# Participating Media Illumination using Light Propagation Maps

## Raanan Fattal<sup>1</sup> Hebrew University of Jerusalem, Israel

Light traveling through semi-transparent media such as smoke and marble is absorbed and scattered. To achieve proper realistic visualizations of such media, illumination algorithms must account for these events. In this paper, we present a new method for solving the *Radiative Transport Equation*, which models such evolution of light. The new method falls into the category of the *Discrete Ordinates Method* and inherits its generality and computational lightness. This method is known to suffer from two main shortcomings, namely the *false scattering* and the *ray effect*, which we avoid in our new method. By propagating the light using lower-dimensional maps of rays we detach their transport from the Eulerian grid and use fine angular discretizations. Thus, the scattering effect at each scattering generation is eliminated and the ray effect is significantly reduced at no additional memory requirements. Results demonstrate the new method's efficiency, ability to produce high-quality approximations, and its usefulness for a wide-range of computer graphics applications.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism— Radiosity

General Terms: Algorithms Additional Key Words and Phrases: global illumination, participating media, Discrete Ordinates method, Radiosity, Monte Carlo

#### 1. INTRODUCTION

Generating realistic rendering of natural objects is one of the main concerns in computer graphics. Over the last years there was a consistent trend to use physical models of increasingly higher degrees of complexity and accuracy to achieve this goal. Fluid flow, cloth dynamics, and fire simulations are just a few examples to received such an attention. In order to generate convincing rendering of scenes containing haze, clouds, dust, and translucent solids such as glass and marble, one must account for light absorption, scattering, and emission effects. These phenomena are formalized in an integro-differential equation known as the *Radiative Transport Equation* (RTE) which we present in Section 2. This equation contains, as its unknown, the light intensity at every point in space and along every possible direction, a total of five dimensions in the case of 3D physical space. This fact alone poses major computational challenges for those attempting to approximate the solution on discrete grids of limited resolution. Therefore, aside from the asymptotic complexity of the number of flops they perform, solvers are measured by their running times and the accuracy of their results. One critical factor is the amount of non-smoothness a method can sustain and produce at a given resolution.

As for the RTE, several different methods are suggested by both the thermal engineering and the computer graphics communities, each approach with its own advantages and disadvantages as we discuss in the following section. The *Discrete Ordinates Method* (DOM) is one family of such solvers that is popular among heat transfer engineers for its good trade-off between accuracy and computational cost while not posing any restrictions on the nature of the medium. Yet, it is known to suffer from two types of truncation error: (i) the false scattering, also known as numerical smearing, in which sharp beams are erroneously smoothed and (ii) the ray effect, whereby light that emanates from

bright sources in a set of fictitious directions.

In this paper we propose a new technique for solving the RTE which falls into the category of the DOM. In this method, the solution is constructed using maps of light rays that propagate the light across the domain. Since these maps have one spatial dimension less than the original domain and they contain rays with a restricted set of orientations, they can be discretized at high resolutions. Moreover, their evolution through space is done parametrically and is independent of the discretization used to store the solution. Thus, the new approach allows us to: (i) avoid the false scattering phenomenon at each scattering generation altogether, and (ii) practice fine angular resolutions when propagating the light and thereby significantly reducing the ray effects. This is achieved without compromising the admissive nature of DOM, *i.e.*, allowing unrestricted albedo range, accounting for multiple scattering, using arbitrary scattering distribution, and not posing any smoothness constraints over profile of the media coefficients.

The remainder of this paper is organized as follows. In the next section we review existing methods for solving the RTE. In Section 2 we review the RTE and establish notations. The new method is described in Section 3. In Section 4 we report the tests and comparisons evaluating its performance, and in the last section we draw some conclusions.

#### 1.1 Previous Work

Much effort was put into devising numerical solvers for the RTE by different communities. We will not attempt to list all these reports here but instead, we classify the techniques into groups and mention their respective representatives. The reader is referred to [Cerezo et al. 2005] and [Perez et al. 1997] for a more elaborate survey of these approaches, and to [Sigel and Howell 1992] for an additional summary of the topic.

The Zonal Method [Rushmeier 1988; Hottel and Sarofim 1967] extends the treatment surfaces receive in Radiosity to volume elements discretizing the 3D medium. This requires computing form-factors between every pair of voxels (volume elements) and involves a computational complexity of the order  $O(n^7)$ , where *n* is the number of discrete variables taken along each axis in 3D. In [Sillion 1995] a hierarchical strategy is employed to reduce this complexity under certain assumptions about the structure of the scene. A method for accelerating the computation of the form-factors is proposed in [Arquès and Michelin 1996] and relies on spatial coherence. This approach is extended for non-isotropic scattering in [Bhate and Tokuta 1992] where *m* spherical harmonics coefficients are used to represent the angular light distribution. In this approach the number of form-factors grows, and the order of computational complexity increases to  $O(n^7m^2)$ .

In [Stam 1995] multiple scattering is approximated as a diffusion process to achieve a low-cost solution for scenes with an optically thick medium. In [Premoze et al. 2004] the spread of direct illumination is estimated by a model relating scattering and absorption to the distance light travels from particular light sources.

In [Behrens and Ratering 1998] shadows due to the occlusion of direct light are computed in a texture-based volume rendering. The volume is stored as parallel polygons which carry the volume data as texture maps. These polygons are shifted according to the light direction and thus interpolations between successive layers are avoided. We use a similar mechanism for both the direct and indirect components of the illumination in the context of a more sophisticated illumination model that accounts for emission and scattering on top of absorption. In [Dobashi et al. 2000] a single scattering effect is computed for clouds modeled by metaballs.

Given the high dimensionality of the problem [Pattanaik and Mudur 1993] employ Monte Carlo based stochastic sampling methods, [Lafortune and Willems 1996] use bidirectional tracing, and [Pauly et al. 2000] employ Metropolis sampling. Despite the various variance reduction strategies used, these approaches require a considerable number of samples to be drawn in order to remove the noise in the resulting image. In [Jensen and Christensen 1998], the amount of samples needed to obtain cleaner images is reduced by approximating the direct component of illumination

ACM Transactions on Graphics, Vol. V, No. N, Month 20YY.

2

•



Fig. 1. On the left, a 2D Discrete Ordinates grid is illustrated. In each square element several variables discretizing light's orientation are held. On the right, a light propagation map is shown to sweep a 3D grid. The discrete directions held in the LPM are finer and restricted to one sixth of the unit sphere.

deterministically, by sampling light sources using ray tracing. In general, stochastic photon tracing approaches require the representation of individual spatial coordinates and their evolution. Therefore, they perform more computations per *single* sample than what is required by methods based on stationary variables discretizing space. Restricted models, used to model light transport at a sub-surface layer, are described in [Hanrahan and Krueger 1993] and in [Jensen et al. 2001] for the case of highly scattering media. These methods use a Monte Carlo integration to compute their solution and are very efficient for their purpose.

In the Discrete Ordinates Method (DOM) [Chandrasekhar 1950], as depicted in Figure 1, the quantities appearing in the RTE are discretized in a straightforward fashion both in space and orientation, typically in the original coordinates systems. These discrete values are used to approximate the different terms in the RTE, which become a finite dimensional system of equations expressing the light exchange between neighboring volume elements. In its classic form, the angular discretization is taken *pointwise*, allowing light to travel only at a finite number of directions (the method owes its name to this property). The DOM is the method of choice for radiative transfer engineers and established itself as the most cost effective approach [Coelho 2002; Languenou et al. 1994]. The main reason behind this is that the calculations are based only on *local* interactions and avoid the need to compute form-factors between every pair of elements, thus reducing the complexity of the problem considerably. In fact, this approach corresponds to approximations made by the Finite Element Method [Zienkiewicz and Taylor 2000] and the Finite Volume Method [Leveque 2002], which dominate many computational regimes, such as fluid and solid mechanics as well as hyperbolic and parabolic conservation laws.

Unfortunately this discretization is known [Perez et al. 1997; Coelho 2002; 2004] to suffer from two major truncation errors, to the extent that they can become visually disturbing. Repeated averaging arising from the interpolation, needed to approximate the flux between volume elements, causes light to smear in space. Effectively this does not allow light beams to maintain a sharp profile. Note that this occurs at each light orientation separately and therefore the light direction remains unaltered, and the term scattering is somewhat misleading. As shown in [Coelho 2002], this effect is the counterpart of the artificial viscosity arising in computational fluid mechanics, where the momentum is falsely-diffused in space. Discretizing the directions in which light can travel causes another type of error, the ray effect. This error consist of light that emanates from relatively small and intense regions and reveals this angular discretization as spurious beams of light. This is more pronounced in media that allow the light to propagate without being immediately absorbed or scattered.

In order to reduce this effect [Ramankutty and Crosbie 1997] propose the *Modified DOM*, where the ray intensity is broken into direct and diffuse components. The direct light component is determined analytically under the assumption of a homogeneous isotropic medium with a constant thickness which cannot produce the ray effect. An extension of this method that allows discontinuous medium emission in the presence of non-homogeneous non-isotropic scattering



Fig. 2. The two images on the left (color coded differently) show the result produced by a first-order DOM simulation with 54 discrete direction in each of the  $128^3$  cells. In this example a sharp beam of light enters from the left, penetrates the vase, and is scattered by its walls. The region where the light initially hits the vase indicates the beam's profile got already smeared before reaching the vase. The ray effect is observed when the light scattered from this regions propagates only in certain directions, reaching specific regions such as near the opening and the right side of the vase. A more accurate solution is shown in the two images on the right, which were obtained using LPMs. This involves a much lower angular discretization consisting of one variable per cell in the output grid (enough for isotropic scattering) and six variables to store the unpropagated light during the construction of the solution. The angular discretization used for the LPMs, which are 2D maps, is  $9 \times 9$  different ray orientations and their spatial resolution is  $128^2$  (equals to the output grid resolution).

is described in [Coelho 2004], at the cost of computing additional spatial integrals. In [Kajiya and Herzen 1984] under the assumption of high albedo medium, a truncated spherical harmonics expansion is used to represent angular light distribution for rendering clouds. Another way of propagating light through whole bins is reported in [Max 1994]. In this method inaccuracies are introduced because light is attenuated not only along a straight path but along many other possible routes. In [Stam 2001] the DOM is used to render human skin layer by modeling its boundary as rough surfaces.

A family of *High-Resolution* techniques were developed for approximating the hyperbolic convective terms [Leveque 2002] in computational fluid mechanics. This reduces the numerical viscosity by maintaining a high-order of accuracy and adaptively switching to lower accuracy stencils in order to withstand discontinuities. The use of these techniques to solve the RTE is proposed in [Jessee and Fiveland 1997; Coelho 2002] and [Liu and Pollard 1996]. As we shall show later, this action reduces the numerical smearing but only up to a limited extent. The low-order stencils used to avoid non-physical oscillations, near sharp transitions in the solution, are diffusive and hence cause a noticeable initial smoothing to sharp beams. Also, although the RTE is linear, these operators yield a system of *non-linear* equations, which are harder to solve. These high-resolution schemes do not offer any reduction in the ray effect.

## 2. GOVERNING EQUATIONS

The radiative transport equation models the propagation of radiation through a participating medium which absorbs, scatters, and emits light. For a particular wavelength the RTE is given by

$$(\boldsymbol{\omega} \cdot \nabla) I(\boldsymbol{x}, \boldsymbol{\omega}) = e(\boldsymbol{x}, \boldsymbol{\omega}) - (k(\boldsymbol{x}, \boldsymbol{\omega}) + \boldsymbol{\sigma}(\boldsymbol{x})) I(\boldsymbol{x}, \boldsymbol{\omega}) + \frac{\boldsymbol{\sigma}(\boldsymbol{x})}{4\pi} \int_{\Omega} I(\boldsymbol{x}, \boldsymbol{\omega}') f(\boldsymbol{\omega}, \boldsymbol{\omega}') d\boldsymbol{\omega}',$$
<sup>(1)</sup>

where  $I(x, \omega)$  is the radiation intensity  $(W/m^2 \cdot sr)$ , at a point *x* that propagates in the direction of  $\omega$ . The scalar fields *e* and *k* are the medium self-emission and absorption coefficients, respectively, and describe the production and loss of radiation along  $\omega$  through *x*. The scattering coefficient of the medium  $\sigma$  appears both in the out-scattering term  $\sigma(x)I(x,\omega)$ , where the intensity along  $\omega$  is reduced, and in the in-scattering (last) term, where radiation, arriving from all incident directions  $\omega' \in \Omega$  and scattered to  $\omega$ , is integrated. The *phase function*  $f(\omega, \omega')$  expresses the fraction of

ACM Transactions on Graphics, Vol. V, No. N, Month 20YY.

scattered radiation along  $\omega$  arriving from  $\omega'$ , relative to the isotropic scattering fraction. Predicting light refraction or Fresnel reflection within the context of the RTE requires singular coefficients and phase function. In our scope we do not account for these cases of interfacing media.

#### 3. NEW METHOD

Discretizing the RTE either using the DOM or the Zonal method results in a system of linear equations which is then typically solved via iterative sweeps such as those of Jacobi, Gauss-Seidel, or Southwell. As shown in [Gortler et al. 1994] these iterations have different physical interpretations in terms of light propagation. In [Languenou et al. 1994] a similar observation is made in the context of the DOM, where light rays traveling close to one another induce a sweep direction when using these iterative solvers to compute a solution. This is exploited by [Jessee and Fiveland 1997] where the iterations are ordered in an "upstream to downstream," fashion. Here we take a more explicit advantage of these interpretations. In the course of solution construction, we iteratively propagate light across the domain by solving (1) explicitly along light rays that are close to one another in their direction of propagation. This is done using Light Propagation Maps (LPMs), which are temporary two-dimensional maps of rays containing only a fraction (one sixth in our implementation) of all possible directions. Hence, in contrast to the standard DOM and what is traditionally done in Finite Volume methods, we do not approximate the light flux between adjacent cells based on the discrete variables storing the solution. Instead, the total light reaching a cell is estimated from the LPMs, and the stationary grid variables serve merely for bookkeeping. By that we achieve two goals: (i) since the light stored in the LPMs is unattached to the coarse grid, it is propagated parametrically and independently of its orientation with respect to the axes of the 3D solution grid. Thus, we bypass the need to estimate the spatial derivatives in (1) and avoid interpolations. As a result a major source of numerical smearing, arising from averaging due to the interpolations, is totally eliminated. And (ii) the use of such lower-dimension maps with a reduced set of orientations allows us to practice fine discretizations of space, and more importantly, orientation. As we will show, this allows us to reduce the ray effect significantly at minimal memory costs, since only the LPM resolution is refined.

The numerical method presented in this section can be applied both in two- and three-dimensional space. The scheme is derived for the 3D case and can be straightforwardly interpreted in 2D. Also, for the sake of clarity, we assume the spatial discretization to consist of rectangular cells. In Section 3.2 we simplify the domain even more by taking it to be the unit cube so that formulas take on a simple and clear form. This method can be implemented for arbitrary domains with arbitrary control volume shapes and arrangements, under mild requirements of coherence between the different discretizations involved (described in Section 3.3).

## 3.1 Spatial and Angular Discretizations

In the new approach we discretize both the physical space and the sphere of orientations in the finite volume fashion [Leveque 2002]. We break up the unit sphere  $S^2$  into a set of non-overlapping bins  $\Omega^m$ , such that  $\cup \Omega^m = S^2$ . We also break the spatial domain  $D \subset \mathbb{R}^3$  into a grid of non-overlapping rectangular cells  $C_{i,j,k}$  of lengths  $\Delta x, \Delta y, \Delta z$  along the principle axes, such that  $\cup C_{i,j,k} = D$ . Note that the index *m* conveys a mapping to a discrete 2D manifold. Similarly to [Languenou et al. 1994], our goal here is to estimate the average scattered radiation intensity at each control volume in the complete cross-product space  $\Omega \times D$ , namely

$$I_{i,j,k}^{m} \approx (V_{i,j,k}^{m})^{-1} \int_{C_{i,j,k}} \int_{\Omega^{m}} \frac{\sigma(x)}{4\pi} \int_{S^{2}} I(x,\omega') f(\omega,\omega') d\omega' dx d\omega, \qquad (2)$$

where  $V_{i,j,k}^m = \Delta x \Delta y \Delta z |\Omega^m|$ , the total volume of the m, i, j, k-cell. This choice of a weak-form numerical representation of cell averages is known to be robust to discontinuities [Leveque 2002], such as the ones we expect to encounter here due to occlusions and non-smooth media coefficients. The emission, absorption, and scattering scalar fields are

assumed to be given either as constants in each cell or at a higher order of representation to allow a higher order formal accuracy as discussed below. Throughout this paper we refer the m, i, j, k grid, used to store the input medium coefficients and the solution itself, as the *coarse-grid* for it holds a coarser spatial and angular discretization than the one used in the LPMs.

#### 3.2 Light Propagation Maps

Each LPM consists of a 2D array of light rays passing through the domain D along *similar* directions, and hence similar calculations can be performed collectively. In our implementation we divide  $S^2$  into six subsets induced by the cartesian coordinates,

$$\begin{split} \Omega^{X-} &= \{ \omega \in S^2 : \omega_x < 0, |\omega_x|/\Delta x > |\omega_y|/\Delta y, |\omega_z|/\Delta z \}, \\ \Omega^{X+} &= \{ \omega \in S^2 : \omega_x > 0, |\omega_x|/\Delta x > |\omega_y|/\Delta y, |\omega_z|/\Delta z \}, \end{split}$$

and the analogically defined  $\Omega^{Y-}$ ,  $\Omega^{Y+}$ ,  $\Omega^{Z-}$ , and  $\Omega^{Z+}$ . To simplify the description we will assume that the domain *D* is the unit cube  $[0, 1]^3$  and derive the equations only for the LPM that correspond to the positive *Z* axis. The equivalent treatment along the other principle directions is easily inferred. The LPM that corresponds to this direction consists of the set of rays defined by

$$R_{r,s}^{n}(z) = (x_{r,s} + z \cdot \omega_{x}^{n} / \omega_{z}^{n}, y_{r,s} + z \cdot \omega_{y}^{n} / \omega_{z}^{n}, z),$$

where  $\omega^n$  samples  $\Omega^{Z+}$  and  $x_{r,s}, y_{r,s}$  sample  $[0,1]^2$ . In our implementation we use a uniform sampling for both space and orientation, meaning that  $\omega_x^n/\omega_z^n$  and  $\omega_y^n/\omega_z^n$  are given by a uniform sampling of  $[-1,1]^2$ , *i.e.*, a uniform sampling of the cube faces. Here again the 2D angular discretization is indexed by a single number *n*. We refer to the *r*,*s*,*n* discretization of the LPM as the LPM discretization or the fine-grid discretization in contrast to the *m*,*i*,*j*,*k* grid used for the final solution (and the input media coefficients) which we call the *coarse-grid*.

We use the LPMs to propagate light throughout the domain by sweeping along the consequent set of directions  $(\Omega^{\pm X}, \Omega^{\pm Y}, \Omega^{\pm Z})$  iteratively. The evolution of the *n*, *r*, *s*-ray intensity is derived from (1) by plugging in the parameterized ray position  $R_{r,s}^n(z)$  and intensity  $L_{r,s}^n(z)$  to get

$$\begin{split} \omega_z^n \frac{d}{dz} L_{r,s}^n(z) &= (\omega^n \cdot \nabla) I(R_{r,s}^n(z), \omega^n) = e(R_{r,s}^n(z), \omega^n) \\ &- \left( k(R_{r,s}^n(z), \omega^n) + \sigma(R_{r,s}^n(z)) \right) L_{r,s}^n(z) \\ &+ \frac{\sigma(R_{r,s}^n(z))}{4\pi} \int_{\Omega} I(R_{r,s}^n(z), \omega') f(\omega^n, \omega') d\omega' \,. \end{split}$$

We cannot express the unknown in-scattering term by the current LPM, which contains only a partial set of all ordinates, and use an axillary variable instead. We introduce a new scalar field *u* to account for the *unpropagated* intensity arising from the emission and scattering of previous sweeps. Given this quantity, the equation above becomes

$$\omega_{z}^{n} \frac{d}{dz} L_{r,s}^{n}(z) = -\left(k(R_{r,s}^{n}(z), \omega^{n}) + \sigma(R_{r,s}^{n}(z))\right) L_{r,s}^{n}(z) + u(R_{r,s}^{n}(z), \omega^{n}).$$
(3)

This derivation is quite similar to the one used in the progressive radiosity [Cohen et al. 1988] where the final intensities are calculated by accumulating the exchanged radiation at each step. In Section 5 we discuss the differences between the two approaches.

Equation (3) is integrated numerically by incrementing  $R_{r,s}^n$  in finite steps matching the cell lengths along that direction (e.g.  $\Delta z$  in the case of the Z+ sweep). Such an increment will advance the ray by less than  $\Delta x, \Delta y$  along the X and Y axes respectively. As shown in Figure 3, this allows us to expect intersections with no more than two cell faces before

ACM Transactions on Graphics, Vol. V, No. N, Month 20YY.



Fig. 3. Ray segments between cells.

hitting the face perpendicular to the Z-axis. The Z-coordinates of the corresponding intersections between the ray and cells faces are denoted by  $z_q$  for q = 0...3, provided all four exist. The parameterized lengths of the resulting segments,  $z_q - z_{q-1}$ , are denoted by  $\Delta l_q$  for q = 1, 2, 3. By treating the coupling between rays due to scattering explicitly, *i.e.*, ignoring the scattered light of the current sweep when evaluating u, equation (3) admits the following exact integration formula along the segments

$$L_{r,s}^{n}(z_{q}) = L_{r,s}^{n}(z_{q-1})e^{-\Delta l_{q}(K_{i,j,k}^{m}+S_{i,j,k})/\omega_{z}^{n}} + U_{i,j,k}^{m}\left(1 - e^{-\Delta l_{q}(K_{i,j,k}^{m}+S_{i,j,k})/\omega_{z}^{n}}\right)/(K_{i,j,k}^{m}+S_{i,j,k}),$$

$$\tag{4}$$

for  $2 \le q \le 2,33$ , depending on the number of intersections. The indices i, j, k correspond to the cell containing the ray fragment  $[R_{r,s}^n(z_q), R_{r,s}^n(z_{q-1})]$ . The numbers  $K_{i,j,k}^m, S_{i,j,k}$ , and  $U_{i,j,k}^m$  approximate  $k, \sigma$ , and u respectively and corresponds to a zero-order interpolation. The ordinate index m is taken such that  $\omega^n \in \Omega^m$ , which again corresponds to a constant interpolation. We further discuss the meaning and role of  $U_{i,j,k}^m$  in the next section. We avoid the numerical restriction  $\Delta l_q (K_{i,j,k}^m + S_{i,j,k})/\omega_z^n < 1$  by using the exponentials in formula (4) which keeps positive intensity values independently of the extinction rates. Higher order interpolations combined with a high-order integration scheme along the rays can be used to achieve a higher order of formal accuracy when approximating (4). Also, if  $k, \sigma$ , and e are given more accurately, then they can be evaluated on demand at the LPMs fine resolution as these evolve. We exploit the fact that many rays in the LPM have the same direction and make our calculation faster by choosing the LPMs' fine-resolution to be equal the coarse grid's resolution times some integer factor. This choice yield the same cell-ray intersections and ray fragment lengths such that the same integration weights can be used in all cells.

The variables  $I_{i,j,k}^m$  are set to zero and  $U_{i,j,k}^m$  to  $E_{i,j,k}^m$ , which is the discrete approximation to the cell emission. Light penetrating the domain from its faces, *i.e.*, the boundary conditions associated with the RTE, is translated in this formulation to the initial conditions of  $L_{r,s}^n$ . This process must not be regarded simply as Lagrangian advection of light, since all light rays are coupled by the scattering, which dictates a certain synchronization when storing and evolving light between successive sweeps, as we describe shortly.

#### 3.3 Iterative Radiance Estimation using LPMs

The iterated sweeps of the LPMs provide means for estimating the amount of light traveling and being scattered. This is needed to construct the solution  $I_{i,j,k}^m$ , as defined by (2), by accumulating the fluxes across each of the control volumes. In each of these sweeps some fraction of light is scattered in all directions, including ones which are not represented in the current LPM sweep. In fact, this is the source of the coupling inherent in the RTE, which is in turn responsible for the need for iterative solvers. The fraction of the scattered light flux which cannot be propagated at the current sweep is temporarily stored in the auxiliary field, u, until an adequate LPM sweep arrives. This light is stored across the entire domain at the coarse-resolution, used for storing the solution itself, by averaging it inside the m, i, j, k-cell and denote it by  $U_{i,j,k}^m$ . These variables are updated using the LPM sweeps according the cell-averaged

in-scattering term in (2) which we discretize as following. We replace the light  $I(x, \omega')$ , arriving from directions that are represented via the current LPM, i.e.,  $\omega' \in \Omega^{+Z}$ , by a zero-order approximation based on  $L_{r,s}^n(z)$ . This means that the light within the LPM's control volumes, which are surface elements partitioning  $[0, 1]^2$  and  $\Omega^{Z+}$ , is constant and equals to  $L_{r,s}^n(z)$ . Thus, the amount of light being scattered from a single LPM bin along a single ray-cell intersection segment is given by the following quadrature

$$\frac{S_{i,j,k}}{V_{i,j,k}^{m}} \sum_{n,r,s} A_{r,s} F^{n,m} \int_{z=z_{q-1}}^{z=z_{q}} L_{r,s}^{n}(z) / \omega_{z}^{n} dz$$

where  $\sigma$  is replaced by a zero-order approximation  $S_{i,j,k}$ , the  $A_{r,s}$ s are the areas of the spatial LPM volumes, and  $[z_{q-1}, z_q]$  is the third coordinate of the segment intersecting with the coarse-resolution i, j, k-cell. The light rays  $L_{r,s}^n(z)$  may show a rapid exponential decay, which is typical to highly-scattering media, and thus evaluations of its integral along  $[z_{q-1}, z_q]$  may lead to a discrepancy between this added quantity and the actual light loss in the LPM due to the scattering, as computed by (4). We avoid this by approximating this integral directly from the amount of light scattered in (4) and use

$$(V_{i,j,k}^m)^{-1} \sum_{n,r,s} A_{r,s} F^{n,m} L_{r,s}^n(z_{q-1}) (1 - e^{-\Delta l_q S_{i,j,k}/\omega_z^n}),$$
(5)

where the expected integration weight  $\Delta l_q S_{i,j,k} / \omega_z^n$  is replaced by  $1 - e^{-\Delta l_q S_{i,j,k} / \omega_z^n}$ . The discrete form-factors are given by

$$F^{n,m} = \frac{1}{4\pi} \int_{\Omega^m} \int_{\Omega^n} f(\omega, \omega') d\omega' d\omega,$$

and are computed in advance. Note that the angular areas do not appear in (5) explicitly, but they are factored by the form-factors. As we mentioned earlier, the discretization we use for the LPMs equals to the one used for the course-grid times some integer factor and hence  $A_{r,s}$  equals  $\Delta x \Delta y$  divided by that factor squared.

The angular resolution of the scattered light is reduced when it is transferred from the LPM into  $U_{i,j,k}^m$ . Yet, since the scattering kernel f acts as a low-pass filter in space of directions, we expect this scattered light to be smoother than the distribution of the light that produced it. Hence, in cases of strongly diffusing phase function this will introduce moderate inaccuracies, whereas a temporary storage of a higher resolution will be required for more concentrated scattering functions.

The only discretization restriction posed by this scheme is that all ray directions belonging to a certain LPM must be fully contained in some subset of the angular coarse-grid bins  $\Omega^m$ , and these bins must not contain any other ray direction belonging to a different LPM. This relation must also hold between each of  $\Omega^{\pm X}$ ,  $\Omega^{\pm Y}$ ,  $\Omega^{\pm Z}$  and the coarsegrid angular bins. We use the following simple construction which obeys this rule. We start with some partitioning of the sphere of all orientations, e.g.,  $\Omega^{\pm X}$ ,  $\Omega^{\pm Y}$ ,  $\Omega^{\pm Z}$ , which is dictated by the major directions along which we intend to scan the domain. Then, we take an integer division of each of these bins to define the coarse-grid angular bins, indexed by *m*, which depends on the phase function. For isotropic scattering a single bin is enough, whereas more bins are needed to represent non-isotropy. This is followed by another integer refinement, applied to each of the coarselevel angular (and spatial) bins that defines the LPM bins, which we index by *n*. This refinement is aimed to achieve accuracy and reduce the ray effect. This rule is needed in order to preserve the validity of the unpropagated energy values between successive sweeps.



Fig. 4. The left four images show a beam of light entering the domain from the left, passing through a medium of constant isotropic scattering. The ray effect is clearly evident in the upper two images, where we use 3x3 directions in each LPM. This artifact is considerably reduced when increasing the angular distribution of the LPMs to 9x9 as shown in the lower two images. On the right we show a comparison between different light propagation schemes. The images show the intensity of a light beam passing through a domain that does not absorb, emit, or scatter. Ideally, in this case we expect the beam's profile to remain unchanged. The Top-left image shows the beam generated by a first-order upwind and to its right we show the output of a second-order Lax-Wendroff advection scheme. At the bottom-left image we see the result of using a flux-limiter and to the right we see the result by propagating the light using LPMs.

#### 3.4 Algorithm Summary

As we carry light across the domain using the LPMs, the solution  $I_{i,j,k}^m$  is constructed by accumulating the average intensity scattered in the m, i, j, k-cell. We summarizes the different steps of this algorithm in the following pseudo-code lines:

```
Initialize U_{i,j,k}^m and I_{i,j,k}^m with medium emitted light
repeat until \max_{m,i,j,k} |U_{i,j,k}^m| < \varepsilon
for each map \Omega^d, d \in \{\pm X, \pm Y, \pm Z\}
if first sweep along d, then set the rays' initial conditions
set \tilde{U}_{i',j'}^m = U_{i',j'}^m for each coarse-grid direction contained in map \Omega^d
for each ray in map \Omega^d
for each ray segment
update L_{r,s}^n according eq. (4) given \tilde{U}_{i',j'}^m
add eq. (5) to I_{i,j,k}^m and U_{i',j'}^m for all m
end
end
advance along the domain in the direction d
end
```

The temporary layer of coarse-grid unpropagated light  $\tilde{U}_{i',j'}^m$  is simply used to evaluate the scattered light, integrated in (4), while allowing newly scattered light from the current sweep to be safely added to  $U_{i,j,k}^m$ . As evident from these lines, the LPMs belonging to each direction  $\pm X, \pm Y, \pm Z$  are used one at a time, and can thus be temporarily stored.



Fig. 5. These images show a medium with binary  $\sigma$ , the scattering region consists of two strips at the low and the upper edges plus two circles close to the upper strip. A beam of light enters from the left side of the domain and propagates towards the lower strip, hitting it at the middle. The light scattered from that region serves as an indirect light source illuminating the upper part of the domain. The two thick circles cast shadows on the upper strip. This is generated using LPMs with 9 × 9 bins in their angular discretization.)

As the maps advance, they will cease to cover the entire cross section of the domain, unless they are sufficiently large. For this reason, the rays leaving the domain should be abandoned while new ones must be introduced on the other end. In our implementation we solve this issue using cyclic mappings; the rays exiting the domain are used as the ones entering it by properly initializing them. Also, note that since the unpropagated light cells are overwritten, it means that all the light they contained must have been carried by the current LPM sweep. The validity of this exchange is ensured by the condition we described earlier, requiring that all rays corresponding to the same angular bin in  $U_{i,j,k}^m$  belong to the same LPM.

Once the solution  $I_{i,j,k}^m$  is calculated, similarly to [Languenou et al. 1994], it is used to evaluate the emitted and outscattered light components when rendering the scene using ray casting. More formally, along each visual ray r(t) we approximate the following integral

$$\int \tau(t) \Big( \frac{\sigma(r(t))}{4\pi} \int_{S^2} I(r(t), \omega') f(\hat{r}, \omega') d\omega' + e(r(t), \hat{r}) \Big) dt,$$

where  $\hat{r} = \frac{dr}{dt}$  (assuming  $\|\frac{dr}{dt}\| = 1$ ) and the transmittance function is defined by  $\tau(t) = \exp(-\int^t (k(s) + \sigma(s))ds)$ . Note that the function  $I(r(t), \omega')$  is precisely what  $I^m_{i,j,k}$  approximates. We use a trilinear interpolation to when integrating along the ray casts at equidistant sampling. We do not use any no aliasing preventative steps such as *jittering* [Pauly et al. 2000], which can be further incorporated.

## 4. RESULTS

We implemented a three-dimensional version of our method in C++ and run it on a 2.7GHz Pentium IV machine. Here we report three kinds of testing: one in which we construct artificial scenarios in order to isolate and evaluate different aspects of the method's performance. The second is a comparison of this method with traditional Discrete Ordinates and Monte Carlo techniques. In the third kind of test, we validate the new approach's applicability for a wide range of typical computer graphics scenarios involving different albedo ranges. Unless stated otherwise, we use the same spatial resolution for the LPMs and coarse-grid. The angular resolution of the coarse-grid is a single variable per cell, representing isotropic scattering, and six for the unpropagated light  $U_{i,j,k}^m$ . The LPM angular resolution, which has a critical role in suppressing the the ray effect, differers from one experiment to the next and is reported in each case. We perform three LPM sweeps in each direction to generate the results and account by that for three generations of light scattering.

ACM Transactions on Graphics, Vol. V, No. N, Month 20YY.

10

#### 4.1 Performance Evaluation

We begin by evaluating how the LPM angular resolution effects the resulting solution. In Figure 4, a beam of light enters the domain from the left encountering a constant  $\sigma$  and e, k = 0. The coarse-grid resolution is 64<sup>3</sup> with 6 angular variables for the unpropagated light and a single one for the solution itself, allowing us to simulate isotropic scattering. The top-left two images, color coded differently, show the ray effect resulting from a  $3 \times 3$  LPM angular discretization. Spurious light streaks are apparent and correspond to this coarse angular discretization. For this extremely concentrated intensity, increasing the LPM discretization to 9 × 9 angular variables makes this artifact disappear. Resolving the ray effect by refining the LPM angular discretization *alone* will pose critical memory requirements for these kind of calculations in 3D, as we quantify later. Also, the beam itself, all along its trajectory, remains as sharp as it was when entering the domain, showing a total absence of numerical smearing. The smearing effect is further investigated in Figure 4, where four different methods of propagating light are compared. The first-order upwind shows a significant amount of smearing while the second-order (Lax-Wendroff) generates severe oscillations and introduces negative intensities. Adding a flux limiter (minmod, see [Leveque 2002]) to the second-order scheme manages to avoid these oscillations, yet still shows a considerable amount of numerical smearing to this discontinuous beam profile. When propagating light using the LPMs, no smearing is generated and the light's cross section is perfectly preserved. In Figure 5, we demonstrate the importance of propagating indirect scattered light in *straight* lines and at high precision. The light reaching the upper edges shows a relatively sharp profile due to the shadows casted by the two scattering circles. These variations in indirect illumination show that there are cases where it can not be properly represented using coarse volume photon maps. Finally we evaluate how the proposed solution copes with strong forward-scattering. This is tested on a constant Heyney-Greenstein scattering medium with g = 0.9. In this case there is a concern that the scheme will over scatter the light and will tend toward isotropy, as the unpropagated light is temporarily stored in coarse angular bins. This is investigated by refining both the coarse-grid and LPM angular discretizations independently. The results shown in Figure 6 indicate that indeed both the coarse-grid and LPMs must be refined to achieve a decent convergence. Based on these results, it seems that for visual purposes a resolution of 54 directions for the coarse-grid and LPMs consisting of  $9 \times 9$  directions achieve reasonable convergence (the middle graph).

#### 4.2 Methods Comparison

In order to demonstrate the new method's advantages over a traditional DOM calculation, we test the two on the same input. We do that on a 3D volumetric model of a vase that defines the scattering coefficient  $\sigma$ , and we perturb it slightly to imitate the appearance of marble. The medium does not emit nor does it absorb light. We solve the linear equations arising from a first-order DOM with a spatial discretization of  $128^3$  cells and the total of  $3 \times 3 \times 6 (= 54)$ angular bins. Despite the large amount of memory needed for the grid (0.9GB), the results strongly suffer both from false scattering and the ray effect. Figure 2 shows a beam aimed towards the vase which gets smeared even before hitting the surface of the vase and shows a falsely smooth bright peak. From this lit region, light is scattered only along the discrete directions which proves to be insufficiently fine. The light distributes non-uniformly inside the vase cavity and creates an unrealistic bright peak just on the opposite vase wall (same horizonal level as the primary penetration region). Due to the strong smearing along the other directions, no other light peaks are observed, but nevertheless this illumination is highly inaccurate. Note that the use of higher order propagation scheme, may reduce the smearing, as shown earlier, but may also accentuate the ray effect. Using our approach, light can be propagated in more directions at lower storage costs (100MBs) as we show this figure. Use LPMs with  $9 \times 9$  angular bins we significantly reduce the ray effect without introducing any smearing. In Figure 7, the visual implication of these errors is shown. In this simulation a vase is lit by three light stripes which penetrate the vase walls and scatter. The excessive amount of smearing produced by the ordinary DOM does not allow us to detect the light pattern which is totally smeared. The result obtained using LPMs captures more accurately the light strips on both sides of the vase.

12 ·



Fig. 6. A beam of light entering from the left is passing through a constant Heyney-Greenstein scattering medium with g = 0.9. The coarse-grid angular resolution is refined from top row to bottom, consisting of the following numbers of bins  $1 \times 1(\times 6)$ ,  $3 \times 3(\times 6)$ , and  $5 \times 5(\times 6)$ . The LPM angular resolution is refined from left column to right, multiplying the coarse-grid angular resolution by factors of 1, 3, and 5.



Fig. 7. Perturbed  $\sigma$  vase lit by three bright light stripes. The left two images show the result of using a first-order DOM with a total of 54 directions (shown at two different viewpoints). And the right two images show the result of computing a solution using LPMs with an angular resolution of  $9 \times 9$  bins.

We also compare our method with results obtained from a Monte Carlo photon tracing simulation. We perform this test on an unperturbed vase model, where parallel light is emitted from the upper face of the domain to illuminate it. The light particles' initial position at the upper wall is determined randomly, and the scattering events are implemented via a Russian Roulette strategy. In Figure 8, we show the results obtained using  $10^6$  and  $5 \times 10^6$  particles to illuminate a grid of  $128^3$  cells. It takes 3.5 and 17.6 minutes respectively to perform these simulations. Even in the case where  $5 \times 10^6$  particles are used some amount of noise is still apparent in the final rendering. Only doubling this number of particles fully eliminates the perceptible noise and the solution converges to the one obtained using a  $9 \times 9$  LPM that takes 3.7 minutes to compute using 5 scattering iterations—one order of magnitude difference.



Fig. 8. Left two images show the results obtained using a Monte Carlo simulation with  $10^6$  light particles. The middle images were generated using  $5 \times 10^6$  particles, and the right images show the result obtained using  $9 \times 9$  LPMs. All three calculations were performed on a  $128^3$  cells grid.

#### 4.3 Computer Generated Scenes

As for typical CG scenes, requiring such a simulation, we use our method to illuminate marble, clouds and smoke. In Figure 10 two marble vases are illuminated from the top. One is modeled by a constant  $\sigma$  and a perturbed k, while the other has k = 0 and its  $\sigma$  is perturbed. In both cases, the method captured the complex inner marble structures interactions with the light penetrating the thick walls. This effect appears more dramatically when an intense laser beam hits the third vase. This example is made by two independent simulations, for red and white wavelengths. We used LPMs with the resolution of  $9 \times 9$  LPM angular bins to propagate the light in the red wavelength and  $5 \times 5$  bins for the white. Modeling scattering due to smoke dynamics involve complex  $\sigma$  fields, as shown in Figure 11. Light, radiated from spot projectors, is scattered when hitting the smoke and illuminates the whole domain. Cloud models, involving a denser  $\sigma$  coefficient, are shown in Figure 9. The illumination produced by the new method brings out clouds' structure and is able to generate some light streaks. In both cases the absorption and emission coefficients are set to zero. The running times of this method, with no hardware acceleration, are as following. Computing the solution over a resolution of  $64^3$  cells using LPM with angular resolution of  $5 \times 5$  directions and accounting for three scattering iterations (each involves six LPM sweeps) take 17 seconds. Using a  $9 \times 9$  LPM angular discretization this takes 135 seconds. It takes a minute to simulate a non-isotropic medium with coarse-grid angular discretization of  $3 \times 3 \times 6 (= 54)$  bins and  $64^3$  cells using  $6 \times 6$  LPMs. These numbers grow linearly with the number of variables used, and are comparable with the running times of standard iterative linear solvers.



Fig. 9. The upper-left image shows a layer of clouds lit from above, while in the lower images the clouds lit from their back. Both solutions are computed over a coarse-grid containing  $256 \times 128^2$  cells and LPMs with  $5 \times 5$  discrete angular bins. The upper-right image is made by compositing a layer of clouds with the image shown below it as the background. Here the clouds are modeled by a Henvey-Greenstein scattering with g = 0.25. Coarse-grid resolution consists of  $128^3$  spatial cells and  $3 \times 3(\times 6)$  angular bins. The LPM angular resolution used here consists of  $6 \times 6$  directions.

### 5. CONCLUSIONS

We presented a new method to transport light in the presence of a participating media. In this method, we exploit the simple pattern by which light travels in space to advect light fronts *collectively* and *independently*. This requires fewer computations per single light ray in comparison to stochastic photon-tracing approaches where each light-particle trajectory is calculated *individually*. Since these fronts are of one dimension less than the physical domain and consist of a reduced set of directions, high-resolution angular discretizations can be used in practice to reduce the errors due to angular truncations. The LPMs' detachment from the stationary coarse-grid  $I_{i,j,k}^m$ , inherently avoids interpolation artifacts such as oscillation and smearing.

The method presented here is robust and shows its ability to generate high quality calculations on a wide range of scenes with no special restriction over the media coefficients. Yet, as in the case of restricted subsurface scattering models, specifically developed solutions will be more efficient in their operation. Also, in cases where no such detailed, deep matter models are needed, volumetric representations would be less attractive than polygonal or other parametric surface models due to their low memory requirements and higher resolution.

The computational order of the new method in 3D is given by  $O(n(n^2 + p^2)mq)$  where *n* and *p* are the numbers of variables along each spatial dimension in the coarse-grid and LPM respectively, and *m* and *q* are the total numbers of angular variables in each of the coarse-grid and LPM bins respectively. The factor  $n^2 + p^2$  describes the number of cells visited when scanning each slice of the coarse-grid and the entire LPM. The factor *n* describes the progression



Fig. 10. The vase on the left is made of a constant  $\sigma$  and a perturbed k, while the ones in the middle and the right are modeled by k = 0 and a perturbed  $\sigma$ . In each of these three examples, the coarse-grid resolutions is 128<sup>3</sup> cells and the LPM angular discretization consists of  $5 \times 5$  orientations. Specular highlights were added during rendering.



Fig. 11. The model shown on the left is made of a constant  $\sigma$  and a perturbed *k*. The coarse-grid resolution is 256<sup>3</sup> cells, and the LPM angular discretization is 5 × 5 bins. Hygia sculpture courtesy of Image-based 3D Models Archive, Tlcom. The images on the right show smoke that scatters light and is lit by projectors, computed on a coarse-grid resolution of 64<sup>3</sup> cells and LPM angular discretization of 7 × 7 orientations.

along the third direction of the coarse-grid and the mq factor results from the coupling between the angular variables in each coarse-grid and corresponding LPM cells.

While this new approach allows finer ordinates discretizations to be used, light cannot be propagated in *every* direction. This means that a beam of uniformly oriented light can not be properly propagated unless its orientation coincides with one of the LPMs'. However, in most cases this inherent limitation of the DOM does not have any consequential

#### 16

•

implication; scattered light tends to be distributed over a continuous set of directions, which is shown to be well resolved by the proposed method. And direct light arriving from the boundaries (*e.g.*, infinite light sources) can be propagated in a single pre-sweep at that light's exact direction. Similar solutions can be practiced to illuminate scenes with a point or other types of light sources, adding no major computational costs.

In the case of highly scattering media, which is also the case of a highly-diffusing heat equation, iterative solvers require many iterations until they fully converge. The design of our method does not allow multi-resolution or other form of a change of basis for accelerating the convergence. Yet, for the variety of scenes we reported here, three iterations (along each direction) reached a visually satisfying result since the amount of light involves decays exponentially at each generation.

Iterative distribution of unpropagated light bears some resemblance to the Progressive Radiosity approach [Cohen et al. 1988]. As shown in [Gortler et al. 1994], Progressive Radiosity is equivalent to an iterative scheme known as *Southwell Relaxation*, used for solving linear equations just like Jacobi and Gauss-Seidel iterations. Although the process described here iteratively approximates a solution to a linear system, it does not correspond to a fixed point stationary iterative method. In fact, different orderings of the LPM sweeps result in different solutions (*i.e.*, differing from one another by a magnitude that is proportional to the truncation errors). Also, this process does not follow the strategy of distributing energy from the brightest element to the rest, according element-pair interactions.

Surfaces interchanging light with the medium can be added to the existing approach. Since they will be swept by the LPMs, they must be able to hold a constructed and unpropagated light across them which can be implemented via a mechanism similar to texture-mapping. Future steps related to this method can be to achieve a higher order of formal accuracy by devising the proper approximation formulas at each discretization step. The structure of the proposed algorithm may allow better parallelization than iterative solvers as the LPM sweeps do not dependent on one another. Also, given nowadays graphics processors supporting 3D textures, the possibility of a hardware implementation seems like a promising direction to speed up these calculations.

#### Acknowledgements

The author would like to thank Dani Lischinksi, Tamar Lando, James O'Brien, Jonathan Shewchuk, and Yossi Farjun for their help and comments.

#### REFERENCES

- ARQUÈS, D. AND MICHELIN, S. 1996. Proximity radiosity: Exploiting coherence to accelerate form factor computations. In *Rendering Techniques*. 143–152.
- BEHRENS, U. AND RATERING, R. 1998. Adding shadows to a texture-based volume renderer. In VVS '98: Proceedings of the 1998 IEEE symposium on Volume visualization. ACM Press, New York, NY, USA, 39–46.
- BHATE, N. AND TOKUTA, A. 1992. Photorealistic volume rendering of media with directional scattering. In Proceedings of the 3rd Eurographics Workshop on Rendering. 227–245.
- CEREZO, E., PEREZ-CAZORLA, F., PUEYO, X., SERON, F., AND SILLION, F. 2005. A survey on participating media rendering techniques. *the Visual Computer*.

CHANDRASEKHAR, S. 1950. Radiative Transfer. Dover, Ney York.

COELHO, P. J. 2002. Bounded skew high-order resolution schemes for the discrete ordinates method. J. Comput. Phys. 175, 2, 412-437.

- COELHO, P. J. 2004. A modified version of the discrete ordinates method for radiative heat transfer modelling. *Computational Mechanics 33*, 375–388.
- COHEN, M. F., CHEN, S. E., WALLACE, J. R., AND GREENBERG, D. P. 1988. A progressive refinement approach to fast radiosity image generation. In *In Proc. ACM SIGGRAPH* '88. ACM Press, New York, NY, USA, 75–84.

- DOBASHI, Y., KANEDA, K., YAMASHITA, H., OKITA, T., AND NISHITA, T. 2000. A simple, efficient method for realistic animation of clouds. In SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 19–28.
- GORTLER, S., COHEN, M. F., AND SLUSALLEK, P. 1994. Radiosity and relaxation methods. IEEE Comput. Graph. Appl. 14, 6, 48-58.
- HANRAHAN, P. AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques. ACM, New York, NY, USA, 165–174.
- HOTTEL, H. AND SAROFIM, A. 1967. Radiative Transfer. McGraw-Hill.
- JENSEN, H. W. AND CHRISTENSEN, P. H. 1998. Efficient simulation of light transport in scences with participating media using photon maps. In In Proc. ACM SIGGRAPH '98. ACM Press, New York, NY, USA, 311–320.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques. ACM, New York, NY, USA, 511–518.
- JESSEE, J. P. AND FIVELAND, W. A. 1997. Bounded, high-resolution differencing schemes applied to the discrete ordinates method. Journal of Thermophysics and Heat Transfer 11, 4, 540–548.
- KAJIYA, J. T. AND HERZEN, B. P. V. 1984. Ray tracing volume densities. In *In Proc. ACM SIGGRAPH '84*. ACM Press, New York, NY, USA, 165–174.
- LAFORTUNE, E. P. AND WILLEMS, Y. D. 1996. Rendering participating media with bidirectional path tracing. In *Proceedings of the eurographics* workshop on Rendering techniques '96. Springer-Verlag, London, UK, 91–100.
- LANGUENOU, E., BOUATOUCH, K., AND CHELLE, M. 1994. Global illumination in presence of participating media with general properties. Proceedings of the 5th Eurographics Workshop on Rendering, 71–86.
- LEVEQUE, R. J. 2002. Finite Volume Methods for Hyperbolic Problems, 1st ed. Cambridge Texts in Applied Mathematics.
- LIU, F., B. H. A. AND POLLARD, A. 1996. Spatial differencing schemes of the discrete-ordinates method. Numerical Heat Transfer B 30, 23.
- MAX, N. L. 1994. Efficient Light Propagation for Multiple Anisotropic Volume Scattering. In *Fifth Eurographics Workshop on Rendering*. Darmstadt, Germany, 87–104.
- PATTANAIK, S. N. AND MUDUR, S. P. 1993. Computation of global illumination in a participating medium by Monte Carlo simulation. *The Journal of Vis. and Comp. Animation 4*, 3 (July–September), 133–152.
- PAULY, M., KOLLIG, T., AND KELLER, A. 2000. Metropolis light transport for participating media. In *Rendering Techniques 2000 (Proceedings of the Eleventh Eurographics Workshop on Rendering)*, B. Peroche and H. Rushmeier, Eds. Springer Wien, New York, NY, 11–22.
- PEREZ, F., PUEYO, X., AND SILLION, F. X. 1997. Global illumination techniques for the simulation of participating media. In In Proc. Rendering Techniques '97. Springer-Verlag, London, UK, 309–320.
- PREMOZE, S., ASHIKHMIN, M., RAMAMOORTHI, R., AND NAYAR, S. 2004. Rendering of Multiple Scattering Effects in Participating Media. In *Eurographics Symposium on Rendering 2004*.
- RAMANKUTTY, M. A. AND CROSBIE, A. L. 1997. Modified discrete ordinates solution of radiative transfer in two-dimensional rectangular enclosures. Journal of Quantitative Spectroscopy and Radiative Transfer 57, 107–140.
- RUSHMEIER, H. 1988. Realistic image sythesis for scenes with radiatively participating media. Ph.D. thesis, Cornell University.

SIGEL, R. AND HOWELL, J. 1992. Thermal Radiation Heat Transfer, Third ed. Hemisphere Publishing Corp., Washington.

- SILLION, F. X. 1995. A unified hierarchical algorithm for global illumination with scattering volumes and object clusters. *IEEE Transactions on Visualization and Computer Graphics 1*, 3, 240–254.
- STAM, J. 1995. Multiple Scattering as a Diffusion Process. In *In Proc. Rendering Techniques '95*, P. M. Hanrahan and W. Purgathofer, Eds. Springer-Verlag, New York, NY, 41–50.
- STAM, J. 2001. An illumination model for a skin layer bounded by rough surfaces. In Proceedings of the 12th Eurographics Workshop on Rendering Techniques. Springer-Verlag, London, UK, 39–52.
- ZIENKIEWICZ, O. AND TAYLOR, R. 2000. The Finite Element Method. Butterworth-Heinemann.