f_{\text{hg}}$  is the HG function and  
172  $\mu = \omega \cdot \mathbf{d}$  is the cosine of the angle between a direction  $\omega$  and the  
173 principal light direction  $\mathbf{d}$ . We assume that the medium is further  
174 characterized by its (spatially-varying) scattering coefficient  $\sigma_s$   
175 and absorption coefficient  $\sigma_a$ ; these two quantities as well as the  
176 spatially-varying anisotropy of the phase function defined by the  
177 HG parameter  $g$  are wavelength-dependent and stored for every  
178 cell of the medium volume (which exists *independently* of the  
179 propagation volumes).

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

##### 197 4.1. Grid initialization

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

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

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

## 4.2. 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 six direct neighbours in every iteration. Specifically, we perform a more GPU-friendly gathering-type computation of how much radiance flows *into* each grid cell from its neighbours based on their radiance distributions, 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  $\text{src}$ , and the target destination cell with  $\text{dst}$ .

**Radiance magnitude contribution.** We first need to determine the amount of radiant energy that flows from cell  $\text{src}$  towards  $\text{dst}$  according to the radiance distribution in  $\text{src}$ . To this end, we efficiently integrate  $L(\mathbf{x}_{\text{src}}, \boldsymbol{\omega})$  over the solid angle subtended by  $\text{dst}$  (denoted as  $\Omega_{\text{src} \rightarrow \text{dst}}$  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  $\text{src}$  travelling towards  $\text{dst}$  using the transmittance  $T_{\text{src} \rightarrow \text{dst}}$  as

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

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

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

(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




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)  $\text{src}$ . Middle: The HG cumulative function  $F_{\text{hg}}$  is used to integrate the radiance from the source cell flowing towards the destination cells (depicted as coloured patches of  $f_{\text{hg}}$ , for  $g = 0.5$ ). Right: On the way the light undergoes scattering and is possibly reduced by absorption.

transmittance  $T_{\text{src} \rightarrow \text{dst}}$  accounts *just for absorption* that affects the radiance propagation on its way from  $\text{src}$  to  $\text{dst}$ . 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  $\sigma_a$  at the source and destination cells and the distance between their centres  $t$ , and apply the Beer-Lambert-Bouguer law. To avoid aliasing if the resolution of the propagation grid is much smaller than the medium volume resolution, the medium parameters need to be sampled from a downscaled version of the volume. On GPU, downscaling is a very fast operation, as it corresponds to building a small number of MIP map levels (depending on the ratio between the grid resolutions, but usually 1 or 2).

**Radiance anisotropy contribution.** Similarly to absorption attenuating the radiant energy flowing between neighbouring cells, the anisotropy of the energy propagated from  $\text{src}$  to  $\text{dst}$  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 [24]: 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  $\sigma_s$  along this path). We obtain  $\sigma_s$  and  $t$  the same way as for computing  $T_{\text{src} \rightarrow \text{dst}}$  above. The change of radiance anisotropy from  $\text{src}$  to  $\text{dst}$  is therefore

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

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 directionality will be preserved perfectly.

**Combining contributions from neighbours.** Updating the radiance distribution at the cell  $\text{dst}$  entails accumulating the contributions from its six neighbours (indexed by  $\text{src}$ ) as

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

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

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  $a_{\text{dst}}$  within the valid range. We discuss implications of Eq. 6 in Sec. 4.5.

### 245 4.3. Iterating the solution

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

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

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

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

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

### 253 4.4. Upsampling and rendering

254 When the solution has converged after a sufficient number of  
 255 iterations, using it for rendering is relatively straightforward.  
 256 We employ ray-marching to integrate the incoming radiance  
 257 for every camera ray using the common front-to-back emission-  
 258 absorption model [23]. In this case the emission term corre-  
 259 sponds to the scattered radiance accumulated in  $\mathbf{L}_{acc}$ .

260 As we discuss in Sec. 7, the typical resolutions used for the prop-  
 261 agation grids need to be rather small (in most of our examples  
 262  $20^3$  or less) for performance reasons. In order to improve the  
 263 rendering quality with such low grid resolutions it is desired to  
 264 upsample them prior to their visualization. We use a 3D version  
 265 of the joint bilateral upsampling [18] where the density field of  
 266 the medium (i.e. the spatially varying scattering coefficient) is  
 267 used as a guidance signal. Typically, the density field is signifi-  
 268 cantly more detailed than the propagation volumes; this detail  
 269 is “transferred” to the solution by the upsampling. According to  
 270 our experiments, low-resolution propagation grids are usually  
 271 sufficient for plausible results.

### 272 4.5. Discussion of the propagation scheme

273 Using the unimodal HG function with a single parameter to rep-  
 274 resent the directional distributions in light transport obviously  
 275 means that there are distributions in a cell that cannot be repre-  
 276 sented well. On the other hand, we compensate for this by using  
 277 multiple grids (see Sec. 5), which in turn can handle anisotropic  
 278 phase functions significantly better than previous work thanks  
 279 to the proposed propagation scheme. In comparison, an exceed-  
 280 ingly large number of SH coefficients is required to represent  
 281 highly anisotropic distributions, and this still does not prevent  
 282 false scattering issues if a local propagation scheme is employed.

283 The most heuristic step of our scheme is the recombination of  
 284 reduced anisotropies from neighbouring cells in Eq. 6 (which  
 285 we further elaborate on in the supplementary materials). The  
 286 logic behind this formulation is that the radiance distribution




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.

287 at  $dst$  will result from superposing the neighbouring distribu-  
 288 tions according to how much energy they contribute to  $dst$ . The  
 289 main limitation of this approach lies in the fact that combin-  
 290 ing multiple HG distributions with different anisotropy values  
 291 cannot generally be represented by any single HG distribution.  
 292 Although we have experimented with fitting the resulting HG  
 293 distribution to the combination of its neighbours in terms of  
 294 least square error, we found that the simple weighted arithmetic  
 295 average produces comparable results while keeping the compu-  
 296 tational cost of this core operation minimal. In addition, Eq. 6  
 297 very well preserves the anisotropy of light transported along the  
 298 principal direction, thus greatly reducing false scattering effects.

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

## 312 5. Multiple propagation grids for environment lighting

313 In the previous section we have described our approach for a  
 314 single directional light source. In order to account for environ-  
 315 mental lighting (typically modelled by an environment map),  
 316 we need to use multiple grids oriented along different principal  
 317 directions.

318 In this section we discuss how to choose these directions and,  
 319 as every grid accounts for light from a finite solid angle, how  
 320 to prefilter the respective incident radiance to avoid singularity  
 321 artifacts (see Fig. 6). We further describe how the multiple  
 322 propagation grids are combined together for rendering. Finally  
 323 we present an additional (optional) step in the pipeline of our  
 324 algorithm, which allows splitting the propagation into two stages,

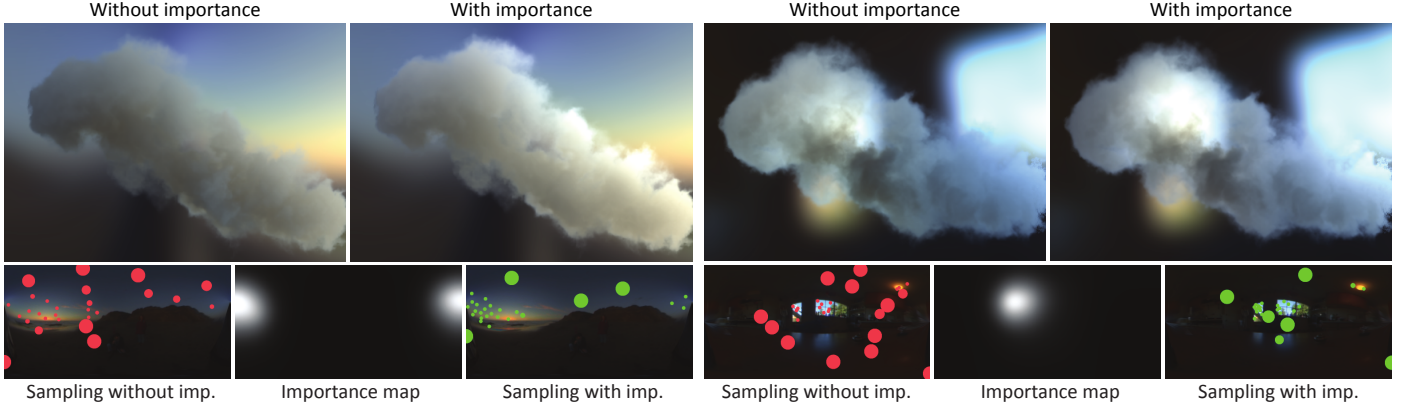


Figure 7: Importance propagation improves overall radiance distribution across the medium and visibility of bright regions behind. This especially holds for high-albedo media with strong scattering anisotropy (here  $g = 0.98$ ) and when using a low number of ordinates (27 here).




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).

anisotropic and isotropic, greatly improving the convergence for media with very high albedo values.

### 5.1. Prefiltering

A straightforward approach is importance-sampling the environment map to obtain  $N$  directions,  $\mathbf{d}_n$ , each carrying an energy corresponding to its associated portion of the directional domain  $\Omega_n$ . We can account for the shape of  $\Omega_n$  when determining the initial directional radiance distributions (parameter  $a_i$  in Sec. 4.1). Recall that the anisotropy parameter of  $f_{\text{hg}}$  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  $\Omega_n$  and use this value for the grid initialization. In practice,  $a_{n,i}$  can be approximated without the integration over  $\Omega_n$  for each ordinate or without even knowing the shape of  $\Omega_n$ . 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  $\Omega_n$ . 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 angles

and high anisotropy, while less important ordinates are more sparse, and 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 experience this is a simple and quick task.

### 5.2. 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 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, mainly when a low number of propagation grids is used (see Sec. 7 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  $32 \times 16$  respectively).

### 5.3. Merging multiple grids

Computing the propagation for each of the  $N$  principal ordinates yields a separate, view-dependent accumulation grid  $\mathbf{L}_{\text{acc},n}$  (Sec. 4). Although it is possible to visualize these directly, this is very inefficient as each grid in the set would have to be accessed at every ray-marching step.

Because of this we instead opt to combine all  $\mathbf{L}_{\text{acc},n}$  into a single medium-aligned grid, prior to upsampling and visualization




Figure 8: Evaluation of the isotropic residuum (IR) propagation. We show the dragon dataset illuminated by two different environment maps, from two opposing viewpoints for each to demonstrate the view dependence of the resulting light distributions. The used medium (milk [26],  $\sigma_s = \{0.91, 1.07, 1.25\} \text{ m}^{-1}$ ,  $g = 0.95$ ) has an albedo over 99.9% across the visible spectrum, making it a difficult material to render because of the high number of iterations necessary to converge. The full propagation (64 ordinates,  $16^3$  grid each) requires about 100 iterations to converge, taking 120 ms (#1) and 90 ms (#2) respectively. In comparison, our heuristic (Eq. 6) switches to isotropic propagation after  $m' = 20$  full anisotropic iterations ( $\epsilon_a = 0.1$ ,  $\Delta x = 0.25 \text{ m}$ ). The IR propagation requires additional 20 iterations in a single  $32^3$  propagation grid. The combined propagation time in this case was 28 ms (25 ms anisotropic and 3 ms isotropic) for illumination #1 and 24 ms (20.5 ms anisotropic and 3.5 ms isotropic) for #2; this is about a 4-fold speed-up compared to the full propagation, with negligible visual difference.

380 (Sec. 4.4). We usually use double the resolution of the individual  
 381 propagation grids, since these are oriented arbitrarily in space  
 382 (and therefore not increasing the resolution would result in undersampling).  
 383 This is however still a very fast step that also  
 384 allows the remainder of the pipeline to stay virtually identical to  
 385 the single-ordinate setting.

#### 386 5.4. Isotropic residuum

387 We assume the solution to be converged when all  $\mathbf{L}_n^m$  are below  
 388 a small threshold  $\epsilon_L$ . This can however take a large number  
 389 of iterations for high-albedo media, a problem inherent to all  
 390 finite-element transport methods. On the other hand, we can  
 391 observe that scattering reduces the anisotropy of the radiance  
 392 distribution and we can treat the propagation as (near-)isotropic  
 393 as soon as  $|a_{n,i}| < \epsilon_a \forall i, \forall n$ , for a small anisotropy threshold  $\epsilon_a$ .

394 As soon as all propagation grids fulfil this criterion the energy  
 395 from them can be merged into a single grid aligned with the  
 396 medium, as there is no directionality present anymore. This is  
 397 similar to merging the accumulation grids (Sec. 5.3), except that  
 398 here the propagation grids are merged as well and the propaga-  
 399 tion process switches to isotropic scattering (i.e. the anisotropies  
 400  $a_{n,i}$  need not be maintained anymore). This decreases the propa-  
 401 gation costs tremendously, as from this point it is performed just  
 402 for a single global grid instead of one grid per principal ordinate.

In practice, we determine the iteration  $m'$  when we can switch to  
 the cheaper isotropic propagation based on the maximum radiance  
 anisotropy,  $\hat{a}$  (for a directional light  $\hat{a} = 1$ , but prefiltering  
 can lower it, to our benefit), and phase function anisotropy  $g$ .  
 Both of these parameters are determined at the initialization.  
 Making use of the the HG self-convolution property (Sec. 4.2),  
 from these values we can approximate the distance  $t$  that light  
 has to travel such that its anisotropy falls below  $\epsilon_a$  with

$$\hat{a} \cdot g \bar{\sigma}_s^t = \epsilon_a, \quad (9)$$

where  $\bar{\sigma}_s$  is the average scattering coefficient in the medium.  
 Furthermore, we know that the travel distance depends on the  
 average grid spacing  $\Delta x$  and the number of iterations, i.e.  $t =$   
 $m \cdot \Delta x$ , and obtain:

$$m' = \frac{\ln \epsilon_a - \ln \hat{a}}{\ln g} \cdot \frac{1}{\bar{\sigma}_s \cdot \Delta x}. \quad (10)$$

403 It is however usually a good idea to use  $m'$  at least equal to the  
 404 grid resolution along the propagation direction, to allow for light  
 405 even from the first row of cells to sufficiently penetrate into the  
 406 rest of the volume (cf. [15]). The subsequent propagation then  
 407 operates on the residual isotropic radiance in the merged grid,  
 408 iterating until the residual energy falls below  $\epsilon_L$ .

409 Reasonable values for  $\epsilon_a$  are around 0.1, or even higher. In fact  
 410 this decision is not unlike the one made in similarity theory [38].

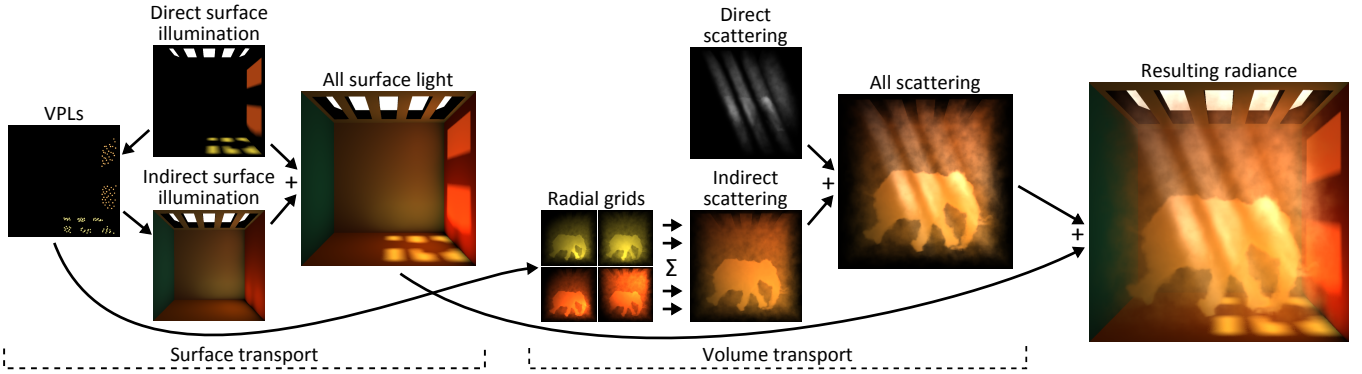


Figure 9: Workflow of the presented algorithm for a single directional light. For distant environment illumination the volumetric part of the pipeline is very similar, with the exception of rectilinear grids being used to propagate illumination from distant ordinates instead of the combination of VPLs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




Figure 10: 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.

436 The resulting grid is thus topologically equivalent to rectilinear  
 437 grid (except for being cyclic in the two angular dimensions) and  
 438 albeit not being uniform, it allows us to approximately treat the  
 439 space as locally Euclidean and obtain plausible results again  
 440 using virtually the same propagation scheme as before. The  
 441 main difference in the propagation is that we have to account  
 442 for the quadratic fall-off : although we base our propagation on  
 443 radiance, we have to explicitly compensate for the varying grid  
 444 cell sizes resulting from the non-uniform shell spacing. To this  
 445 end, we scale the radiance when propagating along the principal  
 446 direction in proportion to the radial coordinate spacing. A sam-  
 447 ple demonstration of this propagation type for a point light in a  
 448 simple homogeneous spherical medium is shown in Fig. 10.

411 Here, after a certain number of scattering events the propagation  
 412 switches to isotropic scattering, which is accompanied by a  
 413 switch to a so-called *reduced scattering coefficient*. This is  
 414 usually done on an empirical basis and despite the fact that using  
 415 the Henyey-Greenstein phase function allows us to quantify the  
 416 decision better (cf. [6]) this approach is still an approximation.  
 417 Similarity theory also does not apply well to heterogeneous  
 418 media. Thanks to the fact that we treat scattering as a gradual  
 419 decrease of anisotropy we can transit to isotropic propagation  
 420 in a well controlled manner, without changing the propagation  
 421 parameters or compromising the solution accuracy (aside from  
 422 small geometric misalignments caused by the grid merging). We  
 423 demonstrate this in Fig. 8.

## 424 6. Radial grids for local light sources

425 In order to extend our method to local light sources, we use  
 426 spherical grids with two angular coordinates and a radial coordi-  
 427 nate which is again aligned with the initial principal directions of  
 428 the point source (Fig. 2, bottom). To obtain more isotropic cell  
 429 shapes, the spacing of shells along the radial coordinate grows  
 430 exponentially (in proportion to the radial segment length at a  
 431 given radius). For parametrizing the spherical domain we use the  
 432 octahedron parametrization [30] mainly as it is simple, provides  
 433 reasonably uniform sampling, and above all, it discretizes the  
 434 domain into a 2D square where every cell has four natural neigh-  
 435 bours (plus two along the radial axis), similar to rectilinear grids.

449 *Instant radiosity*. Given the ability to use local point lights, we  
 450 can use instant radiosity [16] methods, which represent complex  
 451 illumination as a collection of point lights, to simulate surface-  
 452 to-volume light transport. Normally these VPLs are obtained  
 453 from random walks through the scene. In our interactive setting,  
 454 we generate VPLs using a reflective shadow map (RSM) [4]  
 455 for every primary light. We importance-sample these RSMs  
 456 according to surface albedo and (attenuated) irradiance, aiming  
 457 at keeping the total number of VPLs low. The reflected radiance  
 458 is then used to initialize the radial propagation grids. Prefiltering  
 459 can be done in the same way as for environment maps: VPLs  
 460 with a large importance have a high initial anisotropy and vice  
 461 versa. Similar to surface lighting, we can use clamping to reduce  
 462 any remaining singularities [8]. Fig. 9 depicts the pipeline of  
 463 the algorithm when propagating scattering from one directional  
 464 light and VPLs generated from its RSM.

## 465 7. Results and analysis

466 All results were computed on a PC with a 3.7 GHz Intel Xeon  
 467 CPU, 16 GB of RAM and an Nvidia GeForce GTX 770 GPU  
 468 with 2 GB of VRAM. Our implementation is written in C++,  
 469 using OpenGL and GLSL for the GPU code. In all our measure-  
 470 ments we use the framebuffer resolution of  $800 \times 600$  in order  
 471 to let the computation time be dominated by the propagation  
 472 rather than ray-marching. Resolutions of the medium density






Figure 11: Comparison of our Principal-Ordinates Propagation (POP) to SHDOM and a Monte-Carlo reference (light tracing), 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 POP 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 it fails to reproduce the high scattering anisotropy. Our method compares well to the reference solution, and even with real-time settings it qualitatively matches the overall appearance.

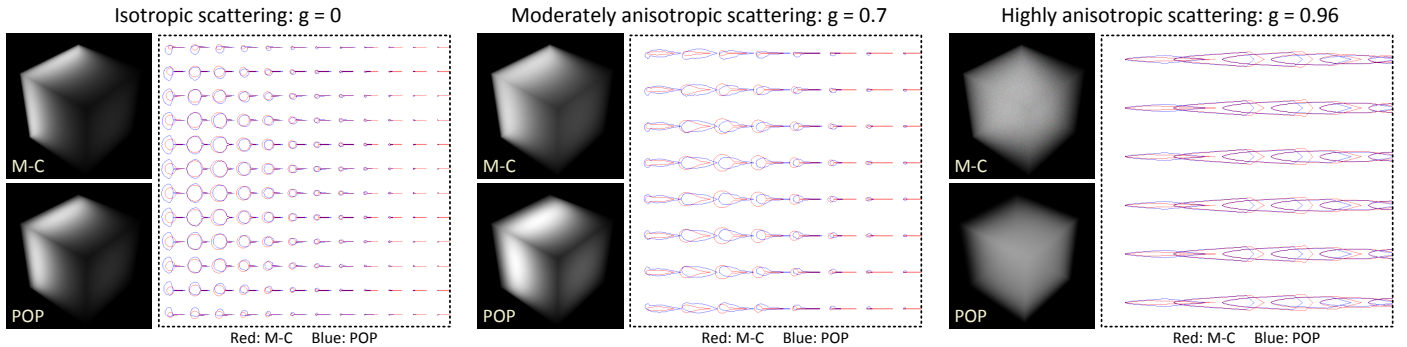


Figure 12: Detailed analysis of our method (POP) in comparison to Monte-Carlo reference (light tracing). We use a sample single-ordinate scenario, with a  $16^3$  propagation domain aligned with the medium (unit homogeneous cube,  $\sigma_s = 4 \text{ m}^{-1}$ , 100% albedo), using three representative anisotropy values. The plots compare the converged (incident) radiance distributions within a 2D horizontal slice in the middle of the domain.

473 datasets are typically in the order of tens in each dimension  
 474 (but effectively enhanced by procedural noise). Although the  
 475 number of propagation iterations needs to be chosen empirically  
 476 at the moment, in general we found that amounts similar to the  
 477 propagation grid resolution along the propagation dimension is  
 478 sufficient (around 10–20 in our examples). Please note that we  
 479 blur the environment maps only for presentation purposes (so  
 480 that the medium lighting features can be examined better) – the  
 481 actual illumination is in fact sampled from the full-resolution  
 482 maps. Other specific scene details are provided in the caption of  
 483 each discussed figure.

484 *Reference comparisons.* We first compare our approach to an  
 485 unbiased Monte-Carlo reference (light tracing), as well as spher-  
 486 ical harmonics (SH) DOM, in Fig. 11. It is apparent that the  
 487 described artifacts prevent SHDOM from handling anisotropic  
 488 media correctly, despite being theoretically capable to do so. In  
 489 contrast POP, despite being biased, reproduces the qualitative  
 490 appearance well.

491 A simpler analysis for a single directional propagation is pre-  
 492 sented as well. Fig. 12 shows comparisons to the reference for a  
 493 single ordinate, propagating in a simple cubic medium. Again,  
 494 despite some differences in the appearance, we can see the di-  
 495 rectional radiance distributions match well. We note that this  
 496 is actually a difficult case for our method, because the medium  
 497 regularity and the high density gradient on the faces parallel to  
 498 the propagation direction violate the alignment assumption, as  
 499 discussed in Sec. 4.5. However, observe that the directional dis-

500 tributions produced by POP have more complex shapes that the  
 501 simple Henyey-Greenstein ellipsoid lobes, since they are consti-  
 502 tuted by a superposition of such lobes. In addition, Fig. 10 shows  
 503 a simplified analysis similar to this, for a point light source.

504 *Propagation behaviour.* We examine the convergence of our  
 505 method in Fig. 13. The setting is identical to the second case  
 506 in Fig. 12. Notice that because of an absence of absorption the  
 507 propagation takes a significant number of iterations, even for the  
 508 small  $16^3$  grid. That is the main motivation for introducing the  
 509 isotropic residuum propagation (Sec. 5.4).

510 The effect of using different numbers of principal ordinates is  
 511 shown in Fig. 14. It can be seen that the discretization becomes  
 512 apparent only with very few ordinates. The importance propaga-  
 513 tion usually helps to alleviate this by sampling those directions  
 514 which will influence the solution most significantly. As Fig. 7  
 515 demonstrates, this is most likely the opposite side of the medium,  
 516 suggesting that a simpler empirical heuristic could potentially  
 517 work in certain cases.

518 One of the main shortcomings of the importance propagation is  
 519 its potential temporal incoherency, mostly manifested by tempo-  
 520 ral flickering. For this reason we filter the importance map both  
 521 spatially and temporarily, which, however, is not a fully robust  
 522 solution to the issue. One of our main targets for future work  
 523 is therefore improving this by explicitly enforcing temporal co-  
 524 herence when the sampled light sources relocate due to camera  
 525 movement or illumination changes, similarly to, e.g. [31].

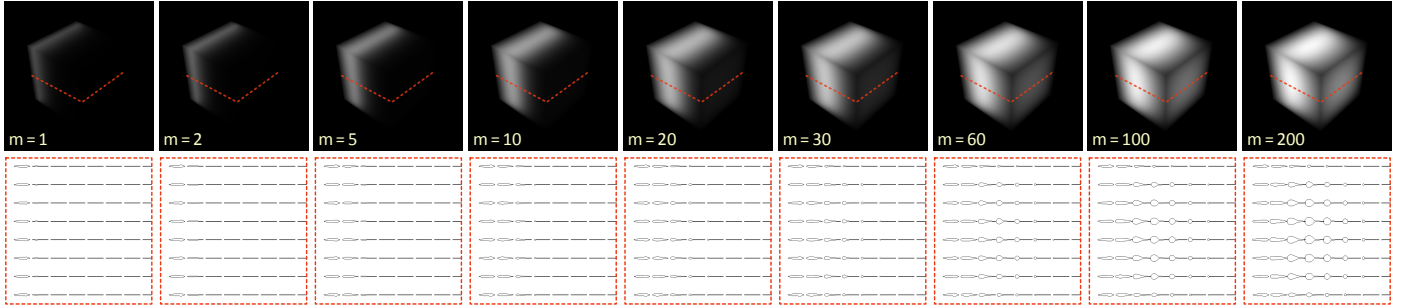


Figure 13: Convergence of our propagation scheme for the setting described in Fig. 12, with scattering anisotropy  $g = 0.7$ . The plots show the respective incident radiance distributions within a 2D horizontal slice in the middle of the domain (marked by the red dashed line). The observed strong forward peaks represent the unscattered energy which did not (yet) interact with the medium.

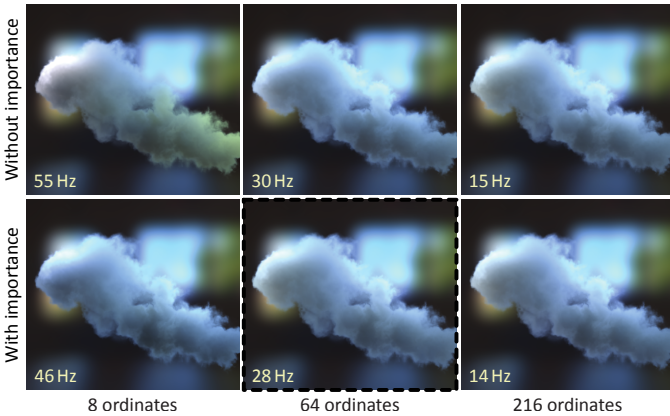


Figure 14: The smoke dataset with an increasing number of ordinates using the “kitchen” environment map ( $g = 0.9$ ,  $20^3$  grid resolution, 10 propagation iterations). Accounting for importance improves the results, mainly if low numbers of principal ordinates are used. The typical setting we use (shown in the bottom-centre) takes 5 ms for importance propagation, 2 ms for determining the ordinates, 2 ms for grid initialization, 12 ms for propagation, 2 ms for residuum propagation, 4 ms for grid merging and upsampling and 5 ms for ray-marching.

526 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. Our 530 prefiltering removes these artifacts but still allows perceiving 532 features of the background illumination, thanks to its adaptivity 533 (as opposed to a naïve prefiltering of the source illumination).

534 *Scattering anisotropy.* The shortcomings of current methods in 535 handling highly anisotropic scattering were the main motivation 536 for our work, as by far the majority of both natural and artificial 537 media exhibit anisotropic scattering (cf. [26]).

538 We tested our method for clouds with naturally very high scatter- 539 ing anisotropy in comparison to their isotropic versions (Fig. 15). 540 It can be seen that our propagation scheme handles both cases 541 well, and that correctly handling anisotropic scattering is a key 542 to reproducing such media. The same can be observed in Fig. 1, 543 since steam has properties similar to clouds. Interestingly, grid 544 resolutions as well as computation times required to render plau-

545 sible participating media are rather insensitive to its anisotropy, 546 i.e. anisotropic media render roughly as fast as isotropic media. 547 Although a larger number of ordinates might be required to re- 548 produce high-anisotropy effects, this additional effort is usually 549 compensated by a decreased complexity of the spatial radiance 550 distribution, which enables using coarser propagation grids.

551 Our dual propagation scheme also efficiently handles optically 552 thick anisotropic media, as seen in Fig. 8. The initial, full prop- 553 agation handles the directionally-dependent portion of radiant 554 energy, while the remaining isotropic residuum is rapidly propa- 555 gated in the second stage.

556 *Animation.* Thanks to the fully dynamic nature of our approach 557 we can seamlessly handle animated media without any precom- 558 putations or performance penalty. Fig. 16 shows several frames 559 of an animated smoke plume coherently rendered at real-time 560 framerates. In Fig. 17 we then demonstrate the dynamic interac- 561 tion of POP with surfaces, as described in Sec. 6.

562 In general, we believe to have demonstrated the versatility of 563 our method. Our propagation is capable of computing direct 564 illumination and low-order scattering effects (light shafts), as 565 well as arbitrary multiple scattering from directional, local and 566 environment illumination. Orthogonal to this, POP is capable of 567 simulating media with wide ranges of optical thickness, albedo 568 and most importantly, scattering anisotropy.

## 569 8. Discussion and conclusion

570 We propose a novel discrete ordinates method capable of comput- 571 ing light transport in heterogeneous participating media exhibit- 572 ing light scattering of virtually arbitrary anisotropy. The method 573 does not require any precomputations, which makes it well suit- 574 able for simulating dynamic and evolving media without extra 575 considerations. Our representation also adapts to and prefilters 576 the incident lighting. Radiance is represented by the Henyey- 577 Greenstein distribution, and propagated by our novel scheme in 578 volumes oriented along estimated principal light directions.

579 In general the steps of the proposed method are physically- 580 plausible (please refer to the supplementary materials for further




Figure 15: 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 35 Hz using 64 ordinates and  $20^3$  grid resolution for each of them, with 15 propagation iterations. The scattering anisotropy was set to  $g = 0.96$ .




Figure 16: Animated sequence of an expanding smoke plume, rendered dynamically at 30 Hz by our method while maintaining good temporal coherence.

581 details). The employed empirical heuristics introduce a certain  
 582 bias but allow us to make design decisions that result in a real-  
 583 time performance on contemporary graphics hardware.

584 The decomposition into a finite number of directions for distant  
 585 light can only be successful if the variation of the initial light  
 586 distribution is not too high; this however holds for the HDR  
 587 environment maps we used in our experiments (e.g. Fig. 15).  
 588 In addition our prefiltered initialization can be used to avoid  
 589 discretization artifacts in favour of a smooth approximation  
 590 (Fig. 6).

591 Since the presented method directly relates to DOM it shares  
 592 some of its basic limitations, such as handling of (surface) bound-  
 593 aries. In volumes with high density gradients (close to opaque  
 594 surfaces) the light distribution might not be faithfully reproduced  
 595 by the HG basis aligned with the initial light direction. Also  
 596 the resolution of every principal grid is limited and the general  
 597 limitations of discrete sampling apply: for finer details higher  
 598 resolutions are required. However, the upsampling (Sec. 4.4)  
 599 and prefiltering (Sec. 5.1) steps help to defer these issues and for  
 600 typical volume data sets moderate propagation grid resolutions  
 601 of  $8^3$ – $20^3$  have shown to be sufficient to handle a wide range of  
 602 illumination conditions and medium properties.

603 Another characteristic inherent to all finite-element methods is  
 604 that their convergence rate depends not only on the propagation  
 605 domain resolution but also on the optical thickness of the simu-  
 606 lated medium; especially for high-albedo media the number of  
 607 iterations required for producing a converged solution might be

608 prohibitively high. Our approach deals with this issue by using  
 609 multiple superimposed, relatively small propagation grids, in  
 610 which a low number of iterations is sufficient to propagate most  
 611 of the radiant energy (cf. Fig. 8). Media with higher optical  
 612 thickness also decrease the anisotropy of the propagated light  
 613 faster, allowing us to switch to the cheaper isotropic propaga-  
 614 tion mode earlier (Sec. 5.4). The lighting frequencies resulting  
 615 from the isotropic transport are by definition low and therefore a  
 616 lower-resolution propagation domain is sufficient here as well.

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

624 *Acknowledgements.* We would like to thank Anton Kaplanyan  
 625 for inspiring discussions, Chuong Nguyen, Peter Vangorp and  
 626 Yulia Gryaditskaya for reviewing early drafts of the article, and  
 627 the anonymous reviewers for their valuable feedback.

628 [1] M. Billeter, E. Sintorn, and U. Assarsson. Real-time multiple scattering  
 629 using light propagation volumes. In *Proc. 13D*, 2012.  
 630 [2] A. Bouthors, F. Neyret, N. Max, E. Bruneton, and C. Crassin. Interactive  
 631 multiple anisotropic scattering in clouds. In *Proc. 13D*, 2008.  
 632 [3] S. Chandrasekhar. *Radiative transfer*. Dover Publications, 1960.  
 633 [4] C. Dachsbacher and M. Stamminger. Reflective Shadow Maps. In *Proc.*  
 634 *13D*, pages 203–213, 2005.  
 635 [5] O. Elek, T. Ritschel, C. Dachsbacher, and H.-P. Seidel. Interactive light  
 636 scattering with principal-ordinate propagation. In *Proc. Graphics Interface*,  
 637 2014.  
 638 [6] O. Elek, T. Ritschel, A. Wilkie, and H.-P. Seidel. Interactive cloud render-  
 639 ing using temporally coherent photon mapping. *Computers & Graphics*,  
 640 36:1109–1118, 2012.  
 641 [7] T. Engelhardt and C. Dachsbacher. Epipolar sampling for shadows and  
 642 crepuscular rays in participating media with single scattering. In *Proc.*  
 643 *13D*, 2010.  
 644 [8] T. Engelhardt, J. Novák, T.-W. Schmidt, and C. Dachsbacher. Approximate  
 645 bias compensation for rendering scenes with heterogeneous participating  
 646 media. *Comp. Graph. Forum*, 31(7):2145–2154, 2012.




Figure 17: Scene with both animated medium and illumination, combining scattering from directional and local virtual light sources, running at 16 Hz (including the generation of the 125 VPLs and rendering the indirect illumination from surfaces, which alone takes 30 ms). The medium has a size of  $20^3$  m with  $\sigma_s = \{3.2, 3.3, 3.4\} \text{ m}^{-1}$ ,  $\sigma_a = \{1.15, 1.2, 1.3\} \text{ m}^{-1}$ ,  $g = 0.7$ , and is very heterogeneous. 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.

- 647 [9] R. Fattal. Participating media illumination using light propagation maps. *ACM Trans. Graph.*, 28:7:1–7:11, 2009.
- 648 [10] R. Geist, K. Rasche, J. Westall, and R. Schalkoff. Lattice-Boltzmann lighting. In *Proc. EGSR*, pages 355–362, 2004.
- 649 [11] L. G. Henyey and J. L. Greenstein. Diffuse radiation in the Galaxy. *Astrophysical Journal*, 93:70–83, 1941.
- 650 [12] W. Jarosz, C. Donner, M. Zwicker, and H. W. Jensen. Radiance caching for participating media. *ACM Trans. Graph.*, 27:7:1–7:11, 2008.
- 651 [13] 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.
- 652 [14] 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.
- 653 [15] A. Kaplanyan and C. Dachsbacher. Cascaded light propagation volumes for real-time indirect illumination. In *Proc. 13D*, 2010.
- 654 [16] A. Keller. Instant radiosity. In *Proc. SIGGRAPH*, pages 49–56, 1997.
- 655 [17] J. Kniss, S. Premoze, C. Hansen, and D. Ebert. Interactive translucent volume rendering and procedural modeling. In *Proc. Visualization*, pages 109–116, 2002.
- 656 [18] J. Kopf, M. F. Cohen, D. Lischinski, and M. Uyttendaele. Joint bilateral upsampling. *ACM Trans. Graph.*, 26(3):96, 2007.
- 657 [19] J. Krüger, K. Bürger, and R. Westermann. Interactive screen-space accurate photon tracing on GPUs. In *Proc. EGSR*, 2006.
- 658 [20] E. P. Lafortune and Y. D. Willems. Rendering participating media with bidirectional path tracing. In *Proc. EGWR*, pages 91–100, 1996.
- 659 [21] G. Liktov and C. Dachsbacher. Real-time volume caustics with adaptive beam tracing. In *Proc. 13D*, 2011.
- 660 [22] T. Lokovic and E. Veach. Deep shadow maps. In *Proc. SIGGRAPH*, pages 385–392, 2000.
- 661 [23] N. Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995.
- 662 [24] 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.
- 663 [25] A. Musbach, G. W. Meyer, F. Reitich, and S. H. Oh. Full wave modeling of light propagation and reflection. *Comp. Graph. Forum*, 32(6), 2013.
- 664 [26] 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.
- 665 [27] J. Novák, D. Nowrouzezahrai, C. Dachsbacher, and W. Jarosz. Progressive virtual beam lights. *Comp. Graph. Forum*, 31:1407–1413, 2012.
- 666 [28] 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.
- 667 [29] M. Pauly, T. Kollig, and A. Keller. Metropolis light transport for participating media. In *Proc. EGWR*, pages 11–22, 2000.
- 668 [30] E. Praun and H. Hoppe. Spherical parametrization and remeshing. *ACM Trans. Graph.*, 22(3):340–49, 2003.
- 669 [31] R. Prutkin, A. Kaplanyan, and C. Dachsbacher. Reflective shadow map clustering for real-time global illumination. In *Proc. Eurographics (Short Papers)*, pages 9–12, 2012.
- 670 [32] 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, Microsoft Research, 2008.
- 671 [33] 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.
- 672 [34] M. Salvi, K. Vidimče, A. Lauritzen, and A. Lefohn. Adaptive volumetric shadow maps. In *Proc. EGSR*, pages 1289–1296, June 2010.
- 673 [35] 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.
- 674 [36] J. Stam. Multiple scattering as a diffusion process. In *Proc. EGWR*, pages 41–50, 1995.
- 675 [37] 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.
- 676 [38] D. R. Wyman, M. S. Patterson, and B. C. Wilson. Similarity relations for the interaction parameters in radiation transport. *Applied Optics*, 28(24):5243–5249, 1989.
- 677 [39] 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.