

Global Illumination Techniques for the Simulation of Participating Media

Frederic Pérez^{*,**}, Xavier Pueyo^{*}, François X. Sillion^{**}

^{*} GGG/IIiA-UdG, ^{**} iMAGIS/GRAVIR-INRIA

Abstract: This paper surveys global illumination algorithms for environments including participating media and accounting for multiple scattering. The objective of this survey is the characterization of those methods: Identification of their base techniques, their assumptions, limitations and range of utilization. To this end, the algorithms are grouped into functional categories and each method is briefly reviewed, with a discussion of its complexity and its pros and cons. We finish by discussing some applications as well as remaining areas for investigation.

1 Introduction

The rendering of images containing participating media is important for a certain number of applications [15]. Simulations of interest can be made for the following areas:

- Safety analyses: smoke filled rooms (visibility of exit signs); foggy environments (roadway lighting, relative contrast of objects like traffic signs in foggy driving).
- Military: Remote sensing (atmospheric effects attenuate and blur images of land surfaces acquired by distant sensors); underwater vision; battlefield smoke plumes.
- Industrial: Design of efficient headlamps for foggy driving.
- Commercial: Entertainment, virtual reality.
- Visual simulation systems for the training of drivers of cars or ships for which optical effects in participating media are important; also fire fighter training.

We present in this paper a study of the different methods that have been proposed to address the global illumination problem of scenes including participating media, where the radiance field is modified at any point of the space, not only at surface points.

In the remainder of this section we review the transport equation and the simple scattering case. In Section 2 we present a general overview of the different algorithms, which are further discussed according to our classification. Section 3 deals with the deterministic methods, and Section 4 with the stochastic ones. A discussion of all methods is done in Section 5 and finally open lines of research are listed.

1.1 The transport equation

A quick review of the fundamentals of how light interacts with participating media is presented here [17]. As light travels through a participating medium, some phenomena take place: there can be absorption, emission and scattering of radiant energy. Absorption consists of a transformation of radiant energy into other energy forms, resulting in a reduction of radiance. For a differential distance $d\mathbf{x}$, this reduction is given by $\kappa_a(\mathbf{x})d\mathbf{x}$, being $\kappa_a(\mathbf{x})$ the *coefficient of absorption* of the medium at point \mathbf{x} . Emission refers to the

^{*} Postal address: Lluís Santaló s/n, E-17003 Girona, Spain.

E-mail: {frederic,xavier}@ima.udg.es

^{**} iMAGIS is a joint research project of CNRS/INRIA/INPG/UJF. Postal address: B.P. 53, F-38041 Grenoble Cedex 9, France. E-mail: Francois.Sillion@imag.fr

process of creation of radiant energy. Scattering means a change in the radiant propagation direction, reducing (out-scattering) the radiance (in this direction) along $d\mathbf{x}$ by the factor $\kappa_s(\mathbf{x})d\mathbf{x}$, being $\kappa_s(\mathbf{x})$ the *scattering coefficient*. Also radiance along the same direction is augmented because of the in-scattering, i.e. because of light coming to \mathbf{x} scattered into the studied direction. The spatial distribution of the scattered light is modeled by the *phase function* $p(\omega_o, \omega_i)$. Physically it expresses the ratio of scattered radiance in direction ω_o to the incoming radiance from direction ω_i by the radiance that would be scattered if the scattering were isotropic (i.e. independent of the direction).

Light interaction with a participating medium is governed by the *transport equation*, describing the variation of radiance in $d\mathbf{x}$ around \mathbf{x} in direction ω_o :

$$\begin{aligned} \frac{dL(\mathbf{x})}{d\mathbf{x}} &= \kappa_t(\mathbf{x})J(\mathbf{x}) - \kappa_t(\mathbf{x})L(\mathbf{x}) \\ &= \kappa_a(\mathbf{x})L_e(\mathbf{x}) + \frac{\kappa_s(\mathbf{x})}{4\pi} \int_{\mathcal{S}^2} L(\mathbf{x}, \omega_i) p(\omega_o, \omega_i) d\omega_i - \kappa_a(\mathbf{x})L(\mathbf{x}) - \kappa_s(\mathbf{x})L(\mathbf{x}) \end{aligned} \quad (1)$$

where $\kappa_t = \kappa_a + \kappa_s$ is the *extinction coefficient*, and $J(\mathbf{x})$ is the *source radiance*, which describes the local production of radiance, i.e. the radiance added to the point \mathbf{x} due to self-emission and in-scattering. Concretely,

$$J(\mathbf{x}) = \underbrace{(1 - \Omega(\mathbf{x}))L_e(\mathbf{x})}_{J_e(\mathbf{x})} + \frac{\Omega(\mathbf{x})}{4\pi} \int_{\mathcal{S}^2} L(\mathbf{x}, \omega_i) p(\omega_o, \omega_i) d\omega_i \quad (2)$$

where $\Omega = \frac{\kappa_s}{\kappa_t}$ is the *scattering albedo*. The solution of Equation 1 is the *integral transport equation*:

$$L(\mathbf{x}) = \underbrace{\tau(\mathbf{x}_0, \mathbf{x})L(\mathbf{x}_0)}_{L_{ri}(\mathbf{x})} + \underbrace{\int_{\mathbf{x}_0}^{\mathbf{x}} \tau(\mathbf{u}, \mathbf{x}) \kappa_t(\mathbf{u}) J(\mathbf{u}) d\mathbf{u}}_{L_m(\mathbf{x})} \quad (3)$$

being $\tau(\mathbf{x}_0, \mathbf{x}) = e^{-\int_{\mathbf{x}_0}^{\mathbf{x}} \kappa_t(\mathbf{u}) d\mathbf{u}}$ the *transmittance* from \mathbf{x}_0 to \mathbf{x} , $L_{ri}(\mathbf{x})$ the *reduced incident radiance*, due to the radiance of a background surface (if any), and $L_m(\mathbf{x})$ the *medium radiance*, due to the contribution of the source radiance within the medium [21]. Note that the source radiance can be decomposed into three terms, accounting for self-emission, for the (first) scattering of reduced incident radiance, and for the scattering of the medium radiance:

$$J(\mathbf{x}) = J_e(\mathbf{x}) + \underbrace{\frac{\Omega(\mathbf{x})}{4\pi} \int_{\mathcal{S}^2} L_{ri}(\mathbf{x}, \omega_i) p(\omega_o, \omega_i) d\omega_i}_{J_{ri}(\mathbf{x})} + \underbrace{\frac{\Omega(\mathbf{x})}{4\pi} \int_{\mathcal{S}^2} L_m(\mathbf{x}, \omega_i) p(\omega_o, \omega_i) d\omega_i}_{J_m(\mathbf{x})}$$

Most of the studied methods that render scenes including participating media are view independent, and thus they use two stages: the *Illumination Pass*, in which the source radiance $J(\mathbf{x})$ (or other equivalent function) is computed, and the *Visualization Pass*, in which Equation 3 is solved for the points of the image plane, using the results of the Illumination Pass.

When the participating medium has a low albedo or is optically thin (i.e. the transmittance through the entire medium is nearly one) then the source radiance can be simplified to not take into account the *multiple scattering* within the medium, considering this term negligible. Therefore, at point \mathbf{x} , the contribution of the scattering of the medium radiance $J_m(\mathbf{x})$ to the source radiance is set to zero, considering that \mathbf{x} is the first scattering point of the radiance coming from the background surfaces (*single scattering case*).

In this paper we will only focus on methods that deal with the more complex problem of multiple scattering (solving Equations 2 and 3).

2 Overview

We classify the existing methods into two main categories: deterministic and stochastic methods. Deterministic methods are further classified according to the space of directions, discerning between isotropic and anisotropic methods. All isotropic methods use constant basis functions for the computation of form factors. The very first of these methods is the zonal method [16], which is an extension to the classical radiosity method. This has been improved by using hierarchies within the context of the progressive refinements method [2, 19] and also of the hierarchical radiosity (HR) [18].

Deterministic methods can deal with anisotropy by means of spherical harmonics (P-N methods), discrete ordinates, or some implicit representation. Kajiya et al. [8] expand the radiance in a truncated spherical harmonic basis and construct a system of partial differential equations. Bhate et al. [3] expand the zonal method by using a spherical harmonic basis. Discrete ordinates refers to the discretization of the direction space into a set of bins [13, 10, 11]. Using a grid of voxels to model the participating media, the transport equation can be solved locally per voxel, thus updating its exiting radiance, i.e. changing the incoming radiance of its neighbors, which must in turn solve their exiting radiances [13, 10]. Alternatively the energy exchange between all pairs of elements can be considered, as zonal method’s extension, setting and solving a system of equations whose coefficients are form factors. Max [11] approximates the effects of the form factors avoiding their computation.

Finally some methods use an implicit representation of the directional distribution of radiance (encoded either in scattering patterns [12] or a diffusion equation [20, 21]). In the method of Nishita et al. [12], the contributions to the radiance $L(\mathbf{x})$ (for the second and third orders of scattering) in the viewing direction in a point \mathbf{x} interior of a participating medium come in the form of a set of extended form factors in a grid that forms a 3d-filter. These form factors must be multiplied by the energy at the related points and accumulated to get $L(\mathbf{x})$. Stam [20, 21] uses a “diffusion approximation” to solve the multiple scattering between blobs modeling the media. Restricting the medium source radiance of a blob to be of a simple form, a diffusion equation can be written as a system of linear equations allowing the calculation of the source radiance for each blob.

Stochastic methods solve the transport equation by means of random sampling, using random paths along interaction points. We distinguish between the methods that set the interaction points by using a constant step distance [4, 5], from those that sample a function of κ_t [14, 9]. Another categorization is made according to the view dependency of the methods. We tag a method as view dependent if it is image based or if in the Visualization Pass to solve the pixel radiance (Equation 3) to get the value of $L(\mathbf{x}_0)$ some extra process is needed (e.g. to use a ray tracing).

The studied papers are summarized and categorized as discussed above in Table 1. Entries in italic style denote methods that don’t solve the global illumination problem, in the sense that in the scene there is only a single volume to illuminate. Symbols and terms used in this paper are provided in Table 2.

3 Deterministic methods

3.1 Constant basis functions

Zonal method. The zonal method [16] is the extension to the classical radiosity method including isotropic participating media, modeled by voxels. The voxel radiosity is defined to include only the self-emitted plus the scattered energy (source function for a

Space of directions						
Isotropic		Anisotropic				
Constant basis functions		Spherical harmonics	Discrete ordinates		Implicit representation	
Zonal method: [16]		[8]	Based on local solutions (local interactions)		Global interactions	3d-filter: [12]
Hierarchy		[3]	Progr. Ref.: [13]	Sweeps: [10]	Sweeps: [11]	Diffusion: [20, 21]
Progr. Ref.: [2, 19]	HR: [18]					

Distance sampling		
	Constant	Random
view independent	Light tracing: [4]	Light tracing: [14]
view dependent	Light tracing: [5]	Bidirectional path-tracing: [9]

Table 1. Deterministic (top) and stochastic (bottom) methods.

Symbol	Units	Meaning
$J(\mathbf{x}, \omega), L(\mathbf{x}, \omega)$	$\text{W m}^{-2} \text{sr}^{-1}$	Source radiance and radiance
M	—	Number of direction bins / harmonic coefficients
n	—	Number of elements per axis in a regular grid
$p(\omega_o, \omega_i)$	—	Phase function
$s, v \quad (v = n^3)$	—	No. of patches / volume elements (voxels or blobs)
\mathbf{x}	—	A point in \mathcal{R}^3
$\kappa_a(\mathbf{x}), \kappa_s(\mathbf{x}), \kappa_t(\mathbf{x})$	m^{-1}	Absorption, scattering and absorption coefficient
$\tau(\mathbf{x}_1, \mathbf{x}_2)$	—	Transmittance
ω	— / sr	Direction / Solid angle
$\Omega(\mathbf{x})$	—	Scattering albedo

Table 2. Table of terms.

voxel). Form factors between volumes, and between volumes and surfaces are defined, and the form factors between surfaces are redefined to include a transmittance factor. They are computed by extending the hemicube technique. A system of s (for surfaces—patches) plus v (for volumes—voxels) related equations is constructed, and solved by the Gauss-Seidel iterative method. The direct application of the zonal method has a prohibitive cost: In a regular cube of n^3 voxels there are n^6 form factors; approximating them by the 1-d integral along the centers of each pair of voxels in time $\mathcal{O}(n)$ (i.e. number of intervening voxels) the computation of all form factors takes $\mathcal{O}(n^7)$. Coherence between form factors has been exploited to compute them with lower cost [1].

Progressive refinement approach. These methods [2, 19] establish a fixed hierarchy in a preprocessing step and thereafter use it in a shooting strategy. There aren't further refinements of the hierarchy that would enable the computation of a cheaper coarse solution that could be iteratively improved by refining it [7].

Bhate's method [2] is a progressive refinement version of the zonal method, using hierarchies. These hierarchies are computed in two preprocessing steps, which consist in the subdivision of volumes and surfaces and the creation of links between volumes and volumes and also between surfaces and volumes, determining the level at which a pair of elements must interact. The proposed heuristics for the volume-volume refinement are: Total form factor (when a rough estimate of the form factor is below some specified threshold, then the related elements can interact at the current level), estimated visibility between the volumes, and the optical depth of the intervening medium. Basically, in

these latest heuristics, when the transmittance between two elements is too high, there is no need for further refinements. The volume-surface refinement also include a brightness factor heuristic (for light sources).

Some particular observations are:

- *Self-refinement*. Participating media are modeled by a set of *global* volumes. Each global volume is refined against each other, but not against itself (there is no self-refinement). Therefore links must be set between each pair of the smallest volume element considered within a given global volume, otherwise its self-illumination won't be correctly computed. A self-refinement strategy would be better since it would reduce the number of interactions.
- *Push-Pull*. There is no Push-Pull procedure to set correctly the values of the radiosity at all levels of the hierarchies. Such a procedure would be needed after each shooting iteration to obtain a correct solution.

In the Sobierajski's method [19] the volumetric data is represented by voxels, that can model lambertian surfaces, isotropic media, or a combination of both. Thus each voxel's BRDF is the sum of ideal diffuse reflection plus isotropic scattering. Depending on the specific coefficients for each component, a voxel can have a more translucent volumetric appearance or resemble more an opaque surface. Therefore, each voxel has a *diffuse* plus an *isotropic* radiosity. Form factors are defined to take into account the relationships between diffuse and isotropic components of the voxel's BRDF, and the surfaces.

The presented technique is an iterative shooting algorithm using hierarchies which—for the case of volumes—are built in a preprocessing step by combining eight neighboring voxels at a certain level to form one voxel of the parent level. A criterion is defined to decide if a parent voxel can be a good approximation of its descendents. The interaction between nodes depends on their levels, their averaged values and the amount of energy transferred between them. There aren't explicit links between nodes, instead at each shooting iteration the best highest possible levels of interaction are found *on-the-fly*. After each shooting iteration a Push-Pull procedure assures the correct representation of the energies of all the nodes in the hierarchies.

Some specific aspects are:

- *Self-interaction*. A volume object cannot self-interact. The self-lighting should be considered but will be missed in volumes that aren't convex solid objects.
- *Brightness-weighted interaction*. A hierarchy is computed independently for each "volumetric object", grouping the low level voxels recursively to form upper level voxels, and no links exist between different hierarchies. Those links are implicitly computed when finding the interaction levels at which two elements are allowed to interact. A test is done to check if the amount of radiosity shot from an element e_j to another element e_i is below a given threshold. This could be modified to check if the estimated energy shot from e_j arriving to e_i is under some threshold (using an estimated form factor). This technique would decrease dramatically the number of interactions when dealing with optically thick media.
- *Push-Pull*. The Push-Pull procedure presented limits the application of the method to media with constant albedo (pushing irradiance instead of radiosity would solve this drawback [18]). It should be noticed that maybe it would be more interesting to use a gathering technique instead of a shooting technique. Performing the Push-Pull procedure after each shoot might imply too high a cost.

Hierarchical radiosity. Sillion [18] presents a hierarchical radiosity algorithm [7] adapted to include isotropic volumes. To represent energy exchanges within a volume, the self-link (link from the volume to itself) is introduced. This link is subdivided in a different way from links between different elements, since each child must include a self-link apart from the usual links between each pair of children. Furthermore, to avoid the quadratic cost of the initial link phase of the classical hierarchical method, that can be overwhelming in complex scenes, the transfer of energy between groups of objects (i.e. sets of surfaces and volumes) is allowed. These groups of objects compose abstract objects (clusters) that exchange energy as a whole. A hierarchy is created above the surface level, and then the initial linking phase is reduced to the creation of a single self-link from the top of the hierarchy to itself, representing the interactions taking place inside the global volume enclosing the scene. Once the initial link is refined by a recursive procedure, gathering and Push-Pull steps are performed until there is no significant change in the radiosities of any element. Care must be taken to perform correctly the Push-Pull procedure when dealing with inhomogeneous media and textured surfaces. Refinement of the links is done by bounding the radiosity transfer.

3.2 Spherical harmonics.

Kajiya et al. [8] present two methods. The first one deals with single scattering, and the second with multiple scattering within the participating media. The radiance is expressed in a truncated spherical harmonics basis, and a system of partial differential equations (PDE's) is constructed for the spherical harmonics coefficients. The system of PDE's is set and solved by relaxation. Only the constant phase function and the Rayleigh phase function are considered, and the expansion in spherical coordinates is truncated after the fourth coefficient (because "only the first few spherical harmonics are necessary for a convincing image", but obviously the cost of the method depends largely on the number of coefficients). The effects between surfaces and volumes are not taken into account.

The method by Bhate et al. [3] deals with the effects between surfaces and volumes missed in [8], being an extension to the zonal method, in which the assumption of the isotropy of the medium is eliminated (through a representation of the phase function and radiance by using spherical harmonics) and the surfaces remain ideal diffuse. The phase function is approximated by the first M terms of its spherical harmonics expansion (approximation to Mie scattering). Note that a large number of form factors will have to be computed, taking into account the spherical harmonics. These are calculated with the extended hemicube technique (cf. 3.1). Finally, a system of Mv equations for v volumes plus s related equations for s surfaces is set and solved using a Gauss-Seidel iterative technique. The direct application of this method is impractical because of its prohibitive cost: In a regular grid of $v = n^3$ voxels the cost to compute the form factors is $\mathcal{O}(n^7 + M^2n^6)$.

3.3 Discrete ordinates

Another possibility to account for directional functions is the use of discrete ordinates, i.e. a discretization of the full 4π solid angle into a set of bins. These represent particular directions, and it is supposed that, for sufficiently small volume elements, the properties are constant for each direction within each volume. The main problem of the discrete ordinates is the "ray effect" problem since the energy is propagated through discrete directions instead of into the whole discrete solid angle.

Local interactions. Patmore [13] formulates the local solution of the transfer equation for the discrete directional model resulting of the subdivision of the volume (resulting in a cubic lattice) and the angular spaces (using in practice 6 or 26 directions). The participating medium considered is non-emitting, since the objective is to render clouds. A global solution of the transfer equation is obtained through iteratively obtaining local solutions (related to points of the cubic lattice). As a consequence of a local solution the unshot energies of the related point are updated. A new local solution is computed for the lattice point adjacent and in the direction of the highest unshot energy of the previous one, thus effectively following importance-based paths, until the unshot energy is below some threshold or the path exits the volume. This method computes directly the radiances exiting the volume, so no integration of source radiances are needed in the visualization pass.

The method by Languénoü et al. [10] follows a progressive refinement approach. The usual shooting method for surfaces is extended to account for the transmittance through the media, and also source terms within the media are updated accordingly. The radiosities of the boundaries of the media are computed propagating the radiance (coming from the previously accumulated source terms) along all the discrete ordinates and using as many iterations as necessary to converge. Each iteration consists in a loop for each direction, in which a complete sweep of the voxel grid is performed to propagate the accumulated energy through adjacent voxels, starting from a convenient boundary voxel (related to the direction considered), where $\mathcal{O}(Mv)$ is the cost per iteration. Finally, the radiance of the boundary faces of the medium is shot, using hemisphere interpolation. The whole process is repeated until convergence is met. The Visualization Pass computes the pixel radiances by using the source radiances of the voxels.

Global interactions. Max’s method [11] is devoted to render clouds. The computation of the M^2v^2 form factors of the finite elements formulation is avoided by approximating their effects as the energy is propagated across the grid. For each bin, this propagation is made distributing the flux to the related neighbor voxels simultaneously for all voxels belonging to a layer, in time $\mathcal{O}(v \log n)$, for v^2 interactions. The “ray effect” is reduced because the energy is propagated through the whole bin, not only through a single direction. The attenuation between two voxels is not accumulated along the straight line joining them, but along a set of possible propagation paths. The multiple scattering events produced within a single receiving element are accounted for. Since the time to scatter the received flux of a voxel to the direction bins is $\mathcal{O}(Mv)$, the final cost per iteration is $\mathcal{O}(Mv \log n + M^2v)$. Thus when the number of iterations required to converge is small compared to v , this method is better than computing the whole set of form factors (with a cost of $\mathcal{O}(v^2n + v^2M)$) and solving the resulting system.

3.4 Implicit representation

The directional distribution of radiance can be represented implicitly by a scattering pattern [12] or by a diffusion equation [20, 21].

3d-filter. Nishita et al. [12] propose a method to display clouds taking into account multiple scattering and sky light (light reaching the cloud due to the atmosphere’s scattering plus the reflected light from the earth’s surface). Radiance from a cloud reaching the eye is computed from the sunlight multiple scattered plus the sky light single scattered by the particles of the cloud. Of the multiple scattering of the sunlight only the three first orders of scatterings are considered (to save computation time), computing separately the single scattering. For the contribution of the second and third order of scattering to

the radiance in the viewing direction, the space including the cloud is subdivided into voxels, with the viewing direction as a principal axis. Instead of computing form factors between each pair of voxels, a smaller space with the mean density of the cloud is set, and the contribution ratios to the radiance of the center voxel (in the viewing direction) from the other voxels are computed, taking into account the sunlight direction. This is the *contribution-ratio pattern*, or 3d-filter. Since the scattering in clouds is mainly forward, most of the energy scattered at a point will lie within a relatively small solid angle. Using this fact it's possible to compute faster the extended form factors, concentrating the effort in those voxels which will effectively contribute to the center voxel, for paths having one or two scatterings (and using an stochastic method to select those voxels). The filter is applied to the voxels in the whole space storing for each voxel the light scattered due to the second and third order scattering in the viewing direction.

Diffusion. Stam [20, 21] solves the global illumination by progressive refinements, by using shooting operations between patches, and between patches and blobs (which model the media). The shooting between blobs (that could be very expensive if the number of blobs v is large) is avoided by a set of v linear equations representing a diffusion equation. This is obtained by a "diffusion approximation" of the source radiance (due to the scattering of the medium radiance), i.e. it is characterized by only two functions: $J_m(\omega) = J^0 + \mathbf{J}^1 \cdot \omega$. Solving the linear system allows the computation of the coefficients J^0 and \mathbf{J}^1 for each blob, and thus the multiple scattering between blobs. When v is not too large ($v < 1000$) the system can be solved with a direct LU-decomposition; for larger systems a relaxation scheme can be used, although the convergence is not guaranteed (but note that with a relatively small number of blobs good looking results are obtained). The proposed method uses far less memory and computation time than would be required by a grid method. Being a progressive method, when it deals with complex scenes composed of lots of surfaces, the cost of the progressive shooting of energy from the surface patches is quite expensive. A hierarchical approach would become necessary in such a case.

4 Stochastic methods

The global illumination stochastic methods basically trace random rays within the environment. The interaction points that limit the rays can be obtained by using a constant step distance [4, 5] or sampling a cumulative density function [14, 9].

4.1 Constant distance sampling

Blasi et al. [4, 5] describe methods to deal with participating media by using a simulation of the particle model of light (Monte Carlo light tracing). The first [4] deals with a single participating medium; the second [5] can render mixed scenes. Both take into account the multiple scattering within the media, using the Schlick [4] phase function, specially defined in such a way that the importance sampling using it is quite inexpensive, while maintaining the possibility of approximating other phase functions. In [4] it's used an approximation to the Mie scattering as a combination of isotropic plus forward scattering components. In the scattering events, the scatter direction is given by optimal importance sampling of the scattering component, and the isotropic part is stored in the voxel. This isotropic part of the voxel is not considered for the illumination of the other voxels. Since the directional component is much more important than the isotropic component, it is expected that the resulting error will not be significant. A progressive refinement strategy could be used when this isotropic energy becomes too important.

Bundles progress in steps of constant length. Therefore, at each interaction point there is a sampling process to decide if there is scattering in that point. Absorption is taken into account along the whole path of the bundle, decreasing its flux at each step by the transmittance due to absorption along distance δ .

In [5] a progressive technique is used to render mixed scenes. Surfaces are classified as “diffuse” or “specular” depending on some threshold. In the Illumination Pass, when a bundle hits a diffuse surface, its energy is stored there (and the bundle’s path ends), whereas when it hits a specular one, it is reflected using importance sampling. Within the media the bundles progress as explained above, although only when a bundle exits the media (if it does) its energy is recorded (at the border voxel). Due to this storage scheme, the number of rays traveling through the volumes must be higher than it would be required with a storage per voxel, to get an accurate sampling of the energy leaving the volume.

4.2 Random distance sampling

Light tracing. The Monte Carlo light tracing by Pattanaik et al. [14] uses a sampling process to find the points of interaction (absorption or scattering) of the bundles within the volume, with the expression $1 - \exp(-\int_0^S \kappa_t(u)du)$ as a cumulative distribution function, where S is the distance traveled. At those points, with the *Simple Absorption* method, another sampling process is performed to decide if the interaction is an absorption or a scattering event, based on Ω . On the other hand, with the *Absorption Suppression* method the bundles always scatter but they reduce their flux multiplying it by Ω . Different variance reduction techniques are proposed: Forced interaction of a bundle with each voxel, the yet mentioned Absorption Suppression method, and the Particle Divergence method (in which the outgoing bundle is split into many bundles at the scattering points). The storage scheme presented is suited for isotropic scattering, but can be changed to deal with anisotropic scattering.

Bidirectional path tracing. In [9] a bidirectional path tracing for non-emitting participating media is presented. Random walks are traced both from the light sources (light paths—light shooting) and from the eye point (eye paths—light gathering), being a combination of light tracing and eye tracing. Consequently this is an image-based method. After tracing a light path and an eye path, each intersection point of the respective paths are connected by shadow rays. Those shadow rays that aren’t occluded constitute a part of (complete) transport paths from the light sources to the eye, and an illumination contribution is computed for each transport path. These illumination contributions are combined to obtain an unbiased estimator for the radiance reaching the eye, taking into account the probability densities for generating the transport paths used. Concretely, the balance heuristic [22] is used to obtain the weights of the illumination contributions. Random walks (both light and eye rays) are traced computing interaction points within the media as in [14] (Simple Absorption case). For the scattering direction computation the Schlick phase function is used.

5 Discussion

5.1 Progressive results

The multi-gridding technique can be used to compute a sequence of solutions stopping when a sufficiently accurate solution is obtained. The sequence starts with a very rapid computation of a first coarse solution which is improved in successive steps. This can

be accomplished in hierarchical approaches like in [18]. If a relatively small quantity of time is given to compute a solution, then with the multi-gridding technique a coarse solution could be obtained. The Monte Carlo light tracing method of Pattanaik et al. [14], on the other hand, in a given short time will produce in the Illumination Pass a partial solution far from converged in the illumination elements, so in practice the image related to that partial results won't be of utility. This is due to the fact that each bundle follows its path in the scene until it dies. In the light tracing method by Blasi et al. [5], however, the reflection on diffuse surfaces is eliminated, so that whenever a bundle hits a diffuse surface its path ends, and that diffuse surface accumulates unshot energy. At each iteration a set of bundles representing the unshot energy of the element having the highest value is spread to the environment, thus being a progressive refinement algorithm. Note that this technique also introduces bias since the process of reflection at a given point is substituted by a shooting from a random point within the element. The progressive nature of [5] allows the computation of an iterative sequence of images. However, the illumination of the media will be far from converged unless a very high number of bundles have been used. In the bidirectional ray tracing, since it's a view dependent method where the illumination is solved directly per pixel, the quality of the image can be gradually improved, starting from a very crude approximate image and converging to the solution as the program progresses.

Partial results of the illumination of a volume treated by the methods which use discrete ordinates by sweeping of energy [10, 11] could be given between successive iterations. Methods that do progressive refinements using hierarchies could do a good job if an ambient term is used for display purposes (as a generalization of the ambient term of the classical radiosity) after the shoot of light of the most energetic elements. This can also be used by other progressive refinements methods like the diffusion approximation [20, 21], and the method by Blasi et al. [5].

5.2 Sampling strategies

It should be noted that the sampling strategy of [5] drives to biased results, while that of [14] doesn't. This is because the bundles can only be scattered at distances which are multiple of δ , and thus the expected length before scattering won't be equal to the mean free path without absorption. The error is reduced as long as the value of δ is diminished. Unfortunately to assure results with a variance below some threshold the time required is approximately of the order of the inverse of δ . Moreover, we have checked that, for a same variance threshold and setting a value of δ relatively small to get a tolerable bias, the computation time using the sampling procedure of [14] is always lower than that of [5].

5.3 Isotropic media

It seems clear that for applications in which the isotropic assumption can be used (i.e. non-realistic applications), the obvious choice is the hierarchical radiosity method [18] since it has the best performance in computation time; moreover it is more reliable than the progressive refinements methods, in which a fixed hierarchy is used. Also the diffusion approximation using blobs method [20, 21] could be simplified for the isotropic scattering case, being then a good choice (less memory and computation time than grid based methods) when the number of blobs is relatively small. When that number of blobs is high, then a hierarchical approach becomes necessary.

5.4 Anisotropic media

In the case of anisotropic media, all the existing methods commit errors; only the bidirectional path tracing is unbiased. It's then important to know what type of error can be accepted for each concrete application. Stam [20, 21] utilizes a “diffusion approximation” (cf. 3.4). This is only valid in the case of a high number of scattering events. This condition fails at the boundaries of the media, and thus the results aren't so precise there; however, for certain applications (such as animations in which the objective is that things “look right”) they are accurate enough. Spherical Harmonics and Discrete Ordinates methods approximate directional functions by using a fixed set of bases. It could be interesting to use an adaptive number of bases in function of the accuracy required for the solution. Max [11] propagates the energy inside the medium in a way so that the “ray effect” (present in [13, 10]) is reduced, but computing an approximation of the true attenuation between two voxels; therefore although it produces visually better images it's not clear if that solution is less accurate than the direct Discrete Ordinates method (although for displaying clouds, for example, it's considered to be better). Obviously the bias of the Spherical Harmonics and Discrete Ordinates methods can be reduced by using a higher number of bases, but the computation time augments dramatically doing so. Thus there is a compromise between computation time and image quality.

Comparing the costs of the zonal Spherical Harmonics method [3] (form factor computation: $\mathcal{O}(n^7 + M^2n^6)$) against Discrete Ordinates methods (iteration cost: $\mathcal{O}(Mv)$ in [10], $\mathcal{O}(Mv \log n + M^2v)$ in [11]—15 iterations are used approximately; form factor computation for the direct extension of the zonal method with discrete ordinates: $\mathcal{O}(v^2n + v^2M)$), with $v \gg M$, we can order the different methods starting from the cheaper as follows: Languénoü et al.'s discrete ordinates [10], Max's discrete ordinates [11], direct extension of the zonal method with discrete ordinates, and the zonal spherical harmonics by Bhate et al. [3]. However it should be noticed that the accuracy of the results follows the reverse order.

6 Future Research

There are some strategies not addressed in the existing methods which could constitute lines of future research. We discuss below the eventual interest of them:

- The use of clusters/hierarchies for anisotropic media would allow the resolution of the global illumination problem for complex scenes in far less time than that required by Monte Carlo methods. These clusters could contain also diffuse and glossy surfaces, as in [6]. For applications which require dealing with specular surfaces maybe a multi-path algorithm should be explored.
- A multiresolution representation of energies, phase functions, etc., could be also accomplished by utilizing wavelets.
- The notion of importance (or potential) of surfaces can be extended to participating media. Thus methods which make use of importance to accelerate the computation of the radiances for a concrete view could be analyzed to be extended for scenes including participating media.

The possibility of giving bounds or characterizations of the error committed by a method could also be studied (in general, or for concrete executions). Finally, more work should be done in order to deal with dynamic environments, walkthroughs and extensions to other physical phenomena like fluorescence and phosphorescence.

Acknowledgements. The first author has developed part of this work at iMAGIS in Grenoble with a grant from the ERCIM Computer Graphics Network, and part within the context of CICYT's project TIC95-0630-C05-05.

References

1. D. Arquès and S. Michelin. Proximity Radiosity: Exploiting Coherence to Accelerate Form Factor Computations. *6th EG Workshop on Rendering*, pp. 144–153, June 1996.
2. N. Bhate. Application of Rapid Hierarchical Radiosity to Participating Media. *Proc. of ATARV-93: Advanced Techniques in Animation, Rendering and Visualization*, pp. 43–53, July 1993.
3. N. Bhate and A. Tokuta. Photorealistic Volume Rendering of Media with Directional Scattering. *3rd EG Workshop on Rendering*, pp. 227–245, Bristol, UK, May 1992.
4. P. Blasi, B. LeSaëc and C. Schlick. A Rendering Algorithm for Discrete Volume Density Objects. *Computer Graphics Forum*, 12(3):201–210, September 1993.
5. P. Blasi, B. LeSaëc and C. Schlick. An Importance Driven Monte-Carlo Solution to the Global Illumination Problem. *5th EG Workshop on Rendering*, pp. 173–183, June 1994.
6. P.H. Christensen. *Hierarchical Techniques for Glossy Global Illumination*. PhD thesis, Department of Computer Science and Engineering, University of Washington, 1995.
7. P. Hanrahan, D. Salzman and L. Aupperle. A Rapid Hierarchical Radiosity Algorithm. *Computer Graphics*, 25(4):197–206, July 1991.
8. J.T. Kajiya and B.P. Von Herzen. Ray Tracing Volume Densities. *Computer Graphics*, 18(3):165–174, July 1984.
9. E.P. Lafortune and Y. Willems. Rendering Participating Media with Bidirectional Path Tracing. *6th EG Workshop on Rendering*, pp. 92–101, June 1996.
10. E. Languénoù, K. Bouatouch and M. Chelle. Global Illumination in Presence of Participating Media with General Properties. *5th EG Workshop on Rendering*, pp. 69–85, June 1994.
11. N.L. Max. Efficient Light Propagation for Multiple Anisotropic Volume Scattering. *5th EG Workshop on Rendering*, pp. 87–104, June 1994.
12. T. Nishita, Y. Dobashi and E. Nakamae. Display of Clouds Taking Into Account Multiple Anisotropic Scattering and Skylight. *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '96*, pp. 379–386, 1996.
13. C. Patmore. Simulated Multiple Scattering for Cloud Rendering. In *Graphics, Design and Visualization*, pp. 59–70, North-Holland, Amsterdam, The Netherlands, 1993.
14. S.N. Pattanaik and S.P. Mudur. Computation of Global Illumination in a Participating Medium by Monte Carlo Simulation. *The Journal of Vis. and Comp. Animation*, 4(3):133–152, 1993.
15. H. Rushmeier. Rendering Participating Media: Problems and Solutions from Application Areas. *5th EG Workshop on Rendering*, pp. 35–56, June 1994.
16. H.E. Rushmeier and E. Torrance. The Zonal Method For Calculating Light Intensities in the Presence of a Participating Medium. *Computer Graphics*, 21(4):293–302, July 1987.
17. R. Siegel and J.R. Howell. *Thermal Radiation Heat Transfer*, 3rd ed. Hemisphere Publishing, Washington, 1992.
18. F.X. Sillion. A Unified Hierarchical Algorithm for Global Illumination with Scattering Volumes and Object Clusters. *IEEE Trans. on Vis. and Comp. Graphics*, 1(3):240–254, Sept. 1995.
19. L.M. Sobierajski. *Global Illumination Models for Volume Rendering*. Chapter 5: Volumetric Radiosity, pp. 57–83, PhD Thesis, 1994.
20. J. Stam. Multiple Scattering as a Diffusion Process. *6th EG Workshop on Rendering*, Dublin, Ireland, 1995, pp. 41–50.
21. J. Stam. *Multi-Scale Stochastic Modelling of Complex Natural Phenomena*. PhD Thesis, Dept. of Computer Science, University of Toronto, 1995.
22. E. Veach and L. Guibas. Optimally Combining Sampling Techniques for Monte Carlo Rendering. *Computer Graphics Proc., Annual Conf. Series: SIGGRAPH '95*, pp. 419–428, 1995.