

# Real-Time Diffuse Global Illumination Using Radiance Hints

Georgios Papaioannou\*

Athens University of Economics and Business




Figure 1: Examples of indirect diffuse lighting calculated in real time with the Radiance Hints method.

## Abstract

GPU-based interactive global illumination techniques are receiving an increasing interest from both the research and the industrial community as real-time graphics applications strive for visually rich and realistic dynamic three-dimensional environments. This paper presents a fast new diffuse global illumination method that generates a sparse set of low-cost radiance field evaluation points (radiance hints) and computes an arbitrary number of diffuse inter-reflections within a given volume. The proposed approximate technique combines ideas from existing grid-based radiance caching techniques with reflective shadow maps as well as a stochastic scheme for visibility calculations, in order to achieve high frame rates for multiple light bounces.

**CR Categories:** I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Shading, shadowing

**Keywords:** Global Illumination, Real Time Rendering, Programmable Graphics Hardware

## 1 Introduction

Modern interactive graphics applications, such as game engines and virtual reality systems strive for higher visual realism through the inclusion of global illumination effects that are special cases of the solution of the *rendering equation* for a three-dimensional scene [Kajiya 1986]. Specialized algorithms for the real-time rendering of specular global illumination effects such as true and approximate reflections, refractive distortions and caustics have been successfully implemented on graphics processors. Although results

\*e-mail: gepap@aub.gr

© ACM, 2011. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version will be published in the HPG 2011 proceedings.

from the adaptation of generic global illumination methods to the GPU (e.g. photon mapping and path tracing) are promising, current methods cannot yet meet the high frame rate requirements of typical applications for large data sets.

Since on most natural or man-made surfaces a significant amount of light is diffusely reflected, the simulation of diffuse light transport within a virtual environment plays a major role in the realistic image synthesis. The simplified rendering equation for diffuse inter-reflections only has been extensively used in existing interactive global illumination algorithms as it is very efficient to approximate by direct irradiance integral evaluation. Furthermore, diffuse indirect lighting varies slowly across space, thus enabling the sparse evaluation and caching of radiance as a volumetric *radiance field*. Its low-frequency nature facilitates the perceptually convincing blending of direct and indirect illumination even for rough approximations of the latter [Durand et al. 2005].

The *Radiance Hints* method presented in this paper estimates the radiance field within the volume of a three-dimensional environment by directly sampling the *Reflective Shadow Maps* of the light sources (see also the related work section below). A low-cost geometry-less evaluation of the radiance reflected from the first diffuse bounce of light is performed at sparse locations (Radiance Hints). Radiance Hints are subsequently used as radiance cache points to evaluate surface irradiance on nearby surfaces and approximate secondary diffuse inter-reflections, both at a very low cost.

## 2 Related Work

The concept of indirect illumination interpolation from cached values was introduced by Ward et al [1988]. Accurate irradiance estimates are computed using ray tracing on a few surface points (irradiance sample points) and for the remaining surface points fast interpolation is used. Wang et al. [2009] presented a method to calculate the irradiance sample points in advance and implemented the algorithm on the GPU. The method is accurate but it achieves interactive frame rates only in very simple scenes. The Irradiance Volume, which was first introduced in [Greger et al. 1998], regards a set of isolated irradiance samples, parameterized by direction, each one storing the incoming light for a particular point on a two-level volume grid. Surface irradiance was then approximated by interpolating the stored irradiance in the lattice vertices. The work of Nijasure et al. [2005] is based on the same principle but uses a

spherical harmonics representation to store the incident radiance at the vertices of a regular grid. Incident radiance is estimated by rendering a cube map at each grid point sampling the resulting texels and encoding the radiance field as spherical harmonics coefficients. The indirect illumination at the surface points is approximated by interpolating the radiance from the closest grid points. This method supports multiple bounces and indirect occlusion but is very expensive for complex dynamic scenes or many cache points due to the capture and encoding of the cube maps.

*Instant Radiosity* methods, introduced by Keller [1997], approximate the indirect illumination of a scene using a set of *virtual point lights* (VPLs). VPLs are points in space that act as light sources and encapsulate light reflected off a surface at a given location. A number of photons are traced into the scene and VPLs are created at surface hit points, then the scene is rendered, as lit by each VPL. The major cost of this method is the calculation of shadows from a potentially large number of point lights but since it does not require any complex data structures it is a very good candidate for a GPU implementation. Lightcuts [Walter et al. 2005] and more recent work such as [Dong et al. 2009] reduce the number of the required shadow queries by clustering the VPLs in groups and using one shadow query per cluster, but the performance is still low for complex scenes.

*Reflective Shadow maps* (RSMs) [Dachsbacher and Stamminger 2005] consider the pixels of a shadow map as VPLs, but the contribution of these lights is gathered without taking scene occlusion into account. To achieve interactive frame rates, screen space interpolation is required and the method is limited to the first bounce of indirect illumination. An extension of this method by the same authors [Dachsbacher and Stamminger 2006] transforms the gathering operation for the incident radiance on a surface into an image-space VPL splatting procedure which achieves better performance, but still ignores occlusion for the indirect lighting. Recently, Nichols and Wyman [2010] further improved the performance of RSMs by performing multi-resolution splatting. *Imperfect Shadow Maps* [Ritschel et al. 2008] use a point-based representation of the scene to efficiently render extremely rough approximations of the shadow maps for all the VPLs in one pass. The method achieves good results at interactive frame rates for one-bounce inter-reflections.

Jensen [1996] introduced the concept of photon mapping, where in a first pass photons are traced from the light sources into the scene and stored in a k-d tree and in a second pass the indirect illumination of visible surface points is approximated by gathering the k nearest photons. In [McGuire and Luebke 2009] the first bounce of the photons is computed using rasterization on the GPU, continues the photon tracing on the CPU for the rest of the bounces and finally scatters the illumination from the photons using the GPU. Since part of the photon tracing still runs on the CPU, a large number of parallel cores are required to achieve interactive frame-rates. Recently Yao et al. [2010] introduced a photon splatting method that calculates ray-surface intersections in the image space of multiple environment maps (with normal and depth information) carefully placed within the volume of the scene by optimizing volume coverage. The method dispenses with ray-polygon intersections by employing the Distance Impostors method proposed in [Szirmay-Kalos et al. 2005], which performs ray-environment map intersections.

Ritschel et al. [2009] extended previous work for screen space ambient occlusion calculation [Shanmugam and Arikan 2007] and introduced a method to approximate the first indirect near field diffuse bounce of the light by only using information in the 2D frame buffer. This method has a very low computational cost but the resulting illumination is hardly accurate since it depends on the projection of the (visible only) objects on the screen.

Kaplanyan and Dachsbacher introduced a volume-based method, the Cascaded Light Propagation Volumes [Kaplanyan and Dachsbacher 2010], where RSM-generated virtual point lights are injected into a volume texture. Instead of evaluating the low frequency indirect illumination from the VPLs at the surface points though, the method uses an iterative propagation scheme to transfer energy from voxel to voxel, taking into account occlusion caused by blocking voxels. Blocking voxels are marked by storing any available depth information from the camera depth map and the RSMs into a separate occlusion (geometry) volume. The method achieves high performance for a relatively small number of propagation iterations with respect to the volume size, which in most cases is not adequate to model multiple bounces of diffuse inter-reflection. The method also suffers from popping artifacts due to view-dependent occlusion information.

Thiedemann et al. [2011] introduced an interactive volume-based global illumination method, where the spatial occupancy and color data are generated by injecting a geometry texture atlas containing point samples of the polygonal geometry. The authors propose an optimized ray marching scheme for intersecting the gathering rays with the volume data. At each hit voxel, the RSMs are looked up to determine its visibility from the source, and subsequently its normal and the incident light. Their method produces fast and high quality first-bounce diffuse global illumination but requires model preprocessing and extra storage for the texture atlas and is sensitive to the point sampling rate and surface deformations. The authors suggest that secondary indirect light bounces should be handled by iterative path tracing but do not provide a solution. Mavridis and Papaioannou adopted a similar volume population approach [Mavridis and Papaioannou 2011], where occupancy and direct illumination of the geometry are injected into the volume as vertices of the tessellated geometry. Supplementary points are generated by injecting the view camera and RSM G-buffers into the volume, similar to [Kaplanyan and Dachsbacher 2010]. Ray marching at voxel level is subsequently applied to simulated global illumination with multiple bounces at a cost proportional to the volume size and sampling distance. Both methods produce a full-scene occupancy volume and can thus correctly handle secondary light bounce occlusion.

There exist several high performance methods that rely on pre-computed radiance transfer and visibility functions, such as [Sloan et al. 2002], but a detailed discussion on this specific genre of algorithms is outside the scope of this paper, as fully dynamic scenes and lighting conditions are considered.

### 3 Method Overview

The proposed approximate method combines ideas from the grid-based radiance caching of Nijasure et al. [2005] and the Reflective Shadow Maps approach. Similar to [Nijasure et al. 2005], it stores sparse radiance field values at regular volume locations. Instead of constructing cube maps to sample the visibility and incoming radiance, though, it employs reflective shadow map sampling in one pass for all sampling locations. This allows for a significant performance increase in radiance field computation at the selected points and does not require any additional rendering of the scene's geometry, as all necessary information is already stored in the RSMs of the light sources. Furthermore, the radiance cache points are updated and queried using a stochastic approach as discussed in Sections 4 and 5, hence the name *Radiance Hints - RH*.

Multiple light bounces are handled by stochastically exchanging energy among the RH points. To account for occlusion in secondary diffuse inter-reflections, a fast light attenuation scheme is proposed, which takes into account the measured distances from the RSM samples and imposes no additional cost.

Compared to other modern techniques, four major improvements of the Radiance Hints method are the fast estimation of the spatially varying radiance field, the ability to capture multiple diffuse inter-reflections with a very small overhead per light bounce, the physically-based modeling of the light transport and its view- and geometry-independent nature. A comprehensive comparative evaluation of Radiance Hints against the latest real-time diffuse global illumination methods is provided in Section 9. An outline of the proposed method is provided below.

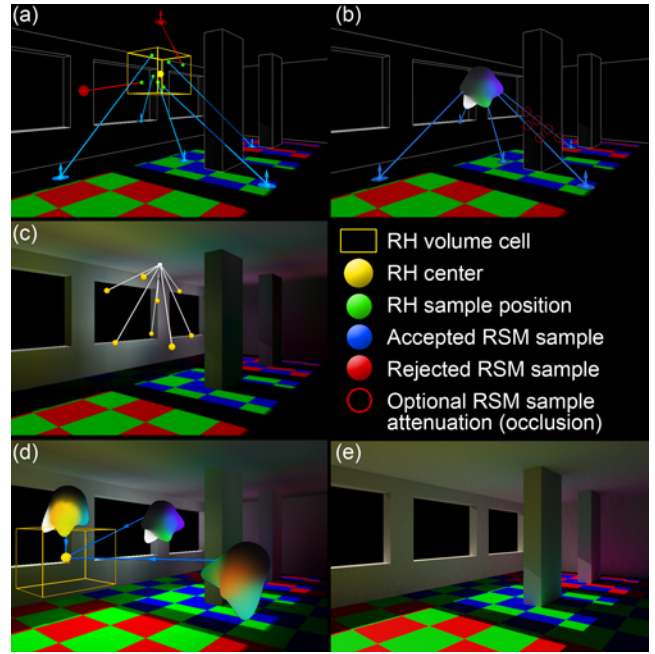
- Prepare the RSMs for each light source.
- Configure a lattice of Radiance Hints as a volume texture. If  $N_b > 1$  light bounces are requested, prepare a second set of RH points for interleaved radiance field update.
- For each Radiance Hint, sample the RSMs from a number of locations inside the RH cell to build its approximate radiance field (Fig. 2(a), Section 4). Encode the result as spherical harmonics coefficients and store them at the RH voxel (Fig. 2(b)).
- For all  $N_b - 1$  bounces, update the Radiance Hints by accumulating radiance from other RH points (Fig. 2(d), Section 6). Use stochastic occlusion in the radiance transfer process (Section 7).
- Estimate the irradiance at each fragment using the contribution of appropriately selected Radiance Hint samples (Fig. 2(c), Section 5)

In the current implementation of the method, the RH locations are selected as the lattice nodes of a uniform partitioning of the scene volume. For each RH and RSM pair, the radiance field is sampled from multiple locations inside the RH volume cell, with an average position equal to the the Radiance Hint location, i.e. the voxel center. The RH position is the representative of the individual jittered RH samples and the approximate radiance field in its volume cell. The resulting radiance field, encoded as Spherical Harmonics coefficients for the three color channels is stored in volume textures, along with the minimum and maximum RSM sample distance. The two distance values are used in the secondary light bounce calculations and the approximate visibility function (see Sections 6 and 7 respectively). The individual RH sampling locations are not maintained.

## 4 Radiance Hint Calculation

In the Reflective Shadow Maps method [Dachsbacher and Stamminger 2005], indirect light at a point  $\mathbf{p}$  is estimated by projecting it on the reflective shadow map of a light source.  $N_s$  normal, position and brightness samples are acquired in a disk centered at the projected coordinate pair in shadow map image space and the contribution of each one is accumulated. RSM samples with a normal vector facing away from  $\mathbf{p}_{RH}$  are rejected. The disk covers a  $c_{RSM}$  fraction of the total RSM.

In Radiance Hints, instead of considering a fixed location, i.e. the RH center  $\mathbf{p}_{RH}$  for all  $N_s$  shadow map samples, different positional samples with a distribution  $D_{RH}$  around each  $\mathbf{p}_{RH}$  are taken (Fig. 2(a)). This results in an estimate of the overall radiance field, representative of the entire space in the vicinity of  $\mathbf{p}_{RH}$  and not just a single location.



**Figure 2:** Estimation and application of the RH radiance field. (a) RSM sampling from randomly placed points inside the effective volume of each Radiance Hint. (b) Optional depth-based occlusion attenuation of RSM samples followed by radiance field encoding. (c) Global illumination reconstruction from near radiance field. (d) RH-to-RH secondary light bounces. (e) Reference path-traced image.

Let  $\Phi_T(\mathbf{x})$  be the incident flux on each RSM texel  $\mathbf{x}$ . For a distant or directional light with total power  $\Phi_L$ , each texel receives a flux  $\Phi_T(\mathbf{x}) = \Phi_L \langle \mathbf{l} | \mathbf{n} \rangle / M_{RSM}$ .  $M_{RSM}$  is the total number of texels in the reflective shadow map,  $\mathbf{l}$ ,  $\mathbf{n}$  are the light vector and the normal at  $\mathbf{x}$ , respectively and  $\langle \cdot | \cdot \rangle$  denotes the dot product. For wide spotlight apertures this power can be also analytically computed. Since the RSM is sampled only at  $N_s$  locations, the incident flux  $\Phi_i(\mathbf{x}_k)$  at each sample  $\mathbf{x}_k$  is the corresponding texel flux scaled by  $M_{RSM}/N_s$ .


Assuming a sufficient sampling of the shadow map, i.e. a small lighting variation near each sample  $\mathbf{x}_k$ , the irradiance in the vicinity of  $\mathbf{x}_k$  becomes  $E_i(\mathbf{x}_k) \simeq \Phi_i(\mathbf{x}_k)/A_k$ , where  $A_k$  is the surface area near  $\mathbf{x}_k$ . Unless a very dense sampling of the RSM is performed,  $A_k$  corresponds to larger area than that of a projected RSM texel. Please note that RSM samples are not regarded as point lights here, but rather as small area lights.

Since only diffuse inter-reflections are considered, the outgoing radiance in any direction in the hemisphere above  $\mathbf{x}_k$  is constant and is given by:

$$L_o(\mathbf{x}_k, \omega) = L_o(\mathbf{x}_k) = \frac{\rho(\mathbf{x}_k)}{\pi} E_i(\mathbf{x}_k) = \frac{\rho(\mathbf{x}_k) \Phi_i(\mathbf{x}_k)}{\pi A_k} \quad (1)$$

Let now  $\mathbf{y}_k$  be the  $k$ -th radiance sample position around the Radiance Hint center  $\mathbf{p}_{RH}$  (Fig. 2(a)). Considering a narrow solid angle from patch  $A_k$  to  $\mathbf{y}_k$ , the cumulative radiance contribution of  $A_k$  to the radiance field at  $\mathbf{y}_k$  around direction  $\mathbf{x}_k \rightarrow \mathbf{y}_k$  is:

$$L(\mathbf{y}_k) = \int_{\Omega_{A_k}} L_o(\mathbf{x}_k, \omega) d\omega = \int_{A_k} \frac{L_o(\mathbf{x}) \langle \mathbf{y}_k - \mathbf{x} | \mathbf{n}_k \rangle}{|\mathbf{y}_k - \mathbf{x}|^2} d\mathbf{x} \quad (2)$$



**Figure 3:** Three-bounce diffuse global illumination with (a) 2<sup>nd</sup> and (b) 3<sup>rd</sup> order spherical harmonics (Sibenik Cathedral model,  $19 \times 15 \times 9$  RH grid, 50 RSM samples per RH). (c) Path-traced image. All images were rendered at  $1024 \times 1024$  pixels and brightened for clarity.

$A_k$  is represented by a single sample and thus, replacing outgoing radiance from Eq. 1 in Eq. 2,  $L(\mathbf{y}_k)$  is approximated as follows:

$$\tilde{L}(\mathbf{y}_k) = \frac{\rho(\mathbf{x}_k) \Phi_i(\mathbf{x}_k) \langle \mathbf{y}_k - \mathbf{x}_k | \mathbf{n}_k \rangle}{\pi |\mathbf{y}_k - \mathbf{x}_k|^2} \quad (3)$$

Equation. 3 can be rewritten as:

$$\tilde{L}(\mathbf{y}_k) = \frac{\rho(\mathbf{x}_k) \Phi_T(\mathbf{x}_k) M_{RMS} \langle \mathbf{y}_k - \mathbf{x}_k | \mathbf{n}_k \rangle}{\pi N_s |\mathbf{y}_k - \mathbf{x}_k|^2} \quad (4)$$

$\rho(\mathbf{x}_k) \Phi_T(\mathbf{x}_k)$ , i.e. the output flux of each texel, is stored in the RSM illumination channel texels. Equation 4 is evaluated for each one of the  $N_s$  selected RSM samples in the volume associated with the Radiance Hint centered at  $\mathbf{p}_{RH}$ .

The radiance field can be now expressed using a convenient set of basis functions, such as Spherical Radial Basis Functions (SRBF) [Leung et al. 2006] and Spherical Harmonics (SH) [Sillion et al. 1991]. In this paper, similar to many other global illumination techniques, the light field is approximated by projecting it on the low-frequency terms (order 2 and 3) of the real Spherical Harmonic basis functions, although any set of orthogonal basis functions can be adopted for this purpose.

By projecting each  $\tilde{L}(\mathbf{y}_k)$  radiance field sample into the low-frequency Spherical Harmonics basis functions  $Y_l^m(\theta, \varphi)$ , the estimated radiance field at  $\mathbf{p}_{RH}$  in the direction  $(\theta, \varphi)$  becomes:

$$\begin{aligned} \tilde{L}(\mathbf{p}_{RH}, \theta, \varphi) &\simeq \sum_l \sum_{m=-l}^l \tilde{L}(\mathbf{p}_{RH}, l, m) Y_l^m(\theta, \varphi) \\ \tilde{L}(\mathbf{p}_{RH}, l, m) &= \sum_{k=1}^{N_s} \tilde{L}(\mathbf{y}_k) Y_l^m(\theta_k, \varphi_k) \end{aligned} \quad (5)$$

where  $\theta_k, \varphi_k$  correspond to the direction of incident light at the  $k$ -th RH sample position from the  $k$ -th RSM sample:  $\mathbf{y}_k \rightarrow \mathbf{x}_k$ . Figure 3 demonstrates the image quality difference in the case of 2<sup>nd</sup> and 3<sup>rd</sup> order Spherical Harmonics.

For multiple light sources and consequently different reflective shadow maps, the superposition of samples from each one of them is trivial. The number of samples per light source can be determined by evenly distributing the total  $N_s$  samples among the sources, by power law or any other importance sampling scheme. The same principle also holds for omnidirectional lights implemented using cube maps. Consecutive passes, one for each RSM, accumulate Spherical Harmonics coefficients (as well as other attributes) in the current frame buffer object, using a pixel shader. When regularly spaced RH locations are used, their data are calculated by

rendering the RH attributes into a volume texture, slice by slice. Alternatively, sparse, unstructured RH points could be created and stored along with their floating point Cartesian coordinates in a single layer bound to a two-dimensional texture.

Typical values for the total number of RSM samples  $N_s$  per Radiance Hint is between 12 (single RSM, spotlight) and 70 (6-8 RSMs or omnidirectional lights). The coverage of the RSM samples on the parametric shadow map area  $c_{RSM}$  should be small (5-15%), if a light source illuminates a large area (i.e. sunlight). In this case, indirect illumination exchange is expected to have relatively localized effect and few, concentrated samples are preferred. Indoor spaces and scenes of small physical scale on the other hand, may use up to 90% of RSM coverage. The number of RH points depends on the desired detail of the reconstructed irradiance and the size of the scene. Typically, scenes with clustered elements (such as the Cornell Box example) require fewer Radiance Hints than environments with many isolated thin structures (such as columns in the Sponza Atrium or the Sibenic Cathedral).

## 5 Reconstruction of Indirect Illumination

In order to reconstruct the reflected radiance at any given surface point  $\mathbf{p}$  with normal  $\mathbf{n}$ , the incident radiance  $L_i(\omega)$  from a direction  $\omega$  is interpolated from Radiance Hints in the vicinity of  $\mathbf{p}$  and integrated over the hemisphere centered at  $\mathbf{n}$ . In contrast to other radiance caching techniques, e.g. [Nijasure et al. 2005], where the integral of the radiance field and the hemisphere above  $\mathbf{p}$  is directly applied as a three-dimensional texture, a more careful sampling was sought to remedy the following issue: Radiance samples must be drawn only from the hemisphere above  $\mathbf{p}$  in order to avoid severe light leaking and ensure that all selected samples have a significant contribution to the reconstructed radiance.

The reflected radiance  $L_r(\mathbf{p})$  for a Lambertian surface is estimated by evaluating the irradiance in the hemisphere  $\Omega_{\mathbf{n}}$  above  $\mathbf{p}$  from the radiance field at  $N_{RH}$  locations:

$$\begin{aligned} L_r(\mathbf{p}) &= \frac{\rho(\mathbf{p})}{\pi} \int_{\Omega_{\mathbf{n}}} L_i(\omega) \cos \theta d\omega \simeq \\ &\frac{\rho(\mathbf{p})}{\pi N_{RH}} \sum_{n=1}^{N_{RH}} \int_{\Omega_{\mathbf{n}}} L(\mathbf{p}_{RH_n}, \omega) \cos \theta d\omega \end{aligned} \quad (6)$$

where  $\rho(\mathbf{p})$  and  $\mathbf{n}$  are the albedo and the normal vector at  $\mathbf{p}$ , respectively.

Having expressed the radiance field as a truncated series of orthonormal basis functions and by doing the same for the oriented hemisphere of Eq. 6, the integral over the hemisphere centered at  $\mathbf{n}$  can be efficiently evaluated by a dot product of the corresponding coefficients. The reconstructed reflected radiance  $L_r(\mathbf{p})$  at  $\mathbf{p}$  then becomes:

$$L_r(\mathbf{p}) = \frac{\rho(\mathbf{p})}{\pi N_{RH}} \sum_{n=1}^{N_{RH}} \sum_l \sum_{m=-l}^l \tilde{L}(\mathbf{p}_{RH_n}, l, m) H_{\mathbf{n}}(l, m) \quad (7)$$

where  $H_{\mathbf{n}}(l, m)$  are the SH coefficients of a hemisphere aligned with  $\mathbf{n}$ .

Ramamoorthi and Hanrahan observed that for the real Spherical Harmonics basis functions, the truncated SH expansion of the irradiance integral with  $l \leq 2$  approximates the slowly varying diffusely reflected light with an accuracy of up to 99% [Ramamoorthi and Hanrahan 2001]. An analytical expression for the Spherical Harmonics coefficients for a directed hemisphere and  $l \leq 2$  can also be found in the same paper.

**Sampling the Radiance Hints.** If the Radiance Hints are the vertices of a volume grid, as is the case in the current implementation, this task is trivial. First, an offset  $r_{offset}$  from the shaded point  $\mathbf{p}$  equal to the diagonal of each voxel is chosen. The  $N_{RH}$  sampling locations  $\mathbf{p}_{RH_n}$  are arranged on a surface-aligned hemisphere of radius  $r_{offset}$  and shifted by  $r_{offset}/2$  along the normal direction  $\mathbf{n}$ . This shifting ensures that no RH points behind  $\mathbf{p}$  are sampled. The radiance field at  $\mathbf{p}_{RH_n}$  is trilinearly interpolated from the nearest Radiance Hints. Practically, even using as few as 4 RH samples produces a good estimate of the smoothly varying radiance field. The sample arrangement on the hemisphere is pre-calculated and rotated to match the normal at  $\mathbf{p}$ . For all the test cases in this paper, the first sample is directly above  $\mathbf{p}$ , while the remaining 3 samples are positioned at equidistant locations on a circle 80 degrees away from the normal direction.

## 6 Approximate Multiple Inter-reflections

Instead of distributing energy among surfaces, radiance from secondary bounces is sampled at the RH locations by drawing radiance from other Radiance Hints (Fig. 2(d)). A uniform Monte Carlo sampling of the radiance field surrounding a RH point  $\mathbf{p}_{RH_n}$  can be achieved by uniformly selecting contributing  $N_b$  RH samples in the storage volume inside a sphere of radius  $r_{max}$  and center  $\mathbf{p}_{RH_n}$ .  $r_{max}$  represents a range beyond which contribution of incoming radiance from secondary bounces is considered negligible and can equal the diameter of the scene. Usually though, a much smaller value is preferred as it helps concentrate few samples in a more significant range in terms of energy exchange.

Since we are interested in the contribution of radiance that is reflected back from surrounding geometry, Radiance Hint samples closer to surfaces must be favored. However, the Radiance Hints method uses no explicit geometric information and therefore, a probabilistic weighting mechanism must be devised that associates a normalized confidence factor  $c_j$  with each sample.  $c_j$  must be higher for RH locations near surfaces and lower for Radiance Hints positioned far from scene geometry.

When sampling the RSMs in order to estimate the radiance field at a Radiance Hint, the minimum and maximum distances  $r_{RH_j, \min}$ ,  $r_{RH_j, \max}$  to the sampled RSM points are maintained and saved in the RH data. Both distances are normalized with respect to  $r_{max}$ . The confidence factor  $c_j$  can be directly associated with the minimum distance  $r_{RH_j, \min}$  at the RH sampling location  $\mathbf{p}_{RH_j}$ .  $r_{RH_j, \min}$  is (trilinearly) interpolated from the corresponding stored distance information at the closest Radiance Hints.

In contrast to the irradiance reconstruction from the Radiance Hints (Section 5), where only the radiance field near the surfaces is sampled, in the case of secondary diffuse inter-reflections the receiving Radiance Hint  $\mathbf{p}_{RH_n}$  can be significantly far from the radiance sampling locations. This necessitates the introduction of a visibility term  $V(\mathbf{p}_{RH_j}, \mathbf{p}_{RH_n})$  between the two points. Since no geometric occlusion data are maintained,  $V(\mathbf{p}_{RH_j}, \mathbf{p}_{RH_n})$  is stochastically estimated from the minimum and maximum distances stored in the Radiance Hints, as described in the next section.

Taking the above two factors into account, the updated radiance field at a Radiance Hint  $\mathbf{p}_{RH_n}$  for the  $k$ -th light bounce is:

$$\begin{aligned} \tilde{L}_o^{(k)}(\mathbf{p}_{RH_n}, l, m) &= \tilde{L}^{(k-1)}(\mathbf{p}_{RH_n}, l, m) + \\ &\frac{4\pi}{V_n} \sum_{j=1}^{N_b} V(\mathbf{p}_{RH_j}, \mathbf{p}_{RH_n}) c_j \tilde{L}_o^{(k)}(\mathbf{p}_{RH_j}, -\theta_j, -\varphi_j) Y_l^m(\theta_j, \varphi_j) \end{aligned} \quad (8)$$

where  $\tilde{L}_o^{(k)}(\mathbf{p}_{RH_j}, -\theta_j, -\varphi_j) = L_i^{(k)}(\mathbf{p}_{RH_n}, \theta_j, \varphi_j)$  is the radiance from the  $j$ -th Radiance Hint sample and  $V_n$  is the sum of the visibility  $V(\mathbf{p}_{RH_j}, \mathbf{p}_{RH_n})$  over the  $N_b$  samples. The radiance is integrated only over visible RH samples since the domain of the approximated integral is the spherical solid angle around  $\mathbf{p}_{RH_n}$  and thus only the nearest (visible) points should contribute to the field at  $\mathbf{p}_{RH_n}$ . This modified Monte Carlo integration formula evaluates to the same result as standard Monte Carlo with rejection sampling. The confidence factors  $c_j$  are easily computed as:

$$c_j = (1 - r_{RH_j, \min}) / \sum_{i=1}^{N_b} (1 - r_{RH_i, \min}) \quad (9)$$

Since no surface normal is available at any of the contributing Radiance Hints, the radiance that is reflected back to  $\mathbf{p}_{RH_n}$  from a distant Radiance Hint sample  $\mathbf{p}_{RH_j}$  emanates from a hemisphere aligned with the incoming sampling direction  $(-\theta_j, -\varphi_j)$ . The outgoing radiance  $\tilde{L}_o^{(k)}(\mathbf{p}_{RH_j}, -\theta_j, -\varphi_j)$  for any direction in the hemisphere due to diffuse inter-reflection is directly proportional to the corresponding radiosity and in particular  $\tilde{L}_o^{(k)}(\mathbf{p}_{RH_j}, -\theta_j, -\varphi_j) = B_j^{(k)}(-\theta_j, -\varphi_j)/\pi$ . The radiosity in this case equals the irradiance on the same hemisphere, which can be evaluated from the spherical harmonic representation of the radiance field at  $\mathbf{p}_{RH_j}$  (Eq. 5):

$$\begin{aligned} \tilde{L}_o^{(k)}(\mathbf{p}_{RH_j}, -\theta_j, -\varphi_j) &= \frac{B_j^{(k)}(-\theta_j, -\varphi_j)}{\pi} = \\ &\frac{1}{\pi} \sum_l \sum_{m=-l}^l \tilde{L}^{(k-1)}(\mathbf{p}_{RH_j}, l, m) H_l^m(-\theta_j, -\varphi_j) \end{aligned} \quad (10)$$

A consequence of the above energy exchange mechanism is that color bleeding takes into account only the outgoing radiance of direct illumination, as no interaction with actual geometry occurs in consecutive bounces. This is manifested as a constant (white) reflectance in Eq. 10. Full color bleeding can be achieved by maintaining and sampling a separate low resolution albedo volume representation for the scene at the additional injection or voxelization cost. The number of  $N_b$  RH sample points can be small. 16-32 points are adequate considering the slow variation of the overall radiance field.

In Eq. 8 and 10 the radiance field in bounce step  $k$  depends on radiance field values of the previous diffuse inter-reflection iteration. This necessitates the use of two identical volume texture buffers with alternate read/write roles (ping pong rendering).

## 7 Handling Occlusion

Correct visibility, i.e. interception of light incident to a point  $\mathbf{p}$  due to the obstruction of the light transfer path by geometry, is one of the main bottlenecks in interactive global illumination rendering. Ambient occlusion [Zhukov et al. 1998] has been extensively used as a statistical measure of light attenuation according to the closeness of geometry in the vicinity of  $\mathbf{p}$ , but unfortunately, it leads to erroneous shadowing of areas with strong indirect light exchange (corners, opposite surfaces etc). In volume-based approaches, blocking voxels are injected in the volume representation of the scene either by full-scene voxelization [Gaitatzes et al. 2010; Thiedemann et al. 2011; Mavridis and Papaioannou 2011] or by injecting view-dependent blocking voxels from the RSM and camera depth information [Kaplanyan and Dachsbacher 2010]. With *Anti-radiance* [Dachsbacher et al. 2007], light that should be reflected by blocking geometry is negatively transmitted through surface boundaries and canceled out by unobstructed light at the next surface.

This is transformed into an iterative process where the number of iterations depends on the depth complexity of the scene along a ray. The complexity of the method is prohibitive for real-time rendering of complex scenes. Instant Radiosity approaches and the Imperfect Shadow Maps method, calculate and maintain directional visibility at the VPLs, a process which can be costly for a large number of VPLs and complex geometry.

Similar to the RSM algorithm [Dachsbacher and Stamminger 2005], for the estimation of the first-bounce incident radiance at the Radiance Hint points, light from the RSM VPLs is not blocked. When the light source RSMs sufficiently sample the scene geometry, this mechanism can produce convincing results for one-bounce indirect lighting. In special cases, where the RSM scene coverage is small, however (see Fig. 5 and Fig. 7), the visibility between the RH sample location and the RSM sample location must be evaluated. Such a visibility measure is proposed in Section 8 as an enhancement of the main method.

For multiple diffuse inter-reflections (see Section 6), the visibility function  $V(\mathbf{p}_{RH_n}, \mathbf{p}_{RH_j})$  between two Radiance Hint sampling locations is stochastically estimated using distance-based attenuation, according to the recorded or interpolated minimum and maximum distances from  $\mathbf{p}_{RH_n}$  and  $\mathbf{p}_{RH_j}$  to the RSM samples. When evaluating the incident radiance at a Radiance Hint  $\mathbf{p}_{RH_n}$  coming from a Radiance Hint sample location  $\mathbf{p}_{RH_j}$ , if  $\|\mathbf{p}_{RH_n} - \mathbf{p}_{RH_j}\|$  is less than  $r_{RH_n, \min} + r_{RH_j, \min}$ , no surface intercepts the path between the two points. If  $\|\mathbf{p}_{RH_n} - \mathbf{p}_{RH_j}\| > r_{RH_n, \max} + r_{RH_j, \max}$ , then the two points are farther from the cumulative recorded distance and thus completely occluded. For the zone of uncertainty between the two marginal cases, occlusion is linearly interpolated. The idea is demonstrated in Fig. 4. The above heuristic visibility function is very fast to compute and is compactly expressed as:


$$\begin{aligned}
 V(\mathbf{p}_{RH_n}, \mathbf{p}_{RH_j}) &= V(\mathbf{p}_{RH_j}, \mathbf{p}_{RH_n}) = \\
 &1 - \text{clamp} \left( \frac{\|\mathbf{p}_{RH_n} - \mathbf{p}_{RH_j}\| - r_{n_j, \min}}{r_{n_j, \max} - r_{n_j, \min}}, 0, 1 \right) \quad (11) \\
 r_{n_j, \min} &= r_{RH_n, \min} + r_{RH_j, \min} \\
 r_{n_j, \max} &= r_{RH_n, \max} + r_{RH_j, \max},
 \end{aligned}$$

Although Eq. 11 does not provide exact visibility,  $V(\mathbf{p}_{RH_n}, \mathbf{p}_{RH_j})$  effectively attenuates light from unreachable Radiance Hints at practically no additional cost. The quality of this technique is affected by the number of RSM samples used and the presence of important occluding surfaces in the RSM.

## 8 Additional Enhancements

**RSM sample occlusion attenuation.** Radiance Hints is a geometry-less method and therefore, RSM sampling cannot take into account secondary occlusion from parts of the scene that are not visible in the shadow maps, as mentioned in the previous section. However, an optional, probabilistic sample visibility step can be introduced that attenuates the radiance from each RSM sample  $\mathbf{x}_k$  according to the intersection of the camera depth map with the line of sight between  $\mathbf{x}_k$  and Radiance Hint sampling position  $\mathbf{y}_k$ .

A small number of evenly-spaced samples is taken along the line segment  $(\mathbf{x}_k, \mathbf{y}_k)$  and their screen-space depth is compared with the depth buffer. Typically, if at least one sample lay beyond the recorded camera depth, the RSM sample should have been rejected as hidden. However this is not a reliable criterion, as an interrupted line segment in screen space does not necessarily correspond to a




**Figure 4:** Stochastic occlusion for secondary light bounces using the minimum and maximum RSM sample distance stored in the RH data. Top: Occlusion principle (see text). Bottom: Attenuation example.

truly blocked light path in object space and can lead to severe view-dependent popping artifacts as the camera moves. Instead, the radiance from an RSM sample  $\mathbf{x}_k$  is attenuated by a constant factor for every occluded path sample (see also Fig. 2b). This way, the radiance field is attenuated according to the view-based blocking of RSM sampling paths, but without introducing a strong bias toward such an unreliable source of occlusion data. As a result, in most cases the radiance field is smoothly modified even for abrupt changes in the screen-space visibility. Typically, 3-8 points along each path  $(\mathbf{x}_k, \mathbf{y}_k)$  are examined, which correspond to the respective number of coherent depth map lookups. This introduces a total overhead between 0.8ms and 1.2ms, depending on the number of RSM samples. An example of this enhancement technique is shown in Fig. 5.

Note that sample occlusion attenuation is proposed as an optional image quality improvement and is not part of the core, view-independent algorithm. Unless explicitly noted, all examples use the basic method.

**Multiple RH volumes.** Adjoining areas of a large scene (e.g. level areas separated by portals) may use separate small Radiance Hints volumes, which can easily interact in a hierarchical manner. RH data can be conventionally pre-filtered to generate one or more additional mipmap levels and therefore, neighboring and distant sections of a scene can be sampled using aggregated (averaged) values for the local radiance field.



**Figure 5:** Optional, depth-based RSM sample occlusion attenuation: An extreme case. The middle wall is not present in the RSM (spotlight shining on red wall) so light passes through the wall to the next room. With depth-based (view-dependent) sample visibility enabled, light is effectively attenuated.

## 9 Implementation and Evaluation

The Radiance Hints method has been implemented on a deferred rendering platform, where visible surface albedo, normals and screen-space surface coordinates are available at the global illumination pass. The reflective shadow maps encode the depth, position, normal and radiant flux at the lit surface points (Eq. 4). The method has been tested with Spherical Harmonics of order 2 and 3 and full color encoding ( $3 \times 4$  and  $3 \times 9$  coefficients respectively). The texture requirements for the 2<sup>nd</sup> order SH coefficients are 4 volume textures multiplied by 2 for ping pong rendering. 3<sup>rd</sup> order Spherical Harmonics require 8 volume buffers. In both cases, the multiple render targets mechanism is used in order to fill the radiance hint data in one pass per RSM.

The rest of this section provides an evaluation of the Radiance Hints method in terms of robustness, performance and usability and compares the proposed method against other modern techniques, as necessary. To this end, Table 1 summarizes some key capabilities and requirements of the proposed method (RH), two geometry injection volume-based methods (VBGI1 and VBGI2), RSM sampling/splatting (RSM) and Light Propagation Volumes (LPV). VBGI1 corresponds to the texture atlas point injection technique by Thiedemann et al. [2011], while VBGI2 is the tessellation-based point injection method described in [Mavridis and Papaioannou 2011]. Please refer to the Related Work section for a brief overview of the above techniques. Radiance Hints is not compared against the method by Nijasure et al. because the latter relies extensively on geometry rasterization from multiple points and heavy texture access for the encoding of the radiance field in spherical harmonic coefficients from the individual buffers. This is not a viable solution for large environments, where typically more than 1000 radiance cache points are required.

### 9.1 Robustness

In terms of temporal coherence, aliasing due to moving viewpoint, light source or both is evident in all volume-based techniques. This is usually manifested as an abrupt change in global illumination (view-dependent methods) and flickering. The global illumination result of the basic Radiance Hints method is view-independent, as indirect illumination depends only on the RSM samples. Therefore, the sudden change in the visibility of geometry in the G-buffer does not affect the result in any way. With the depth-based RSM sample attenuation turned on (Section 4), subtle changes in global illumination will occur when geometry enters or exits the camera view frustum. However, except for extreme cases, the illumination changes gradually, due to the voting mechanism that is used for the radiance attenuation.



**Figure 6:** Example of large open area rendering using sparse Radiance Hints ( $32 \times 3 \times 32$ ). Total GI time: 7.2ms (92,782 triangles,  $1024 \times 1024$  pixels, 2<sup>nd</sup> order SH).



**Figure 7:** Light blocking. In contrast to volume-based techniques with full geometry injection, RH and Light Propagation Volumes cannot handle cases where geometry is not visible to the image buffers used. However, in this example, the RSM sample attenuation mechanism of Section 8 effectively blocks erroneous illumination of the interior walls.

For changes in light source position or highly animated geometry, due to the weighted average radiance field that is used for the reconstruction of the incident light, only significant frame-to-frame changes in the RSM depth values can result in flickering. To counter these potential artifacts, temporal smoothing is employed by gradually updating the global illumination buffer in each frame with a blending factor of 0.5 as in [Kaplanyan et al. 2011].

In light propagation techniques such as the Cascaded Light Propagation Volumes method (LPV) [Kaplanyan and Dachsbacher 2010], which is also the fastest method presented so far – radiance is transferred in a wavefront manner across voxel boundaries. This means that the maximum energy exchange distance is directly dependent on the number of propagation iterations within the volume. On the other hand, the Radiance Hints method estimates light transport from arbitrarily distant VPLs in one iteration and inherently supports multiple bounces with a small cost in a predictable and controlled manner, regardless of propagation distance.

Due to the fact that Radiance Hints represent a stochastic evaluation of the radiance field based on RSM samples, the granularity of their distribution in a scene only affects the coarseness of the resulting indirect illumination and enables the sparse covering of large volumes, as demonstrated in Figure 6. This is in contrast to volume-based geometry injection techniques where a poor discretization can lead to an over-evaluated volume occupancy and therefore severe errors.

Method	Input	Secondary bounces	Secondary occlusion	Memory requirements	Performance bound
RSM	image buffers	no	no	RSMs	RSM samples or splatting fill rate
LPV	image buffers	yes, with enough iterations	view-dependent, partial	RSMs + multi-channel volumes	propagation iterations
VBGI1	points, image buffers	no	yes	RSMs + tex. atlases + occupancy volume	number of injected points, volume resolution
VBGI2	triangles	yes	yes	multi-channel volume	volume size, number of triangles
RH	image buffers	yes	probabilistic + optional view-dependent	RSMs + small multi-channel volume	GI buffer resolution (num. of reconstruction samples)

**Table 1:** Comparison table for recent real-time global illumination techniques.

As RSM sampling only regards lit points, indirect light blocking is not taken into account for geometry that is hidden with respect to the RSMs. The distance occlusion technique described in Section 7 partially alleviates the problem for secondary bounces but cannot handle all cases, especially when the RSM footprint on the scene surfaces is small (e.g. sunlight through a small window) or when critical occlusion information is missing from the RSM image buffers. The depth-based RSM sample attenuation mechanism, effectively addresses the first-bounce occlusion in most cases, but introduces dependence on screen-space data. In Fig. 7, a difficult case is presented, where all light comes from the top of the scene, resulting in the side walls being invisible to both the RSMs and the camera. Buffer-based geometry injection fails to block the light reflected from the ground and the same would happen in the case of Radiance Hints, if the RSM sample attenuation was not enabled. Fig. 5 also presents an extreme case that is hard to address with the basic method.


## 9.2 Performance

Some results of the Radiance Hints Method are presented in Fig. 9 and detailed timings for the various stages are shown in Table 2 and Fig. 10. A direct execution time comparison with other methods was not attempted, as code that is not fine-tuned and optimized for a specific platform would not provide accurate results for any of them. Due to the geometry-less representation, the RH evaluation time depends only on the number of RH points and the total number of Reflective Shadow Map samples. As shown in Fig. 10, the choice of Spherical Harmonics order affects the performance of this first stage, because the data for an additional set of 4 render targets must be calculated and written. Once the first-bounce RH radiance field is evaluated, successive bounces add a small fixed overhead per bounce per RH point. The global illumination buffer resolution only affects the last stage, i.e. the irradiance reconstruction as it is performed per visible fragment. Note that even for complex scenes and 3<sup>rd</sup> order Spherical Harmonics, interactive rates are maintained up to a full high-definition resolution.

## 9.3 Parameterization and Engine Integration

With Radiance Hints, there are many ways to gradually balance performance and visual quality. All parameters are intuitive and can be fine-tuned to match the desired result.

Fewer Radiance Hints result in lower radiance field detail and smoother result, while decreasing the RSM samples increases the radiance and distance estimation variance. For scenes with many lights or heavy texturing, where the precise transport direction cannot be visually tracked, 2<sup>nd</sup> order Spherical Harmonics can be used, as GI rendering time is nearly halved (Fig. 10). The number of



**Figure 8:** Impact of various RH parameters to the visual result. Radiance Hint GI was rendered with three light bounces and 40 RSM samples per RH point at 1440 × 900 pixels. A path-traced rendering of is provided for comparison.


Radiance Hint samples acquired during the reconstruction directly trades smoothness for performance. All test in this paper were performed with 4 samples per fragment, arranged in a rotating cone, as explained in Section 5. Furthermore, at high-definition resolutions, rendering the GI at smaller resolution has no significant impact on visual quality, but doubles or triples the performance. Figure 8 provides an example of gradual parameter tuning and a comparison to a reference image produced by path tracing.

In terms of rendering engine interoperability, the Radiance Hints method relies solely on typical reflective shadow maps. No geometric data or additional textures are required. Furthermore, it makes no use of geometry shaders and requires a small memory overhead, as the RH volumes – the only additional resources – are typically small. For instance, for a grid of 32 × 16 × 32 radiance hints, 2<sup>nd</sup> order spherical harmonics and 16 bit floating point textures, 2 × 512KB are needed for the interleaved (ping pong) rendering of GI with multiple indirect light bounces.

## 10 Conclusion and Future Work

The Radiance Hints method combines RSM sampling and sparse geometry-less radiance caching to provide fast and physically-based irradiance estimation with multiple diffuse inter-reflections.





**Figure 9:** Diffuse global illumination examples (2-bounces). (a) Sibenik Cathedral, (b) Cargo containers, (c) Sponza Atrium with people, (d) Cargo bay, (e) A cellar lit by flashlight and (f) Cornell Box with dragon and happy Buddha models.

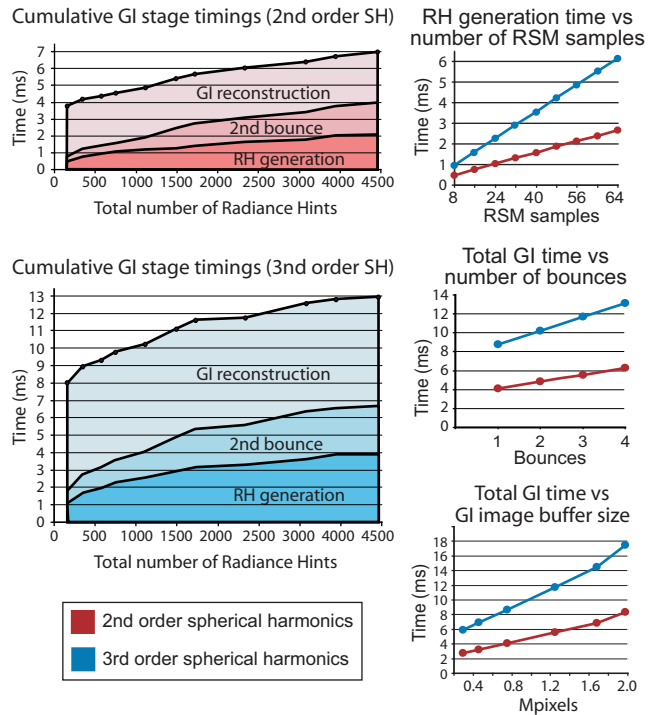
Light transport distance does not affect the method’s performance, as in volume-based methods and the controlling parameters are intuitive, making easy to favor quality over speed and vice versa. The generated radiance field can be exploited in the same manner as in [Kaplanian and Dachsbacher 2010] to render complementary effects such as volumetric lighting. Planned future work will address the issues of unstructured radiance hints and hierarchical sampling of disjoint RH volumes.

## Acknowledgements

The Sponza Atrium model (Crytek version) was downloaded from <http://www.crytek.com/cryengine/cryengine3/downloads> and the Sibenik Cathedral model by Marko Dabrovic was obtained from <http://hdri.cgtechniques.com/~sibenik2>.

## References

DACHSBACHER, C., AND STAMMINGER, M. 2005. Reflective shadow maps. In *13D '05: Proceedings of the 2005 Symposium on Interactive 3D Graphics and Games*, ACM, New York, NY, USA, 203–231.



**Figure 10:** GI timings with respect to various parameters of the Radiance Hints method on an NVIDIA GTX285 card using the Cornell Box scene of Fig 9(f). All measurements refer to 100% image coverage of a  $1024 \times 1024$  GI buffer, 28 RSM samples per RH point and  $16^3$  RH points, unless stated otherwise.

DACHSBACHER, C., AND STAMMINGER, M. 2006. Splatting indirect illumination. In *13D '06: Proceedings of the 2006 Symposium on Interactive 3D Graphics and Games*, ACM, New York, NY, USA, 93–100.

DACHSBACHER, C., STAMMINGER, M., DRETTAKIS, G., AND DURAND, F. 2007. Implicit visibility and antiradiance for interactive global illumination. *ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)* 26, 3 (August).

DONG, Z., GROSCH, T., RITSCHEL, T., KAUTZ, J., AND SEIDEL, H.-P. 2009. Real-time indirect illumination with clustered visibility. In *Vision, Modeling, and Visualization Workshop*.

DURAND, F., HOLZSCHUCH, N., SOLER, C., CHAN, E., AND SILLION, F. X. 2005. A frequency analysis of light transport. *ACM Trans. Graph.* 24 (July), 1115–1126.

GAITATZES, A., MAVRIDIS, P., AND PAPAIOANNOU, G. 2010. Interactive volume-based indirect illumination of dynamic scenes. In *Intelligent Computer Graphics 2010*, vol. 321 of *Studies in Computational Intelligence*. Springer Berlin / Heidelberg, 229–245.

GREGER, G., SHIRLEY, P., HUBBARD, P. M., AND GREENBERG, D. P. 1998. The irradiance volume. *IEEE Computer Graphics and Applications* 18, 32–43.

JENSEN, H. W. 1996. Global illumination using photon maps. In *Proceedings of the eurographics workshop on Rendering techniques '96*, Springer-Verlag, London, UK, 21–30.

Scene	Containers	Cellar	Cargo bay	Sibenik Cathedral	Sponza Atrium	Cornell Box
Triangle count	12,528	15,220	18,297	80,062	267,235	326,383
Volume size	$16 \times 8 \times 16$	$13 \times 5 \times 16$	$16 \times 9 \times 16$	$16 \times 13 \times 7$	$24 \times 11 \times 15$	$10 \times 10 \times 10$
RSM samples per RH	50	20	50	30	30	60
RH estimation time (ms)	1.90 / 3.87	1.25 / 2.81	3.14 / 6.80	1.18 / 1.97	2.26 / 4.13	2.36 / 2.94
Secondary bounce time (ms)	1.55 / 2.70	0.90 / 2.39	1.74 / 2.91	1.23 / 1.63	1.87 / 2.79	0.73 / 1.53
GI reconstruction time (ms)	2.17 / 4.55	2.95 / 6.16	2.95 / 6.16	2.93 / 6.27	2.88 / 6.03	2.57 / 5.58
Total GI time (ms)	<b>5.62 / 11.12</b>	<b>5.10 / 11.36</b>	<b>7.83 / 15.87</b>	<b>5.34 / 9.87</b>	<b>7.01 / 12.95</b>	<b>5.66 / 10.05</b>

**Table 2:** Performance results for 2-bounce indirect diffuse illumination. Global illumination is rendered at full frame resolution ( $1024 \times 1024$  pixels) on an NVIDIA GTX285 graphics card. The two different time measurements correspond to spherical harmonics of order 2 (left) and 3 (right). The scenes are shown in Fig. 9.

- KAJIYA, J. T. 1986. The rendering equation. In *SIGGRAPH '86: Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 143–150.
- KAPLANYAN, A., AND DACHSBACHER, C. 2010. Cascaded light propagation volumes for real-time indirect illumination. In *I3D '10: Proceedings of the 2010 ACM SIGGRAPH symposium on Interactive 3D Graphics and Games*, ACM, New York, NY, USA, 99–107.
- KAPLANYAN, A., ENGEL, W., AND DACHSBACHER, C., 2011. Diffuse global illumination with temporally coherent light propagation volumes, Feb.
- KELLER, A. 1997. Instant radiosity. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 49–56.
- LEUNG, C.-S., WONG, T.-T., LAM, P.-M., AND CHOY, K.-H. 2006. An rbf-based image compression method for image-based rendering. *IEEE Transactions on Image Processing* 15, 4, 1031–1041.
- MAVRIDIS, P., AND PAPAIOANNOU, G. 2011. Global illumination using imperfect volumes. In *GRAPP '11: Proceedings of the international conference on computer graphics theory and applications*.
- MCGUIRE, M., AND LUEBKE, D. 2009. Hardware-accelerated global illumination by image space photon mapping. In *HPG '09: Proceedings of the Conference on High Performance Graphics 2009*, ACM, New York, NY, USA, 77–89.
- NICHOLS, G., AND WYMAN, C. 2010. Interactive indirect illumination using adaptive multiresolution splatting. *IEEE Transactions on Visualization and Computer Graphics* 16, 5, 729–741.
- NIJASURE, M., PATTANAIK, S., AND GOEL, V. 2005. Real-time global illumination on gpus. *Journal of Graphics, GPU, & Game Tools* 10, 2, 55–71.
- RAMAMOORTHY, R., AND HANRAHAN, P. 2001. An efficient representation for irradiance environment maps. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 497–500.
- RITSCHHEL, T., GROSCH, T., KIM, M. H., SEIDEL, H.-P., DACHSBACHER, C., AND KAUTZ, J. 2008. Imperfect shadow maps for efficient computation of indirect illumination. *ACM Trans. Graph.* 27, 5, 1–8.
- RITSCHHEL, T., GROSCH, T., AND SEIDEL, H.-P. 2009. Approximating dynamic global illumination in image space. In *I3D '09: Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games*, ACM, New York, NY, USA, 75–82.
- SHANMUGAM, P., AND ARIKAN, O. 2007. Hardware accelerated ambient occlusion techniques on gpus. In *I3D '07: Proceedings of the 2007 Symposium on Interactive 3D Graphics and Games*, ACM, New York, NY, USA, 73–80.
- SILLION, F. X., ARVO, J. R., WESTIN, S. H., AND GREENBERG, D. P. 1991. A global illumination solution for general reflectance distributions. *SIGGRAPH Comput. Graph.* 25, 4, 187–196.
- SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 527–536.
- SZIRMAY-KALOS, L., ASZODI, B., LAZANYI, I., AND PREMECZ, M. 2005. Approximate ray-tracing on the gpu with distance impostors. *Computer Graphics Forum* 24, 3, 695–704.
- THIEDEMANN, S., HENRICH, N., GROSCH, T., AND MÜLLER, S. 2011. Voxel-based global illumination. In *I3D '11: Proceedings of the symposium on interactive 3D graphics and games 2011*, ACM.
- WALTER, B., FERNANDEZ, S., ARBREE, A., BALA, K., DONIKIAN, M., AND GREENBERG, D. P. 2005. Lightcuts: a scalable approach to illumination. *ACM Trans. Graph.* 24, 3, 1098–1107.
- WANG, R., WANG, R., ZHOU, K., PAN, M., AND BAO, H. 2009. An efficient gpu-based approach for interactive global illumination. In *SIGGRAPH '09: ACM SIGGRAPH 2009 papers*, ACM, New York, NY, USA, 1–8.
- WARD, G. J., RUBINSTEIN, F. M., AND CLEAR, R. D. 1988. A ray tracing solution for diffuse interreflection. *SIGGRAPH Comput. Graph.* 22 (June), 85–92.
- YAO, C., WANG, B., CHAN, B., YONG, J., AND PAUL, J.-C. 2010. Multi-image based photon tracing for interactive global illumination of dynamic scenes. *Computer Graphics Forum* 29, 4, 1315–1324.
- ZHUKOV, S., IONES, A., AND KRONIN, G. 1998. An ambient light illumination model. In *Rendering Techniques - Proc. Eurographics Rendering Workshop*, Springer-Wien, Eurographics, 45–55.