

Radiance Caching for Participating Media

WOJCIECH JAROSZ, CRAIG DONNER, MATTHIAS ZWICKER, and HENRIK WANN JENSEN

University of California, San Diego

In this article we present a novel radiance caching method for efficiently rendering participating media using Monte Carlo ray tracing. Our method handles all types of light scattering including anisotropic scattering, and it works in both homogeneous and heterogeneous media. A key contribution in the article is a technique for computing gradients of radiance evaluated in participating media. These gradients take the full path of the scattered light into account including the changing properties of the medium in the case of heterogeneous media. The gradients can be computed simultaneously with the inscattered radiance with negligible overhead. We compute gradients for single scattering from lights and surfaces and for multiple scattering, and we use a spherical harmonics representation in media with anisotropic scattering. Our second contribution is a new radiance caching scheme for participating media. This caching scheme uses the information in the radiance gradients to sparsely sample as well as interpolate radiance within the medium utilizing a novel, perceptually based error metric. Our method provides several orders of magnitude speedup compared to path tracing and produces higher quality results than volumetric photon mapping. Furthermore, it is view-driven and well suited for large scenes where methods such as photon mapping become costly.

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

General Terms: Algorithms, Performance, Theory

Additional Key Words and Phrases: Global illumination, gradients, irradiance caching, Monte Carlo ray tracing, participating media, ray marching, rendering, spherical harmonics

ACM Reference Format:

Jaroszw, W., Donner, C., Zwicker, M., and Jensen, H. W. 2008. Radiance caching for participating media. *ACM Trans. Graph.* 27, 1, Article 7 (March 2008), 11 pages. DOI = 10.1145/1330511.1330518 <http://doi.acm.org/10.1145/1330511.1330518>

1. INTRODUCTION

Rendering participating media, such as clouds, fire, water, or the atmosphere, has been studied extensively in the computer graphics community. We refer to the excellent survey by Cerezo et al. [2005] for a comprehensive overview of the state of the art in this field. Pioneering contributions include the work of Kajiyama and von Herzen [1984] and Rushmeier and Torrance [1987].

The most general techniques are based on various forms of stochastic path tracing and Monte Carlo integration [Pattanaik and Mudur 1993; Lafortune and Willems 1996; Pauly et al. 2000]. Path tracing is attractive because it is straightforward to include anisotropic phase functions and scattering from surfaces, but it suffers from noise and it is very computationally intensive for all but the simplest scenes.

Several recent techniques make particular assumptions about the medium being rendered in order to gain efficiency. Stam [1995] and Jensen et al. [2001] assume the material is homogeneous and has high scattering albedo, while Premoze et al. [2004] assume it is tenuous and strongly forward scattering. Sun et al. [2005] render single scattering in real time, but without shadowing effects. Our method makes no assumptions about the properties of the medium being rendered. It handles isotropic, anisotropic, homogeneous and heterogeneous media of arbitrary albedo.

Photon maps [Jensen and Christensen 1998], at the cost of introducing bias, provide an elegant solution that improves the efficiency of path tracing. However, photon mapping can be inefficient when large parts of the volume do not contribute much to the rendered view. The car scene shown in Figure 10 is particularly problematic for photon mapping as it shows a dense participating medium with a large extent and many light sources. Our caching approach is similar in spirit to volumetric photon mapping in that we store and reuse illumination information in the volume. However, our method is view-driven and concentrates samples where the illumination is most important to the eye.

Irradiance caching [Ward et al. 1988] is a view-driven technique for computing global illumination on surfaces. It works by storing a sparse set of samples of the incident illumination on surfaces in the scene, and interpolating between these samples. This idea is improved by including gradient information at the sample points to reduce interpolation artifacts [Ward and Heckbert 1992; Arvo 1994]. Krivanek et al. [2005] and Annen et al. [2004] have also used spherical harmonics to represent the full incident hemisphere of radiance as well as the gradient.

Computing the gradient of incident or reflected illumination is of interest for applications other than irradiance caching. Holzschuch and Sillion [1995, 1998] derive radiosity gradients for error estimation in finite element methods for global illumination. Igehy [1999]

This work was supported in part by NSF Grant CPA 0701992 and by the UCSD FWGrid Project, NSF Research Infrastructure Grant EIA-0303622.

Author's address: W. Jarosz; email: wjarosz@ucsd.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2008 ACM 0730-0301/2008/03-ART7 \$5.00 DOI 10.1145/1330511.1330518 <http://doi.acm.org/10.1145/1330511.1330518>

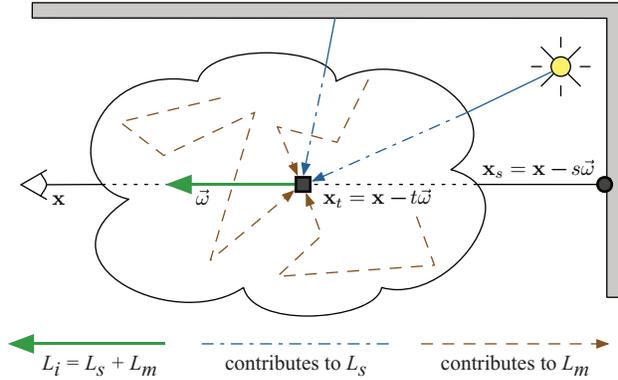


Fig. 1. Radiance in participating media is computed by our method using a combination of ray marching and random walk sampling. We split the computation of in-scattered radiance into single scattering, L_s , and multiple scattering, L_m , terms.

derives ray differentials for antialiasing specular reflections. Durand et al. [2005] and Ramamoorthi et al. [2007] develop general frameworks to analyze local change in reflected radiance, which may have wide applicability for practical problems. However, none of these techniques take into account participating media, which is the focus of our work.

In this article, we present a method for computing, caching, and reusing estimates of the radiance within a participating medium. To estimate the contribution of a sample location, we derive expressions for the gradients of the full path of in-scattered light within a medium including the effect of changing properties of the medium in the case of heterogeneous media. Our radiance caching scheme uses the information in the radiance gradients to sparsely sample and interpolate radiance within the medium. To interpolate radiance we use exponential extrapolation, which is superior to simple linear extrapolation due to the exponential attenuation within participating media. We demonstrate the caching scheme in a Monte Carlo ray tracing system for scenes containing homogeneous and heterogeneous media as well as isotropic and anisotropic scattering.

2. RADIATIVE TRANSFER IN PARTICIPATING MEDIA

Light transport in participating media is described by the Radiative Transport Equation (RTE) [Chandrasekar 1960]. The integral form of the RTE for a nonemissive medium is:

$$L(\mathbf{x}, \vec{\omega}) = \int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega}) dt + T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega}), \quad (1)$$

where $L(\mathbf{x}, \vec{\omega})$ describes the radiance that reaches \mathbf{x} from direction $\vec{\omega}$, s is the depth of the medium to \mathbf{x} from direction $\vec{\omega}$, $\mathbf{x}_t = \mathbf{x} - t\vec{\omega}$ with $t \in (0, s)$, and $\mathbf{x}_s = \mathbf{x} - s\vec{\omega}$ is a point on a surface past the medium (see Figure 1). The transmittance, T_r , is defined as

$$T_r(\mathbf{x}' \leftrightarrow \mathbf{x}) = e^{-\tau(\mathbf{x}' \leftrightarrow \mathbf{x})}, \quad (2)$$

where the optical thickness, τ , is given by:

$$\tau(\mathbf{x}' \leftrightarrow \mathbf{x}) = \int_{\mathbf{x}'}^{\mathbf{x}} \sigma_t(x) dx. \quad (3)$$

We denote the scattering and absorption coefficients by σ_s and σ_a , and $\sigma_t = \sigma_a + \sigma_s$ is the extinction coefficient.

The first term in Equation (1) represents radiance scattered within the medium and the second term is radiance entering at the backside of the medium. The in-scattered radiance is given by

$$L_i(\mathbf{x}, \vec{\omega}) = \int_{\Omega_{4\pi}} p(\mathbf{x}, \vec{\omega}', \vec{\omega}) L(\mathbf{x}, \vec{\omega}') d\vec{\omega}', \quad (4)$$

where p is the normalized phase function describing the angular distribution of scattered radiance at a point \mathbf{x} . We ignore the radiance emitted by the volume as it is trivial to compute.

To compute the in-scattered radiance within participating media we use Monte Carlo ray tracing based on a combination of ray marching and random walk sampling. We use ray marching to numerically integrate the radiance value seen directly by the observer,

$$L(\mathbf{x}, \vec{\omega}) \approx \left(\sum_{t=0}^{S-1} T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega}) \Delta_t \right) + T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega}), \quad (5)$$

where Δ_t is the length of each segment along the ray and $\mathbf{x}_0, \dots, \mathbf{x}_s$ are the sample points for each segment (\mathbf{x}_0 is the point where the ray enters the medium and \mathbf{x}_s is a point on a surface past the medium).

The most expensive part to compute in Equation (5) is the in-scattered radiance L_i , because it involves integrating over all paths that carry radiance from any other point in the scene to the current sample point along the eye ray. Therefore, we focus on the efficient computation of this term. To make the computation more practical, we divide the in-scattered radiance into two components,

$$L_i(\mathbf{x}, \vec{\omega}) = L_s(\mathbf{x}, \vec{\omega}) + L_m(\mathbf{x}, \vec{\omega}), \quad (6)$$

where L_s represents single and L_m multiple scattering respectively.

The *single scattering term* represents radiance that undergoes a single scattering event along its path from a surface to the eye. Surfaces include light sources and any other surface that reflects light due to direct or indirect illumination. Single scattering may be formulated as an integration over the sphere, or over all surfaces.

The *multiple scattering term* integrates radiance that scatters at least once in the medium before it reaches \mathbf{x} . This implies that the path from the surface to the eye includes multiple scattering events. Incident radiance due to multiple scattering is more difficult to evaluate than single scattering, as it leads to a recursive integral.

The distribution of in-scattered radiance is often smooth in large parts of the participating medium. We take advantage of this property by *caching* radiance at a sparse set of locations, and using extrapolation to evaluate the radiance at nearby points in the medium. We show that results without visible errors are achievable with relatively few cached values, resulting in large speedups compared to Monte Carlo ray tracing without caching. Furthermore, the use of caching does not limit the type of media that can be rendered, and our method can easily be incorporated into existing Monte Carlo ray tracers.

We improve the efficiency of the cache by not only storing the value of scattered radiance at each cache point, but also storing its gradient with respect to translation of the point. The gradient is used to estimate the local smoothness and to improve the accuracy of extrapolation of cached values. In the following, we explain two separate methods: one to compute the single scattering contribution and its gradient (Section 3), and another to compute the multiple scattering contribution and its gradient (Section 4). Finally, we detail how to use these quantities in a practical radiance caching algorithm for participating media in Section 5.

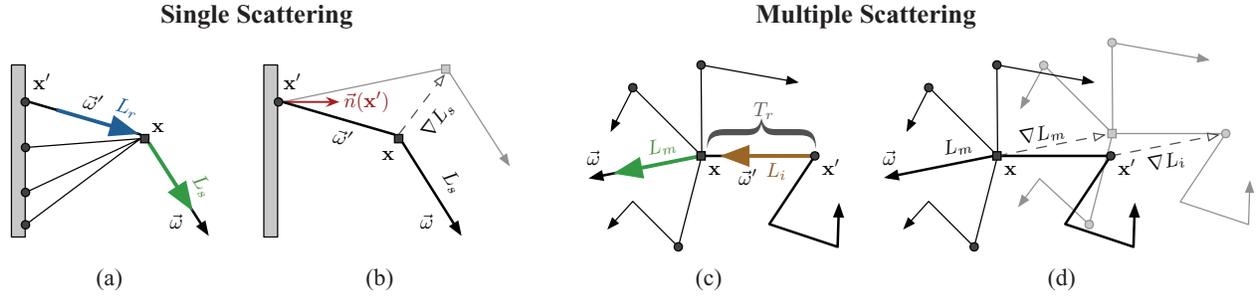


Fig. 2. Computing the single scattering radiance and gradient (a, b) and multiple scattering radiance and gradient (c, d). For single scattering we distribute samples \mathbf{x}' on the area of the light source. The gradient ∇L_s is computed w.r.t. translating the evaluation point \mathbf{x} . For multiple scattering we distribute samples \mathbf{x}' by generating random walk-paths starting at \mathbf{x} . The gradient ∇L_m is computed by considering the translation of \mathbf{x} and the paths *as a whole*.

3. SINGLE SCATTERING

In this section, we describe how to compute the single scattering radiance L_s at a position \mathbf{x} and its gradient ∇L_s with respect to \mathbf{x} . The derivation of the gradient is more convenient if we write L_s as an integral over the surfaces of the scene, instead of the standard integral over the sphere:

$$L_s(\mathbf{x}, \vec{\omega}) = \int_A p(\mathbf{x}, \vec{\omega}', \vec{\omega}) L_r(\mathbf{x}' \rightarrow \mathbf{x}) V(\mathbf{x}' \rightarrow \mathbf{x}) H(\mathbf{x}' \rightarrow \mathbf{x}) d\mathbf{x}', \quad (7)$$

where A denotes surface area and $\vec{\omega}'$ points from \mathbf{x}' towards \mathbf{x} . We illustrate our notation in Figure 2a. The term $L_r(\mathbf{x}' \rightarrow \mathbf{x})$ represents the *reduced radiance*. It is given by the product of the radiance leaving a surface at \mathbf{x}' towards \mathbf{x} and the transmittance through the medium:

$$L_r(\mathbf{x}' \rightarrow \mathbf{x}) = T_r(\mathbf{x}' \leftrightarrow \mathbf{x}) L(\mathbf{x}' \rightarrow \mathbf{x}). \quad (8)$$

We also include a *visibility function* $V(\mathbf{x}' \rightarrow \mathbf{x})$, whose value is unity if \mathbf{x}' is visible from \mathbf{x} , otherwise it is zero. Finally, H is a *geometry term*,

$$H(\mathbf{x}' \rightarrow \mathbf{x}) = \frac{\vec{\mathbf{n}}(\mathbf{x}') \cdot \frac{\mathbf{x} - \mathbf{x}'}{\|\mathbf{x} - \mathbf{x}'\|}}{\|\mathbf{x} - \mathbf{x}'\|^2}, \quad (9)$$

due to the change of the integration variable. Note that negative values of the dot product must be clamped to zero, but we omit this in our notation.

The single scattering gradient follows directly from Equation (7),

$$\begin{aligned} \nabla L_s(\mathbf{x}, \vec{\omega}) &= \int_A \nabla(p L_r V H) d\mathbf{x}' \\ &= \int_A (\nabla p) L_r V H + p(\nabla L_r) V H \\ &\quad + p L_r (\nabla V) H + p L_r V (\nabla H) d\mathbf{x}', \end{aligned} \quad (10)$$

where we omit the function arguments for brevity. The symbol ∇ always denotes a gradient with respect to \mathbf{x} . Intuitively, the gradient corresponds to the direction of maximum change of the scattered radiance at \mathbf{x} under an infinitesimal motion of \mathbf{x} while keeping \mathbf{x}' fixed. This is illustrated in Figure 2b.

In the remainder of this section we examine the single scattering computation in more detail. In Section 3.1 we discuss how to estimate the single scattered radiance and gradient using Monte Carlo

integration and describe how to incorporate point light sources in Section 3.2. Finally, we discuss the gradient of the reduced radiance term in more detail in Section 3.3, while derivations for the gradients of the geometry term H , and the phase function p are in the Appendix.

3.1 Monte Carlo Integration

We compute single scattering radiance values (Equation (7)) using Monte Carlo ray tracing, which leads to the following formula:

$$L_s(\mathbf{x}, \vec{\omega}) \approx \frac{1}{N} \sum_{j=1}^N \frac{p L_r V H}{pdf(\mathbf{x}'_j)}. \quad (11)$$

Similarly, we use Monte Carlo estimation to evaluate the gradient of single scattering radiance from Equation (10):

$$\nabla L_s(\mathbf{x}, \vec{\omega}) \approx \frac{1}{N} \sum_{j=1}^N \frac{(\nabla p) L_r V H + p(\nabla L_r) V H + p L_r V (\nabla H)}{pdf(\mathbf{x}'_j)}. \quad (12)$$

The radiance and the gradient share many of the same terms and are evaluated simultaneously.

Because the visibility function V is discontinuous, its “gradient” is not taken into account. For the purposes of gradient computation we currently assume constant visibility and discuss the implications of this decision in Section 7.

Monte Carlo approximation relies on a set of samples \mathbf{x}'_j that are distributed on the surfaces of the scene according to a probability density function $pdf(\mathbf{x}')$. Surfaces include both area light sources that emit light, and any other surface in the scene that may reflect light.

To gather radiance emitted by area light sources we place samples on their area. Our method allows the use of arbitrary light source sampling techniques to optimize the sample distribution. To account for radiance emitted by all other surfaces we distribute samples uniformly over the sphere of directions. Using the standard solid angle to area conversion, uniform sampling on the sphere corresponds to a distribution

$$pdf(\mathbf{x}') = \frac{\vec{\mathbf{n}}(\mathbf{x}') \cdot \frac{\mathbf{x}_0 - \mathbf{x}'}{\|\mathbf{x}_0 - \mathbf{x}'\|}}{\|\mathbf{x}_0 - \mathbf{x}'\|^2}, \quad (13)$$

where $\vec{\mathbf{n}}(\mathbf{x}')$ is the surface normal at \mathbf{x}' , and \mathbf{x}_0 is the location where to compute the spherical distribution. Conceptually, \mathbf{x}_0 is independent of the location \mathbf{x} where we evaluate in-scattering. Although in

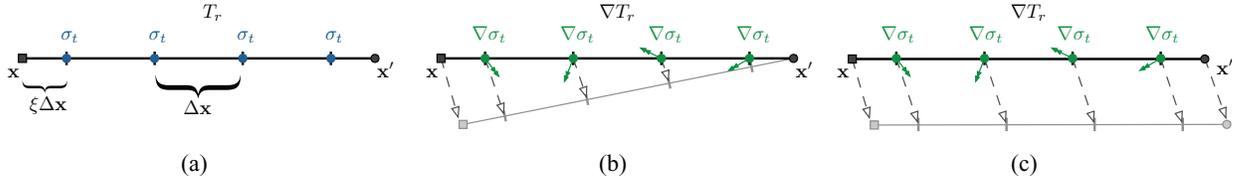


Fig. 3. In heterogeneous media, the transmittance (a) between two points \mathbf{x} and \mathbf{x}' is evaluated using ray marching with a step size Δx and a random offset ξ . Within each step the extinction coefficient of the medium is assumed to be constant. Our gradient computation not only takes into account the local properties of the medium at \mathbf{x} , but also the changing properties along the whole segment between \mathbf{x} and \mathbf{x}' . For single scattering, (b), we compute the gradient w.r.t. translation of \mathbf{x} as \mathbf{x}' is kept fixed. For multiple scattering, (c), the gradient is computed as both endpoints are translated.

practice \mathbf{x}_0 and \mathbf{x} usually coincide, $pdf(\mathbf{x}')$ does not depend on \mathbf{x} . Hence, it does not contribute to the gradient.

3.2 Point Lights

Even though Equation (7) is formulated with respect to surface area, it is straightforward to incorporate point lights. We account for the contribution of point lights to the Monte Carlo estimate of L_s by adding the following summand to Equation (11):

$$L_s(\mathbf{x}, \vec{\omega}) += p \Phi_p T_r V H_p, \quad (14)$$

where Φ_p is the power of the light source and H_p describes the light's fall-off characteristics. For a uniform point light we have

$$H_p^u(\mathbf{x}' \rightarrow \mathbf{x}) = \frac{1}{\|\mathbf{x} - \mathbf{x}'\|^2}.$$

For a Phong light the fall-off term is

$$H_p^p(\mathbf{x}' \rightarrow \mathbf{x}) = \frac{\left(\vec{\mathbf{n}}(\mathbf{x}') \cdot \frac{\mathbf{x} - \mathbf{x}'}{\|\mathbf{x} - \mathbf{x}'\|} \right)^s}{\|\mathbf{x} - \mathbf{x}'\|^2},$$

where s is the Phong exponent.

The contribution of a point light to the single scattering gradient is:

$$\nabla L_s(\mathbf{x}, \vec{\omega}) += \Phi_p ((\nabla p) T_r H_p + p (\nabla T_r) H_p + p T_r (\nabla H_p)). \quad (15)$$

The gradient ∇H_p for uniform and Phong point lights is derived in Appendix.

3.3 Reduced Radiance and Transmittance Gradient

For the sake of computing the gradient of the reduced radiance we assume diffuse surfaces and light sources. In this case the gradient of the surface radiance $L(\mathbf{x}' \rightarrow \mathbf{x})$ with respect to \mathbf{x} vanishes. However, it is straightforward to account for general radiance distributions by adding a term for ∇L . With this simplification, the gradient of the reduced radiance is

$$\nabla L_r(\mathbf{x}' \rightarrow \mathbf{x}) = L(\mathbf{x}' \rightarrow \mathbf{x}) \nabla T_r(\mathbf{x}' \leftrightarrow \mathbf{x}),$$

and the gradient of the transmittance is

$$\nabla T_r(\mathbf{x}' \leftrightarrow \mathbf{x}) = -\nabla \tau(\mathbf{x}' \leftrightarrow \mathbf{x}) T_r(\mathbf{x}' \leftrightarrow \mathbf{x}). \quad (16)$$

Homogeneous Media. If the medium is homogeneous with constant extinction coefficient σ_t , then the optical thickness is

$$\tau(\mathbf{x}' \leftrightarrow \mathbf{x}) = \sigma_t \|\mathbf{x} - \mathbf{x}'\|, \quad (17)$$

and its gradient is

$$\nabla \tau(\mathbf{x}' \leftrightarrow \mathbf{x}) = \sigma_t \nabla(\|\mathbf{x} - \mathbf{x}'\|). \quad (18)$$

Heterogeneous Media. In a heterogeneous medium, to compute optical thickness the extinction needs to be integrated as:

$$\tau(\mathbf{x}' \leftrightarrow \mathbf{x}) = \int_0^1 \sigma_t(\mathbf{y}(t)) \|\mathbf{x}' - \mathbf{x}\| dt, \quad (19)$$

where we have parameterized the line segment between \mathbf{x}' and \mathbf{x} as $\mathbf{y}(t) = \mathbf{x} + t(\mathbf{x}' - \mathbf{x})$.

The corresponding gradient is

$$\nabla \tau(\mathbf{x}' \leftrightarrow \mathbf{x}) = \int_0^1 \nabla(\sigma_t(\mathbf{y}(t)) \|\mathbf{x}' - \mathbf{x}\|) dt. \quad (20)$$

The optical thickness and its gradient can be expressed using Monte Carlo integration as

$$\tau(\mathbf{x}' \leftrightarrow \mathbf{x}) \approx \frac{1}{N} \sum_j \frac{\sigma_t(\mathbf{y}(t_j)) \|\mathbf{x}' - \mathbf{x}\|}{pdf(t_j)}, \quad (21)$$

$$\nabla \tau(\mathbf{x}' \leftrightarrow \mathbf{x}) \approx \frac{1}{N} \sum_j \frac{\nabla(\sigma_t(\mathbf{y}(t_j)) \|\mathbf{x}' - \mathbf{x}\|)}{pdf(t_j)}. \quad (22)$$

We evaluate the two Monte Carlo integrals simultaneously using ray marching with a fixed step size and a single random offset ξ (see Figure 3).

At a high level, the evaluation of transmittance aggregates σ_t along a line segment while its gradient aggregates the contribution of $\nabla \sigma_t$ along the line segment. An important characteristic of this gradient formulation is that it not only takes into account the local properties at \mathbf{x} , it also incorporates how the properties change along the whole segment between \mathbf{x} and \mathbf{x}' . This implies that our gradient computation contains meaningful information about how the transmittance changes even when moving \mathbf{x} out of the line connecting \mathbf{x} and \mathbf{x}' .

3.4 Isotropic and Anisotropic Scattering

Isotropic Media. In isotropic media the phase function is a constant $p = 1/4\pi$. Therefore, the scattered radiance and its gradient are independent of the outgoing direction $\vec{\omega}$, and all terms that include the gradient of the phase function ∇p in Equations (12) and (15) vanish.

Anisotropic Media. For anisotropic media the scattered radiance is a function of the outgoing direction $\vec{\omega}$, depending on the shape of the phase function. As phase functions are smooth, we can compute a compact representation of the directional radiance distribution. We use a spherical harmonic expansion, although other spherical functions could be used. For each sample \mathbf{x}'_j with a fixed incident direction $\vec{\omega}'_j$ we compute the spherical harmonic expansion

of the phase function

$$p(\mathbf{x}, \vec{\omega}'_j, \vec{\omega}) \approx \sum_{k=1}^M c_{j,k} y_k(\vec{\omega}), \quad (23)$$

where y_k are the spherical harmonic functions, and $c_{j,k}$ are the corresponding coefficients. We use a single index for the spherical harmonics to simplify notation.

Applying this formulation to Equation (11) gives the following expression for scattered radiance in anisotropic media:

$$L_s(\mathbf{x}, \vec{\omega}) \approx \frac{1}{N} \sum_{k=1}^M y_k \sum_{j=1}^N \frac{c_{j,k} L_r V H}{pdf(\vec{\omega}'_j)}. \quad (24)$$

We apply the same projection mechanism to the gradient computation. Derivations for the gradients of common phase functions are available in the Appendix. We denote the spherical harmonic expansion of the gradient of the phase function, as a function over outgoing angle $\vec{\omega}$, by

$$\nabla p(\mathbf{x}, \vec{\omega}'_j, \vec{\omega}) \approx \sum_{k=1}^M c'_{j,k} y_k(\vec{\omega}). \quad (25)$$

This leads to the following expression for the gradient:

$$\begin{aligned} \nabla L_s(\mathbf{x}, \vec{\omega}) \\ \approx \frac{1}{N} \sum_{k=1}^M y_k \sum_{j=1}^N \frac{c'_{j,k} L_r V H + c_{j,k} (\nabla L_r) V H + c_{j,k} L_r V (\nabla H)}{pdf(\vec{\omega}'_j)}. \end{aligned} \quad (26)$$

4. MULTIPLE SCATTERING

In this section we describe how we calculate the multiple scattering radiance L_m , and its gradient ∇L_m . We evaluate multiple scattering using standard Monte Carlo path tracing and compute a Monte Carlo estimate of the gradient of the path tracing integral.

The radiance due to multiple scattering is an integral of radiance arriving from all other points within the participating medium. We express this as

$$L_m(\mathbf{x}, \vec{\omega}) = \int_{\Omega_{4\pi}} \int_0^\infty p(\mathbf{x}, \vec{\omega}', \vec{\omega}) T_r(\mathbf{x}' \leftrightarrow \mathbf{x}) \sigma_s(\mathbf{x}') L_i(\mathbf{x}', \vec{\omega}') V(\mathbf{x}' \leftrightarrow \mathbf{x}) dr' d\vec{\omega}', \quad (27)$$

where we parametrize the integration domain using polar coordinates $(r', \vec{\omega}')$ centered at \mathbf{x} so $\mathbf{x}' = \mathbf{x} - r'\vec{\omega}'$. The $L_i(\mathbf{x}', \vec{\omega}')$ term is the in-scattered radiance at \mathbf{x}' towards \mathbf{x} and needs to be evaluated recursively. See Figure 2c for an illustration.

The gradient of multiple scattering is computed by differentiating Equation (27) with respect to \mathbf{x} :

$$\begin{aligned} \nabla L_m(\mathbf{x}, \vec{\omega}) = \int_{\Omega_{4\pi}} \int_0^\infty p(\mathbf{x}, \vec{\omega}', \vec{\omega}) \\ \nabla (T_r(\mathbf{x}' \leftrightarrow \mathbf{x}) \sigma_s(\mathbf{x}') L_i(\mathbf{x}', \vec{\omega}') V(\mathbf{x}' \leftrightarrow \mathbf{x})) dr' d\vec{\omega}'. \end{aligned}$$

Note that with the chosen parametrization the incoming and outgoing directions do not change under translation of \mathbf{x} so there is no gradient term for the phase function. In homogeneous media, the gradient of the transmittance vanishes. For the general case of heterogeneous media, the gradient of transmittance and optical thickness are computed as in Equations (16) and (20). However, in the multiple scattering case $\|\mathbf{x} - \mathbf{x}'\|$ does not change under a translation and so the $\nabla(\|\mathbf{x} - \mathbf{x}'\|)$ term vanishes. This is illustrated in Figure 3c.

4.1 Monte Carlo Integration

Just as with single scattering, we evaluate the radiance due to multiple scattering and its gradient using Monte Carlo integration. The radiance is estimated by

$$L_m(\mathbf{x}, \vec{\omega}) \approx \frac{1}{N} \sum_{j=1}^N \frac{p T_r \sigma_s L_i V}{pdf(r'_j) pdf(\vec{\omega}'_j)}. \quad (28)$$

In practice, at the first scattering event we evaluate this expression using N samples. In order to prevent exponential growth, all subsequent scattering events use only one random sample. Hence, Equation (28) forms the start of N random-walk paths through the medium with the polar coordinates $(r'_j, \vec{\omega}'_j)$ as random variables. We sample the radius according to

$$pdf(r'_j) = \sigma_t(\mathbf{x}) e^{-\sigma_t(\mathbf{x}) r'_j}. \quad (29)$$

We distinguish between isotropic and anisotropic media when sampling directions, discussed in more detail below.

The gradient $\nabla L_m(\mathbf{x}, \vec{\omega})$ is also found using Monte Carlo sampling:

$$\nabla L_m(\mathbf{x}, \vec{\omega}) \approx \frac{1}{N} \sum_{j=1}^N \frac{p ((\nabla T_r) \sigma_s L_i V + T_r (\nabla \sigma_s) L_i V + T_r \sigma_s (\nabla L_i) V)}{pdf(r'_j) pdf(\vec{\omega}'_j)}. \quad (30)$$

Intuitively, the gradient captures the change in radiance as each random-walk path is translated, as a whole, by an infinitesimal amount. This is illustrated in Figure 2d. As with single scattering, we currently ignore the change in the visibility term.

4.2 Isotropic and Anisotropic Scattering

Isotropic Media. For isotropic media we distribute the directions $\vec{\omega}'_j$ uniformly, that is, $pdf(\vec{\omega}'_j) = 1/4\pi$.

Anisotropic Media. For anisotropic media the multiple scattered radiance depends on the outgoing direction $\vec{\omega}$. As with the single scattering contribution, we use spherical harmonics to represent the phase function $p(\mathbf{x}, \vec{\omega}', \vec{\omega})$, which leads to expressions analogous to Equations (24) and (26). However, there is no term for the gradient of the phase function in multiple scattering, as the angle between path vertices does not change under a translation of the whole path.

We apply importance sampling according to the phase function at all but the first scattering event since the outgoing direction is given by the direction to the previous path vertex. However, we always sample the direction at the first scattering event uniformly. The reason is that the outgoing direction $\vec{\omega}$ is not known a priori. It will be determined only when the cache point is evaluated for a certain view direction.

5. RADIANCE CACHING ALGORITHM

Using the groundwork laid out in the previous sections we now describe a radiance caching method for participating media. Our algorithm efficiently computes lighting within participating media by sparsely sampling L_s and L_m and caching them for extrapolation. In addition to computing single and multiple scattering separately as described in Section 2, we separate our cache based on the expected frequency content of the radiance values. Specifically, we maintain distinct cache points for single scattering from lights, single scattering from surfaces, and multiple scattering. We call these *single scattering*, *surface scattering*, and *multiple scattering* cache points.

This allows our system to sample low-frequency components more sparsely.

The core algorithm for radiance caching in participating media is as follows. For each point in the volume where we need to evaluate inscattered radiance, we first query the cache for a set of usable cache points. If we find any points, then we use the cached values to extrapolate the radiance to this point as described in Section 5.3. If there are no usable points, then we compute the radiance and its gradient as described in Sections 3 and 4 and store these values in the cache. The procedure is performed for all three caches. The total inscattered radiance is the sum of the three contributions.

5.1 Cache Entry Storage

For efficient access to cache points, we use an octree. The data structure and lookup procedure is similar to previous work on surface radiance caching techniques [Ward et al. 1988].

All cache entries store the inscattered radiance and its gradients at a location in space, as well as a valid radius to later determine if a cache entry is valid for extrapolation at a given query location. Note that the inscattered radiance is the amount of light that *leaves* a point in a particular direction; thus, we must store the radiance over the full sphere of directions.

Isotropic Media. Since the inscattered radiance does not depend on direction, each cache entry stores it using three scalar values, one for each color channel. In addition, we compute and store the gradient for each channel separately for greater accuracy. The total storage requirement for each isotropic entry is 16 floating point numbers. In total, the size of an isotropic cache entry is 48 bytes, though this could easily be reduced using an RGBE [Ward 1991] or 16-bit floating point representation [Kainz et al. 2006].

Anisotropic Media. In anisotropic media, each Monte Carlo sample of the radiance has a different directional contribution that depends on the shape of the phase function. Because phase functions are generally smooth, we capture this directional dependence using spherical harmonics as described in Section 3.

Our system supports storing spherical harmonic expansions to an arbitrary number of coefficients. However, for scenes with area lights and moderately anisotropic media (e.g., a Henyey-Greenstein phase function with $|g| < 0.6$), we have found that using $M = 9$ coefficients in Equation (23) is sufficient to accurately represent the inscattered radiance. Using 9 coefficients, the full RGB representation of the inscattered radiance requires 27 coefficients. Including the cache sample location and valid radius, this gives 28 floating point values, or 112 bytes per entry assuming 32-bit values.

5.2 Valid Radius and Error Tolerance

Each cache point has a valid radius within which it is a candidate for extrapolation. To compute this radius we have developed a perceptually based error metric which adapts to the local smoothness of the radiance field.

To estimate a valid radius for cache points we define the notion of an “optimal radius.” The optimal radius for a cache point is the largest radius such that the total relative L_2 error over the extrapolated region is below some error threshold ε . More formally, if $\hat{L}_{x_0}(x)$ is the radiance extrapolated from location x_0 to x , then the

optimal radius is given by

$$R_{opt}(\mathbf{x}) = \max_r \left(\frac{\int_{\mathbf{x} \in \Omega_r} |L(\mathbf{x}) - \hat{L}_{x_0}(\mathbf{x})| d\mathbf{x}}{\int_{\mathbf{x} \in \Omega_r} L(\mathbf{x}) d\mathbf{x}} \leq \varepsilon \right). \quad (31)$$

We use the relative error because the human visual system is more sensitive to contrast than absolute differences in intensity [Glassner 1995]. Solving for $R_{opt}(\mathbf{x})$ during rendering would be impractical since it requires knowledge of the true radiance within the valid radius of any cache point. Our goal is therefore to devise an error metric which approximates the optimal radius as closely as possible using only information already available from the Monte Carlo samples at \mathbf{x} .

The ratio of the radiance gradient to the radiance is a local measure of the contrast. It is also equivalent to the log-space gradient of the radiance:

$$\nabla \ln(L) = \frac{\nabla L}{L}. \quad (32)$$

Making the valid radius inversely proportional to the magnitude of the log-space gradient at \mathbf{x} makes sense at an intuitive level: areas where the radiance field has high contrast will be sampled more densely than smoothly varying regions. Furthermore, using the log-space gradient instead of the regular gradient means that the valid radius computation is independent of the absolute intensity of the lighting in the scene.

In our 1D test cases this approach matches the behavior of the optimal radius well at many locations (see Figure 4). Unfortunately the log-space gradient is a completely local quantity and can be a bad predictor of the contrast in the surrounding area. Specifically, the magnitude of the gradient can drop to zero, causing an infinite valid radius. A zero gradient, or saddle-point, can naturally occur in many situations, for example, between two lights. Using the Monte Carlo samples, the reciprocal magnitude of the log-space gradient can be expressed as:

$$\frac{L(\mathbf{x})}{\|\nabla L(\mathbf{x})\|} = \frac{\sum_j L_j(\mathbf{x})}{\left\| \sum_j \nabla L_j(\mathbf{x}) \right\|}, \quad (33)$$

where j represents the Monte Carlo samples over light sources (for single scattering), surfaces (for surface scattering), and random-walk paths (for multiple scattering). However, instead of using the log-space gradient directly, we use a modified metric designed to eliminate saddle-points. The valid radius of cache points is computed as:

$$R(\mathbf{x}) = \varepsilon \frac{\sum_j L_j(\mathbf{x})}{\sum_j \|\nabla L_j(\mathbf{x})\|}, \quad (34)$$

where ε is a user controlled global error tolerance parameter. By summing over gradient magnitudes instead of gradient directions in the denominator we prevent the introduction of infinite radii. In our 1D test cases this error metric matches the optimal radius remarkably well (see Figure 4) and in practice produces very good cache point distributions.

For colored media, we compute the valid radius for the three color channels independently and use the minimum. In anisotropic media, we use the zero-order spherical harmonic coefficients to compute the radius.

In our implementation, the amount of error in the extrapolation is controlled by the tolerance ε . A tolerance of 0 indicates that values

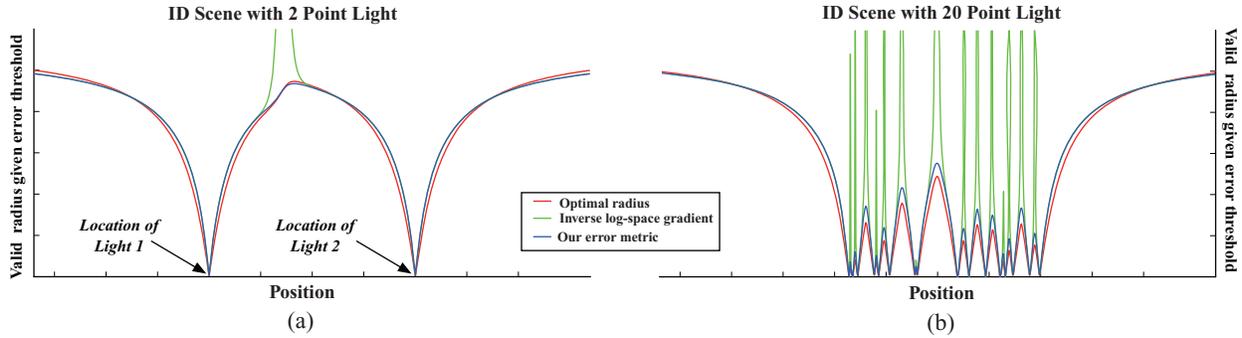


Fig. 4. Experimental validation of our error metric as compared to a numerically computed optimal radius for a 1D scene with (a) two point lights of differing strength and (b) 20 equal strength point lights. Though the log-space gradient matches the optimal radius of (a) in most situations, it produces infinite radii where the gradient vanishes. These singularities become even more prominent with many light sources (b). In both cases, our perceptually based error metric robustly approximates the optimal radius even in the presence of saddle points in the radiance field.

should never be cached or reused. Setting $\varepsilon = 1$ indicates that the extrapolated radiance will be usable until it differs by about a factor of 2 from the cached radiance value.

In areas that are occluded from all light sources, the radiance due to direct illumination and its gradient are zero. This means that the valid radius of the direct cache is undefined. To remedy this problem, we also compute an *unshadowed* gradient for the single scattering cache by assuming that all light sources are visible. We then choose the smaller of the regular and the unshadowed radii as the valid radius.

Finally, we set minimum and maximum radii for the cache points in both world and screen space, and clamp gradients according to standard techniques when using radiance caching techniques [Křivánek et al. 2006].

5.3 The Extrapolated Radiance Estimate

Previous caching techniques for surfaces use constant or linear extrapolation to approximate the radiance at a given location from a set of cached points. Constant extrapolation reuses the value at a cache entry for the new location, while linear extrapolation uses a constant slope to extrapolate the radiance with distance. In participating media, however, light intensity falls off exponentially with distance. We exploit this property in extrapolating the inscattered radiance from cache points by exponentially extrapolating the cached values. We have found that exponential extrapolation provides a smoother reconstruction of the radiance field than linear extrapolation.

Exponential Extrapolation. Exponential extrapolation is equivalent to linear extrapolation in log-space. The radiance at a point \mathbf{x} is approximated from a set of cache points C whose valid radii contain our query location by a weighted sum of extrapolated radiance values:

$$L(\mathbf{x}) \approx \exp \left(\frac{\sum_{k \in C} \left(\ln(L_k) + \frac{\nabla L_k}{L_k} \cdot (\mathbf{x} - \mathbf{x}_k) \right) w(d_k)}{\sum_{k \in C} w(d_k)} \right), \quad (35)$$

where L_k , ∇L_k , \mathbf{x}_k , and r_k are the radiance, gradient, position, and valid radius of cache point k respectively. The weighting function w is a smooth cubic, $w(d) = 3d^2 - 2d^3$. We define $d_k = 1 - \|\mathbf{x}_k - \mathbf{x}\|/r_k$ so that the weighting function has the most influence at the cache

location, and this influence falls off smoothly to zero at the valid radius.

Figure 5 compares constant, linear, and exponential extrapolation in a 2D slice of a scene containing a point source in a participating medium. The images were rendered using the same cache point locations, the only difference is the method used to extrapolate the radiance. Constant extrapolation is clearly a poor choice for extrapolation in participating media as it gives strong discontinuity artifacts from the jumps between points, while linear extrapolation does a somewhat better job at approximating the change in radiance. Exponential extrapolation most precisely matches a reference solution, with almost no visible artifacts, as it more closely approximates the real change in radiance over distance.

6. RESULTS

We implemented the volumetric radiance caching algorithm in a Monte Carlo ray tracer, and all our timings were made on an Intel Core 2 Duo 2.4 GHz machine using only one core. We performed comparisons between path tracing, radiance caching, and photon mapping. Since each rendering technique has its own set of parameters, we performed equal-time comparisons while hand-picking the best parameters for each algorithm. All three methods make full use of standard optimization techniques including Russian roulette, stratification, and importance sampling where applicable. All images were rendered at 1K horizontal resolution with up to 16 samples per pixel.

Figure 6 shows a series of images of a Cornell box filled with smoke. We compare our method to path tracing for a variety of participating media types. Render times range from around 1.5 minutes for homogeneous media with isotropic scattering, to 7.5 minutes in heterogeneous media with anisotropic scattering. In all four examples, path tracing exhibits significant noise for the same render time. In the heterogeneous images the scattering properties are defined procedurally using several octaves of noise [Ebert et al. 2002]. For anisotropic scattering, by projecting the full directional radiance field onto spherical harmonics we are able to correctly reuse cached values even for reflection rays on the mirror spheres. A visualization of the radiance cache point locations from the homogeneous/isotropic image is shown in Figure 7.

Figure 8 shows a frame from an animation moving through light beams in the Sponza atrium. In this scene all lighting is provided using a single point light, so we do not cache single-scattering but

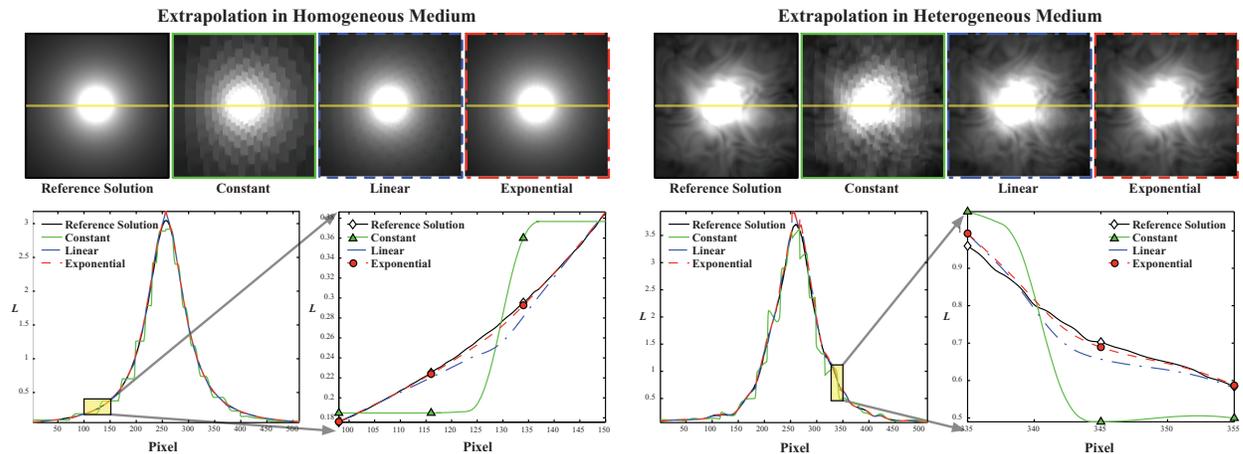


Fig. 5. A comparison of extrapolation methods for a point-light scene in both a homogeneous (left) and heterogeneous (right) medium. The top row shows renderings using three different extrapolation methods. The scanlines highlighted in yellow are plotted in the graphs below. Our proposed exponential extrapolation scheme reconstructs the sampled function more accurately than the other methods.

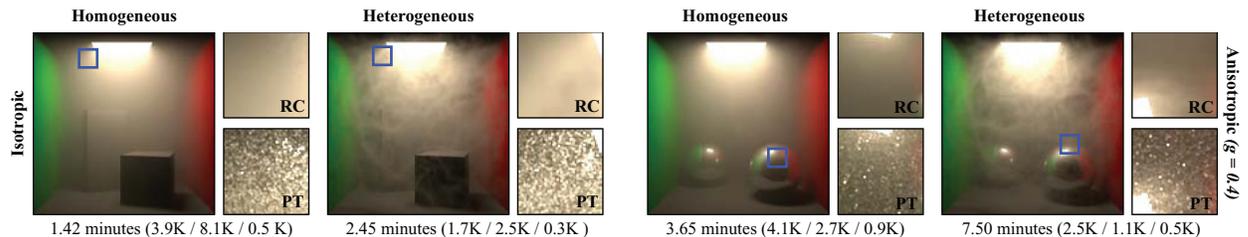


Fig. 6. A series of smoke-filled Cornell boxes showing our method’s ability to handle isotropic (left) and anisotropic (right) scattering in homogeneous or heterogeneous media. All images were rendered at $1K \times 1K$ with up to 16 samples per pixel and include single scattering from lights, single scattering from surfaces, and multiple scattering. Render times are reported underneath each image with the number of cache points in the single, surface, and multiple scattering caches in parenthesis. The zoomed insets provide equal-time comparison to path tracing. A visualization of the cache points used to render the far left image is shown in Figure 7.

compute it directly. Computing multiple scattering is particularly challenging in this scene and requires tracing thousands of paths to estimate the inscattered radiance. Our caching method is able to produce a noise-free result in 19 minutes. For walk-through style animations where only the camera position changes, we are able to reuse the radiance cache for subsequent frames, achieving an even greater speedup. Sponza is an extremely difficult scene for photon mapping. We tried performing an equal-time comparison to photon mapping. However, even by using projection maps, tracing photons for over 40 minutes produced less than 1K usable photons.

Figure 9 shows a frame from an animation of evolving smoke which took less than 6 minutes to generate. In this scene each frame is rendered using a separate cache. Our method is able to produce flicker-free animations even when the cache is recomputed at each frame. With the same render time the path tracing solution exhibits visible noise, which flickers significantly when animated. For this scene, we store about 9K cache points.

Figure 10 shows two cars on a foggy road. This scene has a large extent, and is costly to render with methods such as photon mapping. Here, our caching method fully exploits the smoothness in the radiance across the scene. In the same render time we are able to attain significant noise reduction compared to path tracing and a smoother reconstruction than photon mapping using over 8 million photons.

For all example scenes in this paper, a large portion of render time is spent querying for cache points within the octree (typically between one third and one half of total render time). This is due to the large number of queries performed while ray marching through the medium (in the millions) compared to the number of cache points created (in the thousands). However, for the multiple scattering computation in our example scenes we trace between 2K and 8K random-walk paths. Due to this, it is important to note that computing a new cache point takes about 100 to 1000 times longer than querying within the octree.

7. DISCUSSION AND FUTURE WORK

Many promising avenues remain for combining the strengths of photon mapping and radiance caching. Since photon mapping does not attempt to accelerate the computation of single scattering, radiance caching could effectively accelerate this computation within a photon mapping renderer. Another promising extension is to use our caching technique as a final gather pass for photon mapping. Currently, in highly scattering scenes with many occluders, a large number of multiple scattering rays are necessary using our method. The density estimation in photon mapping provides a natural smoothing, which could reduce the number of multiple scattering rays

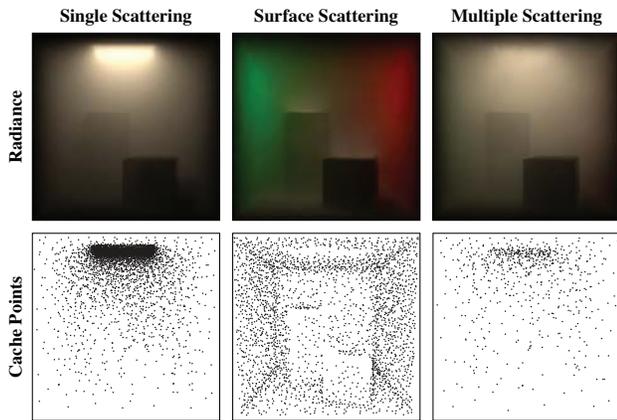


Fig. 7. The bottom row shows a visualization of the single, surface, and multiple scattering cache points using the radiance cache from the homogeneous/ isotropic image in Figure 6. The top row shows the corresponding extrapolated radiance components (the radiance from surfaces past the medium is not included). For display purposes the multiple and surface scattering images have been raised by 1 and 2 F-stops respectively.

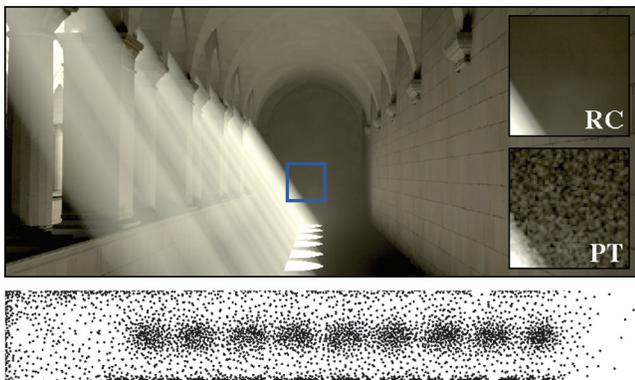


Fig. 8. The Sponza atrium (top) with beams of light and multiple scattering. With volumetric radiance caching, this scene renders in 19 minutes using about 46K cache points and 79M cache queries. Sponza is notoriously difficult for photon mapping: even using projection maps to direct photons towards the windows, tracing photons for over 40 minutes results in less than 1K usable photons. The zoomed insets on the right provide an equal-time comparison to path tracing. The bottom image visualizes the surface cache using a horizontal cross section near the floor of the hallway. The camera is located on the right looking towards the left. Note the concentration of cache points near the brightly illuminated regions of the floor. The radiance caching image was rendered at a horizontal resolution of 1K with up to 16 samples per pixel with error tolerances set at 0.5 and 0.01 for the surface and multiple scattering caches respectively. For multiple scattering, our method traces 512 random-walk paths per cache point.

necessary and improve the cached result. Furthermore, photon mapping is able to efficiently simulate multiple scattering in a variety of scenes because it is well suited for detecting light paths difficult to sample from the eye. However, for smooth reconstruction a large number of photons are required. Performing a final gather would also reduce the necessary number of stored photons and improve the reconstruction quality as compared to regular photon

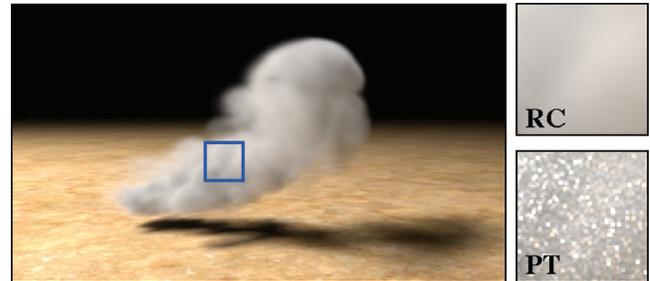


Fig. 9. A still frame from an animation of heterogeneous smoke renders in 5.8 minutes using 9K cache points and 8.2M cache queries. Our method is able to produce flicker-free animation even when recomputing a new cache for each frame. The highlighted region is shown zoomed-in on the right providing an equal-time comparison to path tracing. Error tolerances of 0.05, 1.0, and 1.0 were used for the single, surface, and multiple scattering caches with 2K multiple scattering paths per cache point.

mapping. For isotropic media, final gather optimizations developed for surfaces [Christensen 1999] could be extended for volumetric final gather. Information from the photon map could also be used in other ways, for instance, to provide an importance sampling function for computing the direct and indirect inscattered illumination.

One of the main challenges for our technique stems from ignoring the visibility gradient. Due to this simplification, our error metric is unable to adapt sample density at volumetric shadow boundaries, forcing the user to set a lower global error tolerance to produce artifact-free renderings. Incorporating the visibility gradient, both for extrapolation and in the error metric, could significantly increase quality while reducing render times. We plan to address this limitation in future work.

Caching the radiance and gradient computation using spherical harmonics is effective for moderately anisotropic scattering. For highly directional scattering, this approach becomes less beneficial. Since the cache points can potentially be reused from any viewing direction our error metric does not account for the outgoing direction when computing the valid radius. An alternate approach could be to store radiance values in a 5D cache with associated outgoing directions. Since in this case cache points would only be reused for similar outgoing directions, importance sampling could be used even at the first bounce for multiple and surface scattering. Such an approach would benefit significantly from a full 5D error metric that computes a valid radius independently for the spatial and angular dimensions.

Although our method caches the results of integrating single and multiple scattering, other expensive quantities could also be cached to further improve performance. For scenes with a large visible extent and large image resolutions, cache queries dominate the render time due to the repeated lookups necessary for integrating radiance along eye rays. Caching the radiance integration along eye rays could provide for efficient reuse of the *whole* radiance computation from neighboring pixels, completely eliminating the eye integration for a majority of the image.

Finally, the original surface irradiance caching algorithm could be extended to account for the presence of participating media, and coupled with our method. Since, however, there is no distance to surfaces associated with volumetric cache points, this would be a challenging undertaking. Also, the irradiance gradients would have to account for the change in optical properties along a path as the point moves on the surface, similar to our scattering gradients.



Fig. 10. Two cars in a dense fog on a road illuminated by 60 lights. Our radiance caching method renders the cars scene (a) in 20.35 minutes using only 175K cache points and 27M cache queries. The multiple scattering computation traces 4K random-walk paths per cache point. The highlighted region is shown zoomed-in providing an equal-time quality comparison between radiance caching (RC) and path tracing (PT). This scene has a large extent and is difficult to render with methods such as photon mapping that compute scattering in the entire medium. In a contrast-enhanced, equal-time comparison to photon mapping (b), artifacts remain in the *homogeneous* medium even after tracing 8M photons due to the large scene extent and many light sources. A horizontal cross section of the single scattering cache points (c) as see from above. In this visualization the camera is located at the bottom looking up. Note the concentration of cache points around the street lights and headlights near the camera. The images were rendered with a horizontal resolution of 1K with up to 16 samples per pixel and the error tolerances were set at 0.25 and 0.75 for the single and multiple scattering caches respectively.

8. CONCLUSION

In this paper we have presented a new method for caching radiance in the presence of participating media. The method is general and adapts to scenes with arbitrary scattering and absorbing properties. It works directly in both homogeneous and heterogeneous media, and can handle both isotropic or anisotropic scattering. When rendering anisotropic media, we compactly store the directional distribution of the radiance using spherical harmonics. Our method caches both single and multiple scattered radiance from light sources and surfaces, while providing significant gains in speed over traditional Monte Carlo ray tracing techniques.

Volumetric radiance caching is founded on using the gradients of a Monte Carlo formulation of radiative transport to accurately predict the change in radiance throughout the scene. We have shown how to compute these gradients, and also how to use them to extrapolate the radiance to new locations. This, combined with an exponential extrapolation technique that exploits the properties of light in a medium, allow us to render high quality images of scenes with a relatively small number of cache points. Caches produced by our method handle complex illumination and can be later re-used for animations.

APPENDIX: Single Scattering Gradients

Geometry Term. The gradients of the various geometry terms are

$$\begin{aligned} \nabla H(\mathbf{x}' \rightarrow \mathbf{x}) &= \nabla \left(\bar{\mathbf{n}} \cdot \frac{\mathbf{x} - \mathbf{x}'}{\|\mathbf{x} - \mathbf{x}'\|} \right) \frac{1}{\|\mathbf{x} - \mathbf{x}'\|^2} \\ &\quad + \bar{\mathbf{n}} \cdot \frac{\mathbf{x} - \mathbf{x}'}{\|\mathbf{x} - \mathbf{x}'\|} \nabla \left(\frac{1}{\|\mathbf{x} - \mathbf{x}'\|^2} \right) \end{aligned} \quad (36)$$

$$\nabla H_p^u(\mathbf{x}' \rightarrow \mathbf{x}) = \nabla \left(\frac{1}{\|\mathbf{x} - \mathbf{x}'\|^2} \right) \quad (37)$$

$$\begin{aligned} \nabla H_p^p(\mathbf{x}' \rightarrow \mathbf{x}) &= s \left(\bar{\mathbf{n}} \cdot \frac{\mathbf{x} - \mathbf{x}'}{\|\mathbf{x} - \mathbf{x}'\|} \right)^{s-1} \nabla \left(\bar{\mathbf{n}} \cdot \frac{\mathbf{x} - \mathbf{x}'}{\|\mathbf{x} - \mathbf{x}'\|} \right) \frac{1}{\|\mathbf{x} - \mathbf{x}'\|^2} \\ &\quad + \left(\bar{\mathbf{n}} \cdot \frac{\mathbf{x} - \mathbf{x}'}{\|\mathbf{x} - \mathbf{x}'\|} \right)^s \nabla \left(\frac{1}{\|\mathbf{x} - \mathbf{x}'\|^2} \right) \end{aligned} \quad (38)$$

where:

$$\nabla \left(\bar{\mathbf{n}} \cdot \frac{\mathbf{x} - \mathbf{x}'}{\|\mathbf{x} - \mathbf{x}'\|} \right) = \frac{\bar{\mathbf{n}}}{\|\mathbf{x} - \mathbf{x}'\|} - \frac{(\bar{\mathbf{n}} \cdot (\mathbf{x} - \mathbf{x}'))(\mathbf{x} - \mathbf{x}')}{\|\mathbf{x} - \mathbf{x}'\|^3} \quad (39)$$

$$\nabla \left(\frac{1}{\|\mathbf{x} - \mathbf{x}'\|^2} \right) = \frac{-2(\mathbf{x} - \mathbf{x}')}{\|\mathbf{x} - \mathbf{x}'\|^4}. \quad (40)$$

Phase Function. In isotropic media the phase function is a constant and its gradient vanishes. For anisotropic media, we focus on the Henyey-Greenstein phase function p_{HG} and the Schlick phase function p_s :

$$p_{HG}(\theta) = \frac{1 - g^2}{4\pi(1 + g^2 - 2g \cos \theta)^{1.5}} \quad (41)$$

$$p_s(\theta) = \frac{1 - k^2}{4\pi(1 + k \cos \theta)^2}. \quad (42)$$

These phase functions are 1D, only depending on the cosine of the angle between the incident and outgoing directions. We express $\cos(\theta)$ in terms of $\vec{\omega}$, \mathbf{x} , and \mathbf{x}' as:

$$\cos(\theta) = \vec{\omega} \cdot \frac{\mathbf{x} - \mathbf{x}'}{\|\mathbf{x} - \mathbf{x}'\|}. \quad (43)$$

The gradients of the phase functions with respect to \mathbf{x} are then

$$\nabla p_{HG}(\theta) = \frac{3g(1 - g^2)}{4\pi(1 + g^2 - 2g \cos \theta)^{2.5}} \nabla \cos \theta \quad (44)$$

$$\nabla p_s(\theta) = \frac{k(1 - k^2)}{2\pi(1 + k \cos \theta)^3} \nabla \cos \theta, \quad (45)$$

where

$$\nabla \cos(\theta) = \cos \theta \frac{\mathbf{x} - \mathbf{x}'}{\|\mathbf{x} - \mathbf{x}'\|^2} - \frac{\vec{\omega}}{\|\mathbf{x} - \mathbf{x}'\|}.$$

REFERENCES

- ANNEN, T., KAUTZ, J., DURAND, F., AND SEIDEL, H.-P. 2004. Spherical harmonic gradients for mid-range illumination. In *Proceedings of the Eurographics Symposium on Rendering*.
- ARVO, J. 1994. The irradiance jacobian for partially occluded polyhedral sources. In *Proceedings of SIGGRAPH*.

- CEREZO, E., PÉREZ, F., PUEYO, X., SERON, F. J., AND COIS X. SILLION, F. 2005. A survey on participating media rendering techniques. *Visu. Comput.* 21, 5 (June), 303–328.
- CHANDRASEKAR, S. 1960. *Radiative Transfer*. Dover Publications, New York, NY.
- CHRISTENSEN, P. H. 1999. Faster photon map global illumination. *J. Graph. Tools* 4, 3, 1–10.
- DURAND, F., HOLZSCHUCH, N., SOLER, C., CHAN, E., AND SILLION, F. 2005. A frequency analysis of light transport. *ACM Trans. Graph.* 24, 3 (Aug).
- EBERT, D. S., MUSGRAVE, K. F., PEACHEY, D., PERLIN, K., AND WORLEY, S. 2002. *Texturing & Modeling: A Procedural Approach, 3rd Ed.* (The Morgan Kaufmann Series in Computer Graphics). Morgan Kaufmann.
- GLASSNER, A. S. 1995. *Principles of Digital Image Synthesis*. Morgan Kaufmann.
- HOLZSCHUCH, N. AND SILLION, F. 1995. Accurate computation of the radiosity gradient with constant and linear emitters. In *6th Eurographics Workshop on Rendering*. Dublin, Ireland.
- HOLZSCHUCH, N. AND SILLION, F. 1998. An exhaustive error-bounding algorithm for hierarchical radiosity. *Comput. Graph. For.* 17, 4.
- IGEHY, H. 1999. Tracing ray differentials. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*. ACM Press, 179–186.
- JENSEN, H. W. AND CHRISTENSEN, P. H. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*. ACM Press, 311–320.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*. ACM Press, 511–518.
- KAINZ, F., JAROSZ, W., AND BOGART, R. 2006. *Technical Introduction to OpenEXR*. Lucas Digital Ltd. LLC. <http://www.openexr.com>.
- KAJIYA, J. T. AND HERZEN, B. P. V. 1984. Ray tracing volume densities. In *Proceedings of the 11th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*. ACM Press, 165–174.
- KŘIVÁNEK, J., BOUATOUCH, K., PATTANAİK, S. N., AND ŽÁRA, J. 2006. Making radiance and irradiance caching practical: Adaptive caching and neighbor clamping. In *Eurographics Symposium on Rendering*, T. Akenine-Möller and W. Heidrich, Eds. Eurographics Association, Nicosia, Cyprus.
- KŘIVÁNEK, J., GAUTRON, P., PATTANAİK, S., AND BOUATOUCH, K. 2005. Radiance caching for efficient global illumination computation. *IEEE Trans. Visualiz. Comput. Graph.* 11, 5, 550–561.
- LAFORTUNE, E. P. AND WILLEMS, Y. D. 1996. Rendering participating media with bidirectional path tracing. In *Proceedings of the Eurographics Workshop on Rendering Techniques*. Springer-Verlag, 91–100.
- PATTANAİK, S. N. AND MUDUR, S. P. 1993. Computation of global illumination in a participating medium by Monte Carlo simulation. *J. Visualiz. Comput. Anim.* 4, 3 (July–Sept.), 133–152.
- PAULY, M., KOLLIG, T., AND KELLER, A. 2000. Metropolis light transport for participating media. In *Proceedings of the Eurographics Workshop on Rendering Techniques*. Springer-Verlag, 11–22.
- PREMOZE, S., ASHIKHMIN, M., RAMAMOORTHY, R., AND NAYAR, S. K. 2004. Practical rendering of multiple scattering effects in participating media. In *Rendering Techniques*. 363–373.
- RAMAMOORTHY, R., MAHAJAN, D., AND BELHUMEUR, P. 2007. A First Order Analysis of Lighting, Shading, and Shadows. *ACM Trans. Graph.* 26, 1 (Jan).
- RUSHMEIER, H. E. AND TORRANCE, K. E. 1987. The zonal method for calculating light intensities in the presence of a participating medium. *Proceedings of (SIGGRAPH) Computer Graphics* 21, 4 (July), 293–302.
- STAM, J. 1995. Multiple scattering as a diffusion process. In *Proceedings of the 6th Eurographics Workshop on Rendering*, P. M. Hanrahan and W. Purgathofer, Eds. Springer-Verlag, 41–50.
- SUN, B., RAMAMOORTHY, R., NARASIMHAN, S. G., AND NAYAR, S. K. 2005. A practical analytic single scattering model for real-time rendering. *ACM Trans. Graph.* 24, 3, 1040–1049.
- WARD, G. 1991. Real pixels. In *Graphics Gems II*, 80–83.
- WARD, G. J. AND HECKBERT, P. 1992. Irradiance Gradients. In *Proceedings of the 3rd Eurographics Workshop on Rendering*. Bristol, UK, 85–98.
- WARD, G. J., RUBINSTEIN, F. M., AND CLEAR, R. D. 1988. A ray tracing solution for diffuse interreflection. In *Proceedings of the 15th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*. ACM Press, 85–92.

Received June 2007; accepted November 2007