


PUBLICATION
INTERNE
N° 1623



RADIANCE CACHING FOR EFFICIENT GLOBAL ILLUMINATION
COMPUTATION

JAROSLAV KŘIVÁNEK, PASCAL GAUTRON, SUMANTA PATTANAİK,
KADI BOUATOUCH

Radiance Caching for Efficient Global Illumination Computation

Jaroslav Křivánek^{*}, Pascal Gautron^{**}, Sumanta Pattanaik^{***}, Kadi Bouatouch^{****}

Systèmes numériques
Projet Siames

Publication interne n° 1623 — Juin 2004 — 17 pages

Abstract: In this paper we present a ray tracing based method for accelerated global illumination computation in scenes with low-frequency glossy BRDFs. The method is based on sparse sampling, caching, and interpolating radiance on glossy surfaces. In particular we extend the irradiance caching scheme of [WRC88] to cache and interpolate directional incoming radiance instead of irradiance. The incoming radiance at a point is represented by a vector of coefficients with respect to a spherical or hemispherical basis. The surfaces suitable for interpolation are selected automatically according to the glossiness of their BRDF. We also propose a novel method for computing translational radiance gradient at a point.

Key-words: global illumination, ray tracing, spherical harmonics, directional distribution

(Résumé : tsvp)

^{*} xkrivanj@fel.cvut.cz
^{**} pascal.gautron@irisa.fr
^{***} sumant@cs.ucf.edu
^{****} kadi@irisa.fr

Stockage de luminance pour un calcul efficace d'éclairage global

Résumé : Nous présentons dans ce document une méthode de calcul rapide d'illumination globale dans des scènes comprenant des BRDF non diffuses de basse fréquence. Notre approche, basée sur le lancé de rayons, repose sur un faible taux échantillonnage de la luminance incidente, le stockage des échantillons ainsi que leur interpolation sur des surfaces non diffuses. Plus particulièrement nous proposons d'étendre les travaux de [WRC88] ("irradiance caching") au stockage et à l'interpolation de la luminance incidente (directionnelle). La luminance incidente en un point est représentée à l'aide d'un vecteur de coefficients correspondant à sa projection dans une base de fonctions sphériques ou hémisphériques. Les surfaces pour lesquelles l'interpolation peut être utilisée sont sélectionnées automatiquement selon le caractère spéculaire de leur BRDF. De plus nous proposons un nouveau mode de calcul du gradient de translation.

Mots clés : illumination globale, lance de rayons, harmoniques sphériques, distribution directionnelle

1 Introduction

Monte Carlo ray tracing is the method of choice for computing images of complex environments with global illumination [War94]. It produces high quality images, handles general lighting phenomena and scene descriptions, and scales well with the scene size. Even for the radiosity method, high quality images are created by the final gathering, often making use of Monte Carlo (MC) ray tracing [GD04].

MC ray tracing is expensive when it comes to computing illumination on surfaces with low frequency BRDFs (bi-directional reflectance distribution functions). Too many rays have to be traced to get a reasonably precise estimate of the outgoing radiance at a point. Fortunately, there is a high degree of coherence in the outgoing radiance field on those surfaces [WRC88, Gro92, Guo98, SHS98], which demonstrates itself as a smooth illumination. The coherence of radiance can be exploited by interpolation [WRC88, BDT99] to obtain a significant performance gain.

Our goal is to accelerate global illumination computation in the presence of surfaces with low-frequency glossy BRDFs in the context of MC ray tracing. We achieve it by sparse sampling, caching, and interpolating radiance on those surfaces. In particular we extend Ward's irradiance caching scheme [WRC88, WH92] to glossy surfaces. Ward's work is based on the observation that reflected radiance on diffuse surfaces due to indirect illumination changes very slowly with position. This is true for all surfaces with low-frequency BRDFs, even if they are glossy. We therefore extend Ward's work to cache and interpolate the directional incident radiance instead of the irradiance. This provides the necessary information for the illumination integral evaluation (Eq. 1) in the presence of a glossy BRDF. We dub the new method *radiance caching*.

The incoming radiance at a point is represented by spherical or hemispherical harmonics (HSH, see Appendix A) and radiance interpolation reduces to interpolating the HSH coefficients. The illumination integral evaluation reduces to a dot product of the interpolated incoming radiance coefficients and the coefficients of the BRDF representation. This saves many BRDF evaluations.

The quality of interpolation is enhanced by the use of translational gradients. We propose a novel method for computing gradients that is more general than the method of Ward and Heckbert [WH92].

Radiance caching shares all the advantages of the original Ward's work: computation is concentrated in visible parts of the scene; no restrictions are imposed on the scene geometry; implementation and integration with a ray tracer is easy. Our approach can be directly used with any measured BRDF represented by spherical or hemispherical harmonics.

This paper extends the initial description of radiance caching given in [GKPB04]. The main contributions of this paper are:

- Extension of the irradiance caching algorithm to support glossy surfaces.
- Automatic method for selecting BRDFs for which radiance caching is used.
- New method for computing translational radiance gradient at a point.
- Integration of radiance caching in a distribution ray tracer.

The rest of the paper is organized as follows. Section 2 summarizes the related work. Section 3 gives an overview of how radiance caching works and how it is integrated in a distribution ray tracer. Section 4 details different aspect of radiance caching. Section 5 gives the results. Section 6 discusses various topics not covered in the algorithm description. Section 7 concludes the paper and summarizes our ideas for future work.

2 Related work

Interpolation. Interpolation is used in global illumination where there is a certain level of smoothness in the radiometric quantity being computed. All radiosity methods use interpolation in the form of surface discretization. They adapt to the irradiance smoothness by adaptive geometry subdivision, *e.g.* [HSA91, Hec91, LTG92]

In the context of MC ray tracing many approaches have been proposed for *screen space interpolation* [Guo98, WDP99, WGD02, BWG03]. The goal of these methods is to display an approximate solution quickly, possibly at interactive frame rates. However, they do not accelerate the computation of the final high quality solution. *Object space interpolation* has also been used for the purpose of interactive previewing [SS00, TPWG02]. Sparse sampling and interpolation for high quality rendering was used by [BDT99] and [WRC88]. The approach of Bala et al. is suitable only for deterministic ray tracing. Ward et al. use interpolation only for diffuse surfaces. Our approach extends this work to support caching and interpolation on glossy surfaces.

Caching Directional Distribution. Caching the directional representation of radiometric quantities has been used to extend the radiosity method to glossy surfaces, *e.g.* [SAWG91, AH93, SDS95, PB95, SH95, LF96, CLSS97, SSG⁺00]. Sillion et al. [SAWG91, SDS95] use spherical harmonics to represent directional distributions.

In the context of MC ray tracing, directional distribution caching was used by [SHS98, Kat02, TMS04]. In those papers, the incoming radiance at a point is represented by keeping the list of radiance samples used for hemisphere sampling. Slusallek et al. and Kato use reprojection of the samples for improved interpolation. Tawara et al. selectively update the sample list in time to exploit temporal coherence. Slusallek et al. and Tawara et al. mention the suitability of their representation for accelerating rendering with glossy surfaces. The papers discuss only the diffuse case, though. Storing light particles in the scene can also be thought of as caching the directional distribution, since the incoming direction is usually retained with the particle [Jen01, ZSP98].

Illumination Gradient Computation. Arvo [Arv94] computes the Jacobian of the vector irradiance at a point due to partially occluded polygonal emitters of constant radiosity. Holzschuch and Sillion [HS95] handle polygonal emitters with arbitrary radiosity. Ward and Heckbert [WH92] compute irradiance gradient at a point using the information from stochastic hemisphere sampling. Our gradient computation also uses hemisphere sampling.

Spherical Function Representation. Piecewise constant representation [CLSS97, SHS98, Kat02, TMS04] is simple but prone to aliasing and usually very memory demanding.

Unless higher order wavelets are used, even wavelet representation [SH95, PB95, SS95, LF96, SSG⁺00] does not remove the aliasing problems. But with higher order basis functions the mathematics becomes complicated and hence discourages their use.

Spherical Harmonics [CMS87, SAWG91, WAT92, SDS95, RH02, Ram02, KSS02, SKS02, SHHS03] remove the aliasing problem and are efficient for representation of low-frequency functions. However, representation of sharp functions require large number of coefficients and ringing artifacts might appear. Hemispherical harmonics [GKPB04] are better suited for representing functions on a hemisphere. Basis functions very similar to spherical harmonics are Zernike polynomials [WC92, KvDS96] and hemispherical harmonics of Makhotkin [Mak96]. Unlike for spherical harmonics, the rotation matrices are not available for these basis functions. We choose spherical and hemispherical harmonics for their good anti-aliasing properties, low storage cost and ease of use.

Irradiance Caching. Ward et al. [WRC88] propose irradiance caching as a means of computing indirect diffuse inter-reflections in a distribution ray tracer [War94]. They exploit the smoothness of the indirect illumination by sampling the irradiance sparsely over surfaces, caching the results and interpolating them.

For each ray hitting a surface, the irradiance cache is queried. If one or more irradiance records are available, the irradiance is interpolated. Otherwise a new irradiance record is computed by sampling the hemisphere and is added to the cache. In this way, the cache gets filled lazily, progressively in a view dependent manner. As it gets filled, more and more irradiance computations can be carried out by interpolation. Ward uses an octree for storing the irradiance records. In [WH92] the interpolation quality is improved by the use of irradiance gradients.

We retain the basic structure of the original algorithm, but each record stores the incoming radiance function over the hemisphere. This allows interpolation to be applied to glossy surfaces.

3 Algorithm Overview

Radiance caching is a part of a distributed ray tracing approach to global illumination. Rays are cast from the camera into the scene through pixels. At every ray-surface intersection, the outgoing (reflected) radiance has to be computed by evaluating the *illumination integral*:

$$L(\theta_o, \phi_o) = \int_0^{2\pi} \int_0^{\pi/2} L_i(\theta_i, \phi_i) f(\theta_i, \phi_i, \theta_o, \phi_o) \cos \theta_i \sin \theta_i d\theta_i d\phi_i \quad (1)$$

where L is the outgoing radiance, L_i is the incoming radiance and f is the BRDF. The integral is split into parts and each of them is solved by a different technique:

- Direct illumination uses deterministic method for point light sources and area sampling for area light sources [SW94].
- Perfect specular reflections/refractions are solved by tracing a single deterministic secondary ray.
- Ward's irradiance caching computes the indirect diffuse term for *purely* diffuse surfaces.

- Two different techniques may be used for glossy surfaces.

Low-frequency BRDF. Our radiance caching computes the indirect glossy *and diffuse* terms.

High-frequency BRDF. MC importance sampling computes the indirect glossy term and irradiance caching computes the indirect diffuse term.

Radiance caching is not used for high frequency BRDFs since many HSH coefficients would be needed. High-frequency BRDFs are also well localized and importance sampling provides good accuracy with few secondary rays. The distinction between low- and high-frequency BRDF is done automatically as described in Section 4.1. Steps of the rendering algorithm related to radiance caching are shown in Figure 3.

```

// preprocessing - BRDF conversion
for (every surface in the scene) do
  if (surface suitable for radiance caching) then
    - Compute and store the HSH representation of the surface's BRDF.
  end if
end for

// rendering with radiance caching
for (every ray-surface intersection p) do
  - Retrieve the HSH representation of the BRDF at p.
  if ( HSH representation of BRDF not available) then
    // high-frequency BRDF
    - MC importance sampling computes the indirect glossy term.
    - Irradiance caching computes the indirect diffuse term.
  else
    // low-frequency BRDF; use radiance caching
    if ( one or more radiance cache records exist near p ) then
      - Compute the HSH coefficients of the incident radiance at p by gradient based
        interpolation.
    else
      - Compute incident radiance at p by sampling the hemisphere above p.
      - Compute HSH coefficients of the incoming radiance.
      - Compute the translational gradient of the coefficients.
      - Store the new radiance record in the radiance cache.
    end if
    - Compute the outgoing radiance at p as the dot product of the coefficient vector of
      the incoming radiance and that of the BRDF.
  end if
end for

```

Figure 1: Outline of the radiance caching algorithm.

The i -th radiance cache record contains:

- position \mathbf{p}_i ,
- local coordinate frame $(\mathbf{u}_i, \mathbf{v}_i, \mathbf{n}_i)$,
- HSH coefficient vector Λ_i representing the incoming radiance,
- two derivative vectors $\frac{\partial \Lambda_i}{\partial x}$ and $\frac{\partial \Lambda_i}{\partial y}$ representing the translational gradient,
- harmonic mean distance R_i of objects visible from \mathbf{p}_i .

In the rest of the paper Λ denotes a HSH coefficient vector and λ_l^m denotes the coefficients, that is $\Lambda = \{\lambda_l^m\}$.

Each record stores an incident radiance function which can be reused for different viewpoints. The records are stored in an octree as described by Ward et al. [WRC88].

4 Radiance Caching Details

4.1 BRDF Representation

We precompute the HSH representation of surface BRDFs using the method by Kautz et al. [KSS02] which we briefly describe here.

We discretize the outgoing hemisphere. For each discrete outgoing direction (θ_o, ϕ_o) we use HSH to represent the BRDF values over the incoming hemisphere (*i.e.* one HSH coefficient vector per one discretized (θ_o, ϕ_o)). The n -th order representation of a cosine weighted¹ BRDF $f_{(\theta_o, \phi_o)}$ for an outgoing direction (θ_o, ϕ_o) is

$$f_{(\theta_o, \phi_o)}(\theta_i, \phi_i) \approx \sum_{l=0}^n \sum_{m=-l}^l c_l^m(\theta_o, \phi_o) H_l^m(\theta_i, \phi_i), \quad (2)$$

where

$$c_l^m(\theta_o, \phi_o) = \int_0^{2\pi} \int_0^{\pi/2} f(\theta_o, \phi_o, \theta_i, \phi_i) H_l^m(\theta_i, \phi_i) \sin \theta_i d\theta_i d\phi_i. \quad (3)$$

H_l^m are the hemispherical harmonics basis functions (Appendix A). We sample the outgoing hemisphere for (θ_o, ϕ_o) using the parabolic parametrization proposed by Heidrich and Seidel [HS99].

Adaptive BRDF Representation. BRDFs are converted to HSH before the rendering starts. For each outgoing direction, the representation of $f_{(\theta_o, \phi_o)}$ uses the minimum order n sufficient for not exceeding the user specified maximum error. The error is measured as described in [RH02]. If no order $n < n_{\max}$ is sufficient for the specified error, the HSH representation is discarded: radiance caching will not be used for that BRDF and that outgoing direction. After applying this procedure, only low-frequency BRDFs are represented using HSH and radiance caching is used for them. n_{\max} is user specified; we use $n_{\max} = 10$ for our examples. The aim is to have n_{\max} such that a BRDF is classified as low-frequency if and only if radiance caching is more efficient than MC importance sampling. The required HSH order may vary significantly with the outgoing direction, higher order (more coefficients) is usually required for grazing angles (see Figure 2).

4.2 Incoming Radiance Computation

Whenever interpolation is not possible at \mathbf{p} because none of the radiance cache records meet the error criterion (Section 4.4), a new radiance cache record has to be computed and stored in the cache.

We represent the incoming radiance function L_i at a point \mathbf{p} by a vector of HSH coefficients $\Lambda = \{\lambda_l^m\}$:

$$L_i(\theta, \phi) \approx \sum_{l=0}^n \sum_{m=-l}^l \lambda_l^m H_l^m(\theta, \phi), \quad (4)$$

where n is the order of HSH representation and

$$\lambda_l^m = \int_0^{2\pi} \int_0^{\pi/2} L_i(\theta, \phi) H_l^m(\theta, \phi) \sin \theta d\theta d\phi. \quad (5)$$

We compute λ_l^m by MC quadrature with uniform sampling of the hemisphere of incoming directions:

$$\lambda_l^m = \frac{2\pi}{N} \sum_{k=1}^N L_i(\theta_k, \phi_k) H_l^m(\theta_k, \phi_k), \quad (6)$$

where $L_i(\theta_k, \phi_k)$ is the incoming radiance coming from the sampled direction (θ_k, ϕ_k) and N is the number of sampled directions. Direct light source contribution is excluded from L_i . We use a fixed number N of directions but adaptive hemisphere sampling [SP01, SC94] is desirable.

The HSH order n for the incoming radiance representation is equal to the order of the BRDF representation at \mathbf{p} . This cuts off high frequencies from the incoming radiance. The approach is justified by the fact that low-frequency BRDF acts as a low-pass filter on the radiance [Ram02].

¹All BRDFs are premultiplied by the cosine term $\cos \theta_i$ of the illumination integral (Eq. 1) before computing the HSH representation.

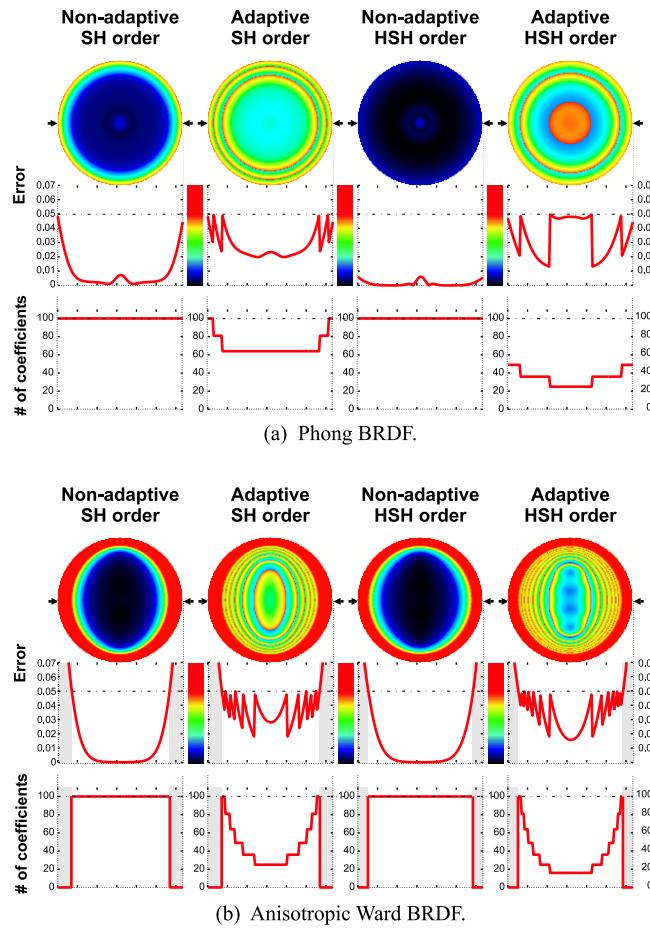


Figure 2: Adaptive BRDF representation for (a) Phong BRDF with exponent $h = 15$ and (b) anisotropic Ward BRDF [Ward 1992] with $k_d = 0$, $k_s = 1$, $\alpha_x = 0.6$, $\alpha_y = 0.25$. The color disks represent BRDF error for different outgoing directions (θ_o, ϕ_o) . Directions are mapped on the disk with the parabolic parametrization. (One can imagine the disks as looking at the hemisphere from the top.) The graphs represent one scanline from the disk images (*i.e.* fixed y and varying x component of the outgoing direction). Adaptive BRDF representation (columns 2 and 4) chooses the minimum number of coefficients sufficient for not exceeding the user specified maximum error (here 5%). For the Ward BRDF, radiance caching is not used for some directions, since the error is too high. Hemispherical harmonics (columns 3 and 4) give lower errors or fewer coefficients than spherical harmonics (columns 1 and 2).

4.3 Translational Gradient Computation

The coefficients λ_l^m of the HSH projection of the incoming radiance are computed with Eq. (6). We want to compute the translational gradient $\nabla \lambda_l^m$ for each λ_l^m , *i.e.* the rate of change of λ_l^m with a differential displacement of the sampling point \mathbf{p} on the surface.

The form of Equation (6) is similar to Ward and Heckbert’s formula for computing irradiance (Equation (2) in [WH92]). The difference is in the choice of the probability density for the hemisphere sampling: we use uniform density whereas Ward and Heckbert use cosine density. Another difference is the weighting function for incoming radiance samples: we use $H_l^m(\theta, \phi)$ while no weight is used in [WH92]. We did not succeed in extending Ward and Heckbert’s gradient computation to handle our case since their method tightly couples the cosine probability density and the weighting.

Instead, we have developed two new methods for computing translational gradient $\nabla \lambda_l^m$. The first, numerical, displaces the center of the hemisphere. The second, analytic, is based on differentiating the terms of Eq. (6). In both cases we compute the gradient $\nabla \lambda_l^m = \left[\frac{\partial \lambda_l^m}{\partial x}, \frac{\partial \lambda_l^m}{\partial y}, 0 \right]$ by computing the partial derivatives $\partial \lambda_l^m / \partial x$ and $\partial \lambda_l^m / \partial y$. The gradient is defined in the local coordinate frame at the point \mathbf{p} . The derivative with respect to z is not computed since we assume displacement only in the tangent plane. We compute the translational gradients and the coefficients λ_l^m simultaneously.

4.3.1 Numerical Gradient Computation

To compute the derivative $\partial\lambda_l^m/\partial x$ numerically, we displace the point \mathbf{p} , along the local X axis, by Δx to \mathbf{p}' (Figure 3). For each Monte Carlo sample $L_i(\theta_k, \phi_k)$ we:

1. Compute the new direction (θ'_k, ϕ'_k) at \mathbf{p}' as $(\theta'_k, \phi'_k) = \frac{\mathbf{q}_k - \mathbf{p}'}{r'_k}$. Here \mathbf{q}_k is the point hit by the ray from \mathbf{p} in direction (θ_k, ϕ_k) and $r'_k = \|\mathbf{q}_k - \mathbf{p}'\|$. We will also denote $r_k = \|\mathbf{q}_k - \mathbf{p}\|$. See Figure 3 for the various terms used here.
2. Compute the solid angle Ω'_k associated with the new direction (θ'_k, ϕ'_k) . The solid angle Ω_k associated with each direction in Equation (6) is uniform and equal to $2\pi/N$. With the displacement of the point \mathbf{p} , the solid angles no longer remain uniform. The change in solid angle is due to the change in distance $r_k = \|\mathbf{q}_k - \mathbf{p}\|$ and orientation of the surface at \mathbf{q}_k , as seen from the hemisphere center \mathbf{p} or \mathbf{p}' . The solid angle before the displacement is

$$\Omega_k = \Delta A_k \frac{\cos \xi_k}{r_k^2} = \frac{2\pi}{N},$$

where ξ_k is the angle between the surface normal at \mathbf{q}_k and the vector from \mathbf{q}_k to \mathbf{p} . The area $\Delta A_k = \frac{2\pi}{N} \frac{r_k^2}{\cos \xi_k}$ is the part of the environment visible through Ω_k . It does not change with the displacement because we assume that the environment visible from \mathbf{p} and \mathbf{p}' is the same. After the displacement the solid angle subtended by ΔA_k becomes

$$\Omega'_k = \Delta A_k \frac{\cos \xi'_k}{r_k'^2} = \frac{2\pi}{N} \frac{r_k^2 \cos \xi'_k}{r_k'^2 \cos \xi_k}.$$

We now estimate the coefficient $\lambda_l^{m'}$ at \mathbf{p}' as

$$\lambda_l^{m'} = \frac{2\pi}{N} \sum_{k=1}^N \frac{r_k^2 \cos \xi'_k}{r_k'^2 \cos \xi_k} L_i(\theta_k, \phi_k) H_l^m(\theta'_k, \phi'_k),$$

and finally we compute $\partial\lambda_l^m/\partial x$ as

$$\frac{\partial\lambda_l^m}{\partial x} = \frac{\lambda_l^{m'} - \lambda_l^m}{\Delta x}.$$

The computation of $\partial\lambda_l^m/\partial y$ proceeds in a similar way. This completes the numerical estimation of the translational gradient $\nabla\lambda_l^m$ at the point of interest.




Figure 3: Gradient computation by displacing the hemisphere center from \mathbf{p} to \mathbf{p}' ((a) before and (b) after the displacement). The quantities changing with the displacement are (shown in red): sample ray direction (θ_k, ϕ_k) ; the solid angle Ω_k associated with this sample; the angle ξ_k between the sample direction and the surface normal at the hit point \mathbf{q}_k . Neither the hit point \mathbf{q}_k nor the area ΔA_k visible through Ω_k change with the displacement.

4.3.2 Analytic Gradient Computation

We rewrite Equation (6) as

$$\lambda_l^m = \sum_{k=1}^N \Omega_k L_i(\theta_k, \phi_k) H_l^m(\theta_k, \phi_k), \quad (7)$$

with $\Omega_k = \frac{2\pi}{N}$ for uniform hemisphere sampling. We have seen that Ω_k does not remain constant with displacement of \mathbf{p} and therefore it has to be included in the sum and differentiated. The notation in this section is as above.

The partial derivative $\partial\lambda_l^m/\partial x$ is computed by differentiating the terms of the sum in Eq. (7)

$$\begin{aligned}\frac{\partial\lambda_l^m}{\partial x} &= \sum_{k=1}^N \frac{\partial}{\partial x} \left(\Omega_k L_i(\theta_k, \phi_k) H_l^m(\theta_k, \phi_k) \right) = \\ &= \sum_{k=1}^N L_i(\theta_k, \phi_k) \left(\frac{\partial\Omega_k}{\partial x} H_l^m(\theta_k, \phi_k) + \Omega_k \frac{\partial H_l^m(\theta_k, \phi_k)}{\partial x} \right)\end{aligned}\quad (8)$$

The derivative of the basis function is

$$\frac{\partial H_l^m(\theta_k, \phi_k)}{\partial x} = \frac{\partial\theta_k}{\partial x} \frac{\partial H_l^m(\theta_k, \phi_k)}{\partial\theta_k} + \frac{\partial\phi_k}{\partial x} \frac{\partial H_l^m(\theta_k, \phi_k)}{\partial\phi_k}, \quad (9)$$

with [Wei04b]

$$\begin{aligned}\partial\theta_k/\partial x &= -\cos\theta_k \cos\phi_k/r_k, \\ \partial\phi_k/\partial x &= \sin\phi_k/(r_k \sin\theta_k).\end{aligned}\quad (10)$$

Those derivatives with respect to y would be

$$\begin{aligned}\partial\theta_k/\partial y &= -\cos\theta_k \sin\phi_k/r_k, \\ \partial\phi_k/\partial y &= -\cos\phi_k/(r_k \sin\theta_k).\end{aligned}\quad (11)$$

Derivatives $\partial H_l^m/\partial\theta_k$ and $\partial H_l^m/\partial\phi_k$ are given in the Appendix B.

The derivative of the solid angle Ω_k is

$$\frac{\partial\Omega_k}{\partial x} = \frac{\partial}{\partial x} \Delta A_k \frac{\cos\xi_k}{r_k^2} = \Delta A_k \frac{\partial \cos\xi_k}{\partial x} \frac{1}{r_k^2},$$

The area $\Delta A_k = \frac{2\pi}{N} \frac{r_k^2}{\cos\xi_k}$ is the part of the environment visible through Ω_k . It does not change with the displacement. The change of $\cos\xi_k/r_k^2$ with the displacement of \mathbf{p} is opposite to its change with the displacement of $\mathbf{q}_k = (q_x, q_y, q_z)$, *i.e.*

$$\frac{\partial \cos\xi_k}{\partial x} \frac{1}{r_k^2} = -\frac{\partial}{\partial q_x} \frac{\cos\xi_k}{r_k^2}.$$

The derivative $\frac{\partial}{\partial q_x} \frac{\cos\xi_k}{r_k^2}$ can be computed with the assumption that \mathbf{p} lies at the origin because only the relative position of \mathbf{p} and \mathbf{q}_k matters (see Figure 4). Since $\cos\xi_k = -\frac{\mathbf{n}_k \cdot \mathbf{q}_k}{r_k}$ and $r_k = \|\mathbf{q}_k\| = \sqrt{q_x^2 + q_y^2 + q_z^2}$ we have

$$\begin{aligned}\frac{\partial}{\partial q_x} \frac{\cos\xi_k}{r_k^2} &= -\frac{\partial}{\partial q_x} \frac{n_x q_x + n_y q_y + n_z q_z}{(q_x^2 + q_y^2 + q_z^2)^{3/2}} \\ &= -\frac{r_k^2 n_x - 3q_x (\mathbf{n}_k \cdot \mathbf{q}_k)}{r_k^5} \\ &= -\frac{r_k n_x + 3q_x \cos\xi_k}{r_k^4}.\end{aligned}\quad (12)$$

Here $\mathbf{n}_k = (n_x, n_y, n_z)$ is the surface normal at \mathbf{q}_k . Combining this result with $\Delta A_k = \frac{2\pi}{N} \frac{r_k^2}{\cos\xi_k}$, we get

$$\frac{\partial\Omega_k}{\partial x} = \frac{2\pi}{N} \frac{r_k n_x + 3q_x \cos\xi_k}{r_k^2 \cos\xi_k}.\quad (13)$$

Plugging Equations (13) and (9) into (8) we get the complete formula for $\partial\lambda_l^m/\partial x$. The formulas for $\partial\lambda_l^m/\partial y$ are similar; only Equation (10) must be replaced by Equation (11).

A similar gradient calculation was also proposed in [ZSP98]. Their method disregards the change of Ω_k and hence does not give good results.




Figure 4: Quantities in the computation of $\frac{\partial}{\partial q_x} \frac{\cos \xi_k}{r_k^2}$.

4.3.3 Discussion

For the derivation of both numerical and analytic methods we assumed:

- The radiance $L_i(\theta_k, \phi_k)$ from the point \mathbf{q}_k incident at \mathbf{p} does not change with the displacement of \mathbf{p} .
- Visibility of ΔA_k , the small area around \mathbf{q}_k , does not change with the displacement of \mathbf{p} .

Though none of these assumptions is necessarily valid in all scenes, they are reasonable for small displacements.

The quantities changing with the displacement are:

- Direction (θ_k, ϕ_k) of the MC sample contributing to the sum (6). As a consequence, the value of $H_l^m(\theta_k, \phi_k)$ changes too.
- The size of the solid angle Ω_k through which the contribution to the sum (6) is brought. The change of Ω_k is due to the change in distance r_k and in the apparent orientation of the surface ΔA_k around the hit point \mathbf{q}_k .

The numerical and analytic methods are equivalent, their results are indistinguishable. The numerical method is easier to implement since we do not need to evaluate the basis function derivatives. The analytic method is numerically more stable near edges and corners and also slightly faster to evaluate.

Unlike Ward's method, ours does not assume any stratification of the hemisphere. This allows one to use our gradient computation with different sampling strategies, *e.g.* quasi Monte Carlo sampling.

4.3.4 Irradiance Gradient Computation

Note that both of the methods we propose can be still used if H_l^m is replaced by any other hemispherical function. We also do not rely on uniform hemisphere sampling. Any probability density $p(\theta, \phi)$ can be used for sampling. The only change is that $\Omega_k = \frac{1}{Np(\theta_k, \phi_k)}$ instead of $\Omega_k = \frac{2\pi}{N}$ used for the uniform sampling.

As an example we compute the irradiance gradient ∇E with a cosine-weighted hemisphere sampling. $H_l^m(\theta, \phi)$ is replaced by $\cos \theta$, the probability density of sampling in direction (θ, ϕ) is $p(\theta, \phi) = \frac{\cos \theta}{\pi}$ and therefore $\Omega_k = \frac{\pi}{N \cos \theta_k}$. The resulting formula for the analytic method is

$$\frac{\partial E}{\partial x} = \sum_{k=1}^N L_i(\theta_k, \phi_k) \left(\frac{\partial \Omega_k}{\partial x} \cos \theta_k - \frac{\pi \sin \theta_k}{N \cos \theta_k} \frac{\partial \theta_k}{\partial x} \right)$$

with

$$\frac{\partial \Omega_k}{\partial x} = \frac{\pi}{N \cos \theta_k} \frac{r_k n_x + 3q_x \cos \xi_k}{r_k^2 \cos \xi_k}.$$

We implemented this irradiance gradient computation method and that of Ward and Heckbert and we compared them on a sample scene (Figure 5). The results were similar for both methods. Ward's method gives better results when there are surfaces seen at very sharp grazing angles from the sampling point \mathbf{p} . Otherwise our method gives slightly smoother results.

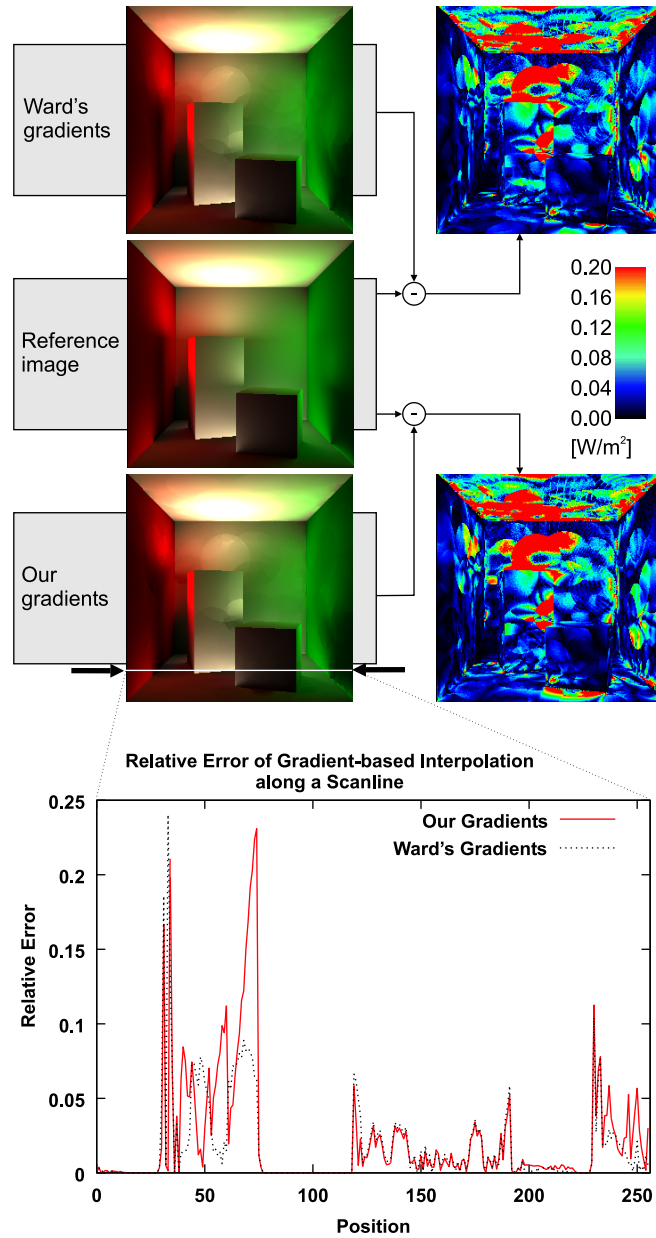


Figure 5: Comparison of irradiance gradient computation. The scene is a diffuse Cornell box; only first bounce indirect illumination is computed. The color-coded images show the difference between the gradient-based interpolation and the reference solution (10,000 samples per hemisphere). RMS error of the images is 0.125 for Ward’s method and 0.131 for our method. The graph shows relative error of the interpolation along a single scanline. Ward’s method gives better results when there are surfaces seen at very sharp grazing angles from the sampling point. Otherwise our method gives slightly lower error.

4.4 Radiance Interpolation

We use a weighted interpolation scheme similar to the one proposed in [WH92] for interpolating the coefficient vectors $\Lambda_i = \{\lambda_{i,m}^m\}$ at any required surface point \mathbf{p} . The difference is that we replace the use of the rotational gradient by applying a rotation to the coefficient vector Λ_i of each cache record i used for interpolation at \mathbf{p} . This aligns the coordinate frame at the position \mathbf{p}_i of the cache record and the frame at \mathbf{p} (see Figure 6). The weight $w_i(\mathbf{p})$ of record i with respect to \mathbf{p} is

$$w_i(\mathbf{p}) = \left(\frac{\|\mathbf{p} - \mathbf{p}_i\|}{R_i} + \sqrt{1 - \mathbf{n} \cdot \mathbf{n}_i} \right)^{-1}, \quad (14)$$

where \mathbf{n} is the surface normal at \mathbf{p} , \mathbf{n}_i is the surface normal at \mathbf{p}_i , and R_i is the harmonic mean distance to objects visible from \mathbf{p}_i . The coefficient vector of the interpolated radiance is:

$$\Lambda(\mathbf{p}) = \frac{\sum_S \mathbf{R}_i \left(\Lambda_i + d_x \frac{\partial \Lambda_i}{\partial x} + d_y \frac{\partial \Lambda_i}{\partial y} \right) w_i(\mathbf{p})}{\sum_S w_i(\mathbf{p})}, \quad (15)$$

where $S = \{i | w_i(\mathbf{p}) > 1/a\}$ and a is a user defined desired accuracy. The definition of the set S is the criterion used to decide which radiance cache records can be used for interpolation. Derivatives $\partial \Lambda_i / \partial x$ and $\partial \Lambda_i / \partial y$ are the translational gradient components and (d_x, d_y) are the displacements of $\mathbf{p} - \mathbf{p}_i$ along the X and Y axes of the record i 's local coordinate frame. \mathbf{R}_i is the HSH rotation matrix [GKPB04] that aligns the coordinate frame at \mathbf{p}_i with the frame at \mathbf{p} . It transforms the whole coefficient vector Λ_i after the translational gradient has been applied.




Figure 6: Rotation \mathbf{R}_i has to align the coordinate frame at \mathbf{p}_i with that at \mathbf{p} before interpolation is possible.

4.5 Outgoing Radiance Computation

In general the outgoing radiance at a point \mathbf{p} is computed by evaluating the illumination integral (1). As both the incoming radiance L_i and the BRDF f in the integral are represented as HSH coefficient vectors, the integral evaluation reduces to the dot product [Sze75]:

$$L(\theta_o, \phi_o) = \sum_{l=0}^{n-1} \sum_{m=-l}^l \lambda_l^m f_l^m(\theta_o, \phi_o).$$

λ_l^m is an interpolated incoming radiance coefficient and $f_l^m(\theta_o, \phi_o)$ is a BRDF coefficient at \mathbf{p} for the outgoing direction (θ_o, ϕ_o) . This computation saves many BRDF evaluations.

5 Results

Figure 7 shows three renderings of a Cornell box with glossy back wall (Phong, exponent 22), taken from three different viewpoints at resolution 850×850 . The number of rays cast to sample each hemisphere is $N = 4000$. The order of HSH used is $n = 10$. Except for the back wall, all objects are Lambertian. Only direct lighting and first bounce reflection for the back wall were computed. Image 7(b) took 25.8 sec. to render² with 1121 radiance cache samples. Images 7(a) and (c) took 11.2 and 10.3 seconds to render since only 250 and 198 additional radiance cache samples were needed. Memory consumption of the radiance cache was 1.5 MB. The timings compare well with 260 seconds rendering time for the images of approximately the same visual quality generated with MC importance sampling.

6 Discussion

The Rotation Loss. There is a loss of information when radiances are interpolated on a curved surface (Figure 8). It is due to the alignment (rotation \mathbf{R}_i) of the radiance records with the coordinate frame at the interpolation point \mathbf{p} .

A part of the radiance incident at \mathbf{p}_i should disappear under the surface (marked ‘a’ in Fig. 8) and should not contribute to the interpolated radiance at \mathbf{p} . A part of the radiance actually incident at \mathbf{p} is not represented by the radiance record at \mathbf{p}_i (marked ‘b’ in Fig. 8) and is therefore missing in the interpolated radiance.

²The timings were measured on a 2.26GHz Pentium 4 with 512MB RAM running Windows XP.




Figure 7: Cornell box with glossy back wall rendered using radiance caching. 1121 radiance cache records are used to render the center view (b). The number of additional records needed when the viewpoint changes to left (a) and right (c) is 250 and 198. The size of each of the images is 850×850 .

This problem is not due to using a hemispherical basis for representing the incoming radiance, but due to the fact the incident radiance at a surface point *is* a hemispherical function. Using spherical harmonics instead of HSH would not solve this problem. In practice the error introduced by this problem is very small because the difference between the normal at \mathbf{p} and the normal at any record used for interpolation at \mathbf{p} is small. Note also that this problem is present in the Ward's irradiance caching as well.




Figure 8: Loss of information when radiances are interpolated on a curved surface. A part of the radiance actually incident at \mathbf{p}_i should disappear under the surface (marked 'a'). A part of the radiance incident at \mathbf{p} is not represented by the radiance record at \mathbf{p}_i (marked 'b').

Global vs. Local Coordinates. Incoming radiance at a point \mathbf{p} can be represented either in the local frame at \mathbf{p} (*i.e.* aligned with the surface normal at \mathbf{p}) or in the global frame. This influences the way the interpolation at \mathbf{p} is performed:

Incoming radiance in the global frame

```

for (each available record  $i$  at  $\mathbf{p}$ ) do
  - Update the interpolation sum.
end for
- Align the interpolation result with the local frame at  $\mathbf{p}$ .
- Compute the dot product.

```

Incoming radiance in the local frame

```

for (each available record  $i$  at  $\mathbf{p}$ ) do
  - Align the local frame at  $\mathbf{p}_i$  with the local frame at  $\mathbf{p}$ .
  - Update the interpolation sum.
end for
- Compute the dot product.

```

The final dot product between the BRDF coefficient vector and the interpolated coefficient vector of the incoming radiance is always carried out in the local frame at \mathbf{p} .

On curved surfaces, less alignment operations (rotations) are needed if the incoming radiance is represented in the global frame than if it is represented in the local frame. On the other hand, if the incoming radiance is represented in the local frame, no alignment (even with the BRDF) is needed on flat surfaces. The best performance (*i.e.* least rotations) is obtained if the incoming radiance is represented in the local frame on flat surfaces and in the global frame on curved surfaces. Note that full spherical function representation (*e.g.* spherical harmonics) is needed to represent the incoming radiance in the global frame.

Suitability of Hemispherical Harmonics. We use (hemi)spherical harmonics since they are simple, computationally efficient (manipulations of vectors of floats) and avoid aliasing. The essential disadvantage is the lack of directional localization. When we create a new radiance cache record, the full hemisphere must be sampled, whatever the incoming ray direction is. The more directional the BRDF is, the more this approach becomes wasteful. With a basis that localizes in directions, only that required part of the hemisphere need to be sampled. For this purpose one can use piecewise constant representation [SHS98, Kat02, TMS04], but it would presumably introduce too much aliasing. Spherical wavelets [SS95] are probably a good choice. This is left for further investigation. Even though we have implemented radiance caching using spherical and hemispherical harmonics only, the approach should be valid with any other set of basis function constructed in a similar way.

7 Conclusion and Future Work

We have presented a method for accelerating computation of the indirect illumination on surfaces with low-frequency glossy BRDFs. Our method is based on sparse sampling, caching, and interpolating incoming radiance on glossy surfaces. We have shown that this approach is more efficient than pure Monte Carlo sampling at every surface point and delivers images of superior quality.

In the future we would like to use adaptive hemisphere sampling to compute the incoming radiance coefficients [SC94, ZSP98, RFS03], use a different representation for the incoming radiance that localizes better in directions and devise interpolation criteria better suited for glossy surfaces. In the long term we also want to investigate the relationship between the frequency content of a BRDF and the suitability of interpolating radiance on surfaces with that BRDF.

References

- [AH93] Larry Aupperle and Pat Hanrahan. A hierarchical illumination algorithm for surfaces with glossy reflection. *Computer Graphics*, 27(Annual Conference Series):155–162, 1993.
- [Arv94] James Arvo. The irradiance jacobian for partially occluded polyhedral sources. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 343–350. ACM Press, 1994.
- [BDT99] Kavita Bala, Julie Dorsey, and Seth Teller. Radiance interpolants for accelerated bounded-error ray tracing. *ACM Trans. Graph.*, 18(3):213–256, 1999.
- [BWG03] Kavita Bala, Bruce Walter, and Donald Greenberg. Combining edges and points for interactive high-quality rendering. *ACM Trans. Graph. (Proceedings of SIGGRAPH 2003)*, 22(3), July 2003.
- [CLSS97] Per H. Christensen, Dani Lischinski, Eric J. Stollnitz, and David H. Salesin. Clustering for glossy global illumination. *ACM Trans. Graph.*, 16(1):3–33, 1997.
- [CMS87] Brian Cabral, Nelson Max, and Rebecca Springmeyer. Bidirectional reflection functions from surface bump maps. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 273–281. ACM Press, 1987.
- [GD04] Xavier Granier and George Drettakis. A final reconstruction approach for a unified global illumination algorithm. *ACM Trans. Graph.*, 23(2):163–189, 2004.
- [GKPB04] Pascal Gautron, Jaroslav Křivánek, Sumanta N. Pattanaik, and Kadi Bouatouch. A novel hemispherical basis for accurate and efficient rendering. In *Eurographics Symposium on Rendering (to appear)*, June 2004.
- [Gro92] Eduard Groeller. *Coherence in Computer Graphics*. PhD thesis, Technische universität Wien, 1992.
- [Guo98] B. Guo. Progressive radiance evaluation using directional coherence maps. In *Computer Graphics (SIGGRAPH '98 Proceedings)*, pages 255–266, 1998.
- [Hec91] Paul S. Heckbert. *Simulating Global Illumination Using Adaptive Meshing*. PhD thesis, University of California, Berkeley, June 1991.

- [HS95] Nicolas Holzschuch and François Sillion. Accurate computation of the radiosity gradient with constant and linear emitters. In *Sixth Eurographics Workshop on Rendering*, June 1995. Dublin, Ireland.
- [HS99] Wolfgang Heidrich and Hans-Peter Seidel. Realistic, hardware-accelerated shading and lighting. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 171–178. ACM Press/Addison-Wesley Publishing Co., 1999.
- [HSA91] Pat Hanrahan, David Salzman, and Larry Aupperle. A rapid hierarchical radiosity algorithm. *Computer Graphics (Proceedings of SIGGRAPH 1991)*, 25(4):197–206, 1991.
- [Jen01] Henrik Wann Jensen. *Realistic Image Synthesis Using Photon Mapping*. AK Peters, July 2001.
- [Kat02] T. Kato. Photon mapping in Kilauea. In *Siggraph 2002, Course Notes No. 43, A Practical Guide to Global Illumination using Photon Mapping*, pages 122–191, 2002.
- [KSS02] Jan Kautz, Peter-Pike Sloan, and John Snyder. Fast, arbitrary BRDF shading for low-frequency lighting using spherical harmonics. In *Proceedings of the 13th Eurographics workshop on Rendering*, pages 291–296. Eurographics Association, 2002.
- [KvDS96] J.J. Koenderink, A.J. van Doorn, and M. Stavridi. Bidirectional reflection distribution function expressed in terms of surface scattering modes. *ECCV*, B:28–39, 1996.
- [LF96] Robert R. Lewis and Alain Fournier. Light-driven global illumination with a wavelet representation. In *Rendering Techniques '96 (Proceedings of the Seventh Eurographics Workshop on Rendering)*, pages 11–20, New York, NY, 1996. Springer-Verlag/Wien.
- [LTG92] Daniel Lischinski, Filippo Tampieri, and Donald P. Greenberg. Discontinuity meshing for accurate radiosity. *IEEE Computer Graphics and Applications*, 12(6):25–39, November 1992.
- [Mak96] Oleg A. Makhotkin. Analysis of radiative transfer between surfaces by hemispherical harmonics. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 56(6):869–879, 1996.
- [PB95] Sumanta N. Pattanaik and Kadi Bouatouch. Haar wavelet: A solution to global illumination with general surface properties. In *Photorealistic Rendering Techniques (Proceedings of Fifth Eurographics Workshop on Rendering)*. Springer-Verlag, Berlin, June 1995.
- [Ram02] Ravi Ramamoorthi. *A Signal-Processing Framework for Forward and Inverse Rendering*. PhD thesis, Stanford University, 2002.
- [RFS03] Jaume Rigau, Miquel Feixas, and Mateu Sbert. Refinement criteria based on f-divergences. In *Proceedings of the 14th Eurographics workshop on Rendering*, pages 260–269. Eurographics Association, 2003.
- [RH02] Ravi Ramamoorthi and Pat Hanrahan. Frequency space environment map rendering. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 517–526. ACM Press, 2002.
- [SAWG91] François X. Sillion, James R. Arvo, Stephen H. Westin, and Donald P. Greenberg. A global illumination solution for general reflectance distributions. In *Proceedings of the 18th annual conference on Computer graphics and interactive techniques*, pages 187–196. ACM Press, 1991.
- [SC94] Peter Shirley and Kenneth Chiu. Notes on adaptive quadrature on the hemisphere. Technical Report TR-411, Dept. of Computer Science, Indiana University, July 1994.
- [SDS95] François Sillion, George Drettakis, and Cyril Soler. A clustering algorithm for radiance calculation in general environments. In *Rendering Techniques*, June 1995.
- [SH95] P. Schröder and Patrick Hanrahan. Wavelet methods for radiance computations. In G. Sakas and P. Shirley, editors, *Photorealistic Rendering Techniques (Proceedings Fifth Eurographics Workshop on Rendering)*, pages 310–326. Springer-Verlag, Berlin, 1995.
- [SHHS03] Peter-Pike Sloan, Jesse Hall, John Hart, and John Snyder. Clustered principal components for precomputed radiance transfer. *ACM Trans. Graph. (Proceedings of SIGGRAPH 2003)*, 22(3):382–391, 2003.
- [SHS98] P. Slusallek, W. Heidrich, and H.-P. Seidel. Radiance maps: An image-based approach to global illumination. *SIGGRAPH '98, Technical Sketch*, 1998.
- [SKS02] Peter-Pike Sloan, Jan Kautz, and John Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. Graph. (Proceedings of SIGGRAPH 2002)*, 21(3):527–536, 2002.
- [SP01] Xavier Serpaggi and Bernard Péroche. An adaptive method for indirect illumination using light vectors. In A. Chalmers and T.-M. Rhyne, editors, *Computer Graphics Forum (EUROGRAPHICS 2001 Proceedings)*, volume 20(3), pages 278–287. Blackwell Publishing, 2001.
- [SS95] Peter Schröder and Wim Sweldens. Spherical wavelets: efficiently representing functions on the sphere. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 161–172. ACM Press, 1995.
- [SS00] Maryann Simmons and Carlo H. Séquin. Tapestry: A dynamic mesh-based display representation for interactive rendering. In *Proceedings of the 11th Eurographics Workshop on Rendering*, pages 329–340, June

- 2000.
- [SSG⁺00] Marc Stamminger, Annette Scheel, Xavier Granier, Frederic Perez-Cazorla, George Drettakis, and François Sillion. Efficient glossy global illumination with interactive viewing. *Computer Graphics Forum*, 19(1), 2000.
- [SW94] Peter Shirley and Changyaw Wang. Direct lighting calculation by Monte Carlo integration. In P. Brunet and F. W. Jansen, editors, *Photorealistic Rendering in Computer Graphics (Proceedings of the Second Eurographics Workshop on Rendering)*, pages 54–59, New York, NY, 1994. Springer-Verlag.
- [Sze75] Gabor Szegő. *Orthogonal polynomials*. American Mathematical Society, Providence, Rhode Island, 4 edition, 1975.
- [TMS04] Takehiro Tawara, Karol Myszkowski, and Hans-Peter Seidel. Exploiting temporal coherence in final gathering for dynamic scenes. In *Proceedings of Computer Graphics International (to appear)*, June 2004.
- [TPWG02] Parag Tole, Fabio Pellacini, Bruce Walter, and Donald P. Greenberg. Interactive global illumination in dynamic scenes. *ACM Trans. Graph. (Proceedings of SIGGRAPH 2002)*, 21(3):537–546, July 2002.
- [War92] Gregory J. Ward. Measuring and modeling anisotropic reflection. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 265–272. ACM Press, 1992.
- [War94] Gregory J. Ward. The radiance lighting simulation and rendering system. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 459–472. ACM Press, 1994.
- [WAT92] Stephen H. Westin, James R. Arvo, and Kenneth E. Torrance. Predicting reflectance functions from complex surfaces. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 255–264. ACM Press, 1992.
- [WC92] James C. Wyant and Katherine Creath. Basic wavefront aberration theory for optical metrology. In *Applied optics and Optical Engineering, Vol XI*, pages 27–39. Academic Press, Inc., 1992.
- [WDGD02] Bruce Walter, George Drettakis, Donald P. Greenberg, and Oliver Deussen. Enhancing and optimizing the render cache. In *Proceedings of the 10th Eurographics Workshop on Rendering*, June 2002.
- [WDP99] Bruce Walter, George Drettakis, and Steven Parker. Interactive rendering using render cache. In *Proceedings of the 13th Eurographics Workshop on Rendering*, pages 19–30, 1999.
- [Wei04a] Eric W. Weisstein. Legendre polynomial. From *MathWorld – A Wolfram Web Resource*. <http://mathworld.wolfram.com/LegendrePolynomial.html>, 2004.
- [Wei04b] Eric W. Weisstein. Spherical coordinates. From *MathWorld – A Wolfram Web Resource*. <http://mathworld.wolfram.com/SphericalCoordinates.html>, 2004.
- [WH92] Gregory J. Ward and Paul S. Heckbert. Irradiance gradients. In *Proceedings of the 2nd Eurographics Workshop on Rendering*, 1992.
- [WRC88] Gregory J. Ward, Francis M. Rubinstein, and Robert D. Clear. A ray tracing solution for diffuse interreflection. In *Proceedings of the 15th annual conference on Computer graphics and interactive techniques*, pages 85–92. ACM Press, 1988.
- [ZSP98] Jacques Zaninetti, Xavier Serpaggi, and Bernard Péroche. A vector approach for global illumination in ray tracing. *Computer Graphics Forum (EUROGRAPHICS '98 Proceedings)*, 17(3), 1998.

A Hemispherical Harmonics

The hemispherical harmonics (HSH) basis functions are defined as

$$H_l^m(\theta, \phi) = \begin{cases} \sqrt{2}\tilde{K}_l^m \cos(m\phi) P_l^m(2\cos\theta - 1) & \text{if } m > 0 \\ \sqrt{2}\tilde{K}_l^m \sin(-m\phi) P_l^{-m}(2\cos\theta - 1) & \text{if } m < 0 \\ \tilde{K}_l^0 P_l^0(\cos\theta) & \text{if } m = 0 \end{cases}$$

where P_l^m are the associated Legendre polynomials [Wei04a] and \tilde{K}_l^m is the following normalization value:

$$\tilde{K}_l^m = \sqrt{\frac{(2l+1)(l-|m|)!}{2\pi(l+|m|)!}}.$$

The definition domain is $(\theta, \phi) \in [0, \pi/2] \times [0, 2\pi]$, $l \in \{0, 1, \dots\}$, $m = \{-l, \dots, 0, \dots, l\}$. Further information on hemispherical harmonics is given in [GKPB04].

B SH and HSH Derivatives

The partial derivatives for spherical harmonics are:

$$\frac{\partial y_l^m}{\partial \theta}(\theta, \phi) = \begin{cases} -\sqrt{2}K_l^m \cos(m\phi) \sin(\theta) \frac{dP_l^m}{dx}(\cos \theta) & \text{if } m > 0 \\ -\sqrt{2}K_l^m \sin(-m\phi) \sin(\theta) \frac{dP_l^{-m}}{dx}(\cos \theta) & \text{if } m < 0 \\ -K_l^0 \sin(\theta) \frac{dP_l^0}{dx}(\cos \theta) & \text{if } m = 0, \end{cases}$$

$$\frac{\partial y_l^m}{\partial \phi}(\theta, \phi) = \begin{cases} 0 & \text{if } m = 0 \\ -my_l^{-m}(\theta, \phi) & \text{otherwise.} \end{cases}$$

The derivative of the associated Legendre polynomials can be found from the recurrence formula:

$$\frac{dP_l^m}{dx}(x) = \begin{cases} \frac{1}{x^2-1} \left(x l P_l^m(x) - (m+l) P_{l-1}^m(x) \right) & \text{if } m < l \\ -(-1)^m x (2m-1)!! (1-x^2)^{\frac{m}{2}-1} & \text{if } m = l, \end{cases}$$

where $x!!$ is the double factorial (product of all odd integers less than or equal to x).

The partial derivatives for hemispherical harmonics are:

$$\frac{\partial H_l^m}{\partial \theta}(\theta, \phi) = \begin{cases} -2\sqrt{2}\tilde{K}_l^m \cos(m\phi) \sin(\theta) \frac{dP_l^m}{dx}(2\cos \theta - 1) & \text{if } m > 0 \\ -2\sqrt{2}\tilde{K}_l^m \sin(-m\phi) \sin(\theta) \frac{dP_l^{-m}}{dx}(2\cos \theta - 1) & \text{if } m < 0 \\ -2\tilde{K}_l^0 \sin(\theta) \frac{dP_l^0}{dx}(2\cos \theta - 1) & \text{if } m = 0, \end{cases}$$

$$\frac{\partial H_l^m}{\partial \phi}(\theta, \phi) = \begin{cases} 0 & \text{if } m = 0 \\ -mH_l^{-m}(\theta, \phi) & \text{otherwise.} \end{cases}$$