Computers & Graphics 45 (2014) 28–39

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Special Section on Graphics Interaction

Principal-Ordinates Propagation for real-time rendering of participating media

Oskar Elek [a,b,*], Tobias Ritschel [a,b], Carsten Dachsbacher [c], Hans-Peter Seidel [a,b]

[a] Max Planck Institut Informatik, Germany
[b] MMCI Cluster of Excellence, Saarland University, Germany
[c] Karlsruhe Institute of Technology, Germany

ARTICLE INFO

Article history:
Received 4 July 2014
Received in revised form
22 August 2014
Accepted 22 August 2014
Available online 6 September 2014

Keywords:
Participating media
Light scattering
Natural phenomena
Real-time rendering
Physically based rendering
Finite elements

ABSTRACT

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

© 2014 Elsevier Ltd. All rights reserved.

1. Introduction

Scattering, or translucency, greatly contributes to the appearance of many natural substances and objects in our surrounding. Albeit the problem can be easily formulated as the radiance transfer equation [3,23], computing a solution can be very costly. Consequently, many existing approaches simplify the problem, e.g. by assuming isotropic scattering or homogeneity of the material, to achieve interactive performance.

In this work we propose a novel algorithm for plausible real-time rendering of heterogeneous participating media with arbitrary anisotropy "... anisotropy (see Fig. 1 for an example).". The core of our approach is to propagate light in propagation volumes oriented along the *principal ordinates* of the source illumination. For this we typically use multiple rectilinear grids to propagate environmental (distant) lighting, and spherical grids to account for point light sources. In both cases, one dimension of the grids is aligned with the prominent directional part of the source radiance for which the grid has been created. In contrast to previous methods (e.g. [15,1]), discretizing the illumination into directional and point light sources enables us to approximately describe the anisotropy (directionality) of light transport by a single scalar value per grid cell. Specifically, this anisotropy value corresponds to a unimodal function implicitly aligned with the respective principal ordinate. In addition to exploiting data locality and the parallelism of GPUs, the benefit of these decisions is a significant reduction of the *false scattering* (numerical dissipation) and *ray effect* (misalignment errors) artifacts arising in many finite-element methods as a consequence of representing the propagated radiance by, e.g. spherical harmonics or piecewise-constant functions. Our main contributions can be summarized as follows:

- We propose the concept of Principal-Ordinates Propagation (POP), a deterministic finite-element scheme suitable for real-time simulation of anisotropic light transport in heterogeneous participating media (Section 3).
- The theory of iterative light propagation in a uniform Euclidean grid using a minimal unimodal propagation basis and explicit alignment with the illumination direction to minimize propagation artifacts and maintain light directionality (Section 4).
- An extension of the propagation scheme to handle environmental illumination by decomposing it in a set of discrete directions. This includes several new steps, namely specialized

* Corresponding author.
E-mail address: oelek@mpi-inf.mpg.de (O. Elek).

http://dx.doi.org/10.1016/j.cag.2014.08.003
0097-8493/© 2014 Elsevier Ltd. All rights reserved.

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

prefiltering, importance propagation, and a separate propagation of isotropic residual energy (Section 5).

- An extension to local light sources via spherical grids, enabling the integration of instant radiosity to simulate light interaction between solid objects and the medium (Section 6).
- Finally we analyse our approach in a number of diverse scenarios, demonstrating its versatility (Section 7).

2. Related work

Offline methods: A range of different approaches has been presented to compute solutions to the radiance transport equation for participating environments [3,23]. However, none of the classic techniques provides a satisfying combination of generality, robustness, and, most importantly in our context, speed. Unbiased Monte-Carlo methods, such as bidirectional path tracing [20] and Metropolis light transport [29], usually require a large number of paths to be traced; in particular in dense media with high scattering anisotropy and albedo (like clouds or milk) the computation time increases tremendously. Caching is often used to speed up the computation, e.g. radiance caching [12], photon mapping [14,13] or virtual point lights [8]. However, these methods typically do not handle highly anisotropic scattering very well, even with recent improvements [27,28], and their performance is often far from interactive.

Finite-element methods: Finite-element methods, including volume radiosity [33], the discrete ordinates method (DOM) [3], light diffusion [36], and lattice-Boltzmann transport [10] handle highly multiple scattering well. However, in practice they allow only isotropic or moderately anisotropic scattering, and usually suffer from false scattering (smoothing of sharp light beams) and ray effects (selective exaggeration of scattered light due to discretized directions). Light propagation maps [9] significantly reduce the artifacts, but are still limited to rather low scattering anisotropy.

It can therefore be seen that strong scattering anisotropy is one of the main limiting factors for existing methods. This is unfortunate, as most real-world media exhibit relatively high anisotropy (Henyey–Greenstein [11] coefficient $g \approx 0.9$ or more [26]). Although isotropic approximations are acceptable in some cases, this is generally not a valid assumption and one of the primary motivations for our work.

Interactive rendering: Numerous works focus on individual optical phenomena to achieve interactive or real-time performance. These phenomena include light shafts [32,7], volume caustics [19,21], shadows [22,34], and clouds [2]. Various approaches can also be found in visualization literature, e.g. half-angle slicing [17] which empirically computes forward scattering for volume visualization. Sometimes precomputation is used to speed up the rendering of heterogeneous translucent objects [35,37] or smoke using compensated ray marching [39]. In contrast, we target general multiple scattering in participating media without any precomputation or focus on a particular phenomenon.

We extend the work of Elek et al. [5], building primarily on the concept of DOM [3] and the more recent light propagation volumes [15,1]. These approaches are attractive for interactive applications as their grid-based local propagation schemes allow for easy parallel implementation on contemporary GPUs. Our work also shares similarities with the finite-difference time domain method [25], however we only consider the radiance amplitude and in general concentrate on efficiency.

Virtually all existing variants or extensions of DOM use a single scene-aligned propagation grid, where every cell stores a representation of the directional radiance function using spherical harmonics (SH) or piecewise-constant functions. This representation is then used to iteratively calculate energy transfer between nearby cells, typically within a local 18- or 26-neighborhood. However, this representation is only suited for moderately anisotropic scattering at best – especially for anisotropic media under complex (high-frequency) illumination such approach causes prominent ray effects and false scattering artifacts (see [9]). We take a different approach and propose to identify the most important light propagation directions (principal ordinates) in the scene and then use *multiple propagation grids* aligned with these directions, instead of a single one. This enables using a unimodal representation of the angular energy distribution around the principal direction in each grid cell.

3. Principal-Ordinates Propagation

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

Furthermore, we use only a single scalar value per grid cell to describe the local anisotropy of the directional light distribution. In our scheme, we use the well-known Henyey–Greenstein (HG) [11] distribution; the aforementioned value, called the *anisotropy coefficient*, is used to parametrize this distribution. Using principal directions implies that for more complex lighting scenarios we have to use multiple grids that sufficiently well approximate their directionality; for local light sources we propose to use spherical grids centered around them.

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

Download English Version:

https://daneshyari.com/en/article/441465

Download Persian Version:

https://daneshyari.com/article/441465

Daneshyari.com