


# HDR Light Probe Sequence Resampling for Realtime Incident Light Field Rendering

Joakim Löw\* Anders Ynnerman Per Larsson Jonas Unger

VITA  
Linköping University, Sweden



**Figure 1:** A sequence of high dynamic range light probes is resampled to a three dimensional grid of radiance maps, which are projected onto spherical harmonics to allow for realtime rendering of incident light fields. The images show realtime rendered geometry with diffuse and glossy materials, lit using spherical harmonic grid representations of the incident light fields.

## Abstract

This paper presents a method for resampling a sequence of high dynamic range light probe images into a representation of *Incident Light Field* (ILF) illumination which enables realtime rendering. The light probe sequences are captured at varying positions in a real world environment using a high dynamic range video camera pointed at a mirror sphere. The sequences are then resampled to a set of radiance maps in a regular three dimensional grid before projection onto spherical harmonics. The capture locations and amount of samples in the original data make it inconvenient for direct use in rendering and resampling is necessary to produce an efficient data structure. Each light probe represents a large set of incident radiance samples from different directions around the capture location. Under the assumption that the spatial volume in which the capture was performed has no internal occlusion, the radiance samples are projected through the volume along their corresponding direction in order to build a new set of radiance maps at selected locations, in this case a three dimensional grid. The resampled data is projected onto a spherical harmonic basis to allow for realtime lighting of synthetic objects inside the incident light field.

**CR Categories:** I.3.6 [Computer Graphics]: Methodology and Techniques—Graphics data structures and data types

**Keywords:** Image Based Lighting, Incident Light Fields, Spherical harmonics

## 1 Introduction

Photo-realistic computer graphics renderings, that cannot be distinguished from photographs, are a fundamental part in many application areas ranging from virtual reality to computer games and special effects in movies. By employing physically based material and lighting models virtual objects can be seamlessly placed into real world environments. For this reason methods for capturing real world scene illumination have been developed. Such techniques are generally known as *Image Based Lighting* (IBL), [Debevec 1998; Sato et al. 1999], and have been successfully employed in industry pipelines. In the traditional IBL technique, scene illumination is measured as an angular radiance distribution at a single position in space. This means that traditional IBL cannot capture or recreate effects such as cast shadows, shafts of light or other important spatial variations in the illumination.

Facilitated by recent developments in camera sensor technology and using light field imaging, traditional IBL techniques have been further developed to capture the full spatial and angular dimensionality of the scene illumination. This is usually done by capturing sequences of illumination measurements located at many positions in space. Such captured real world lighting data set is called an *Incident Light Field* (ILF) [Unger et al. 2003]. This generalization of IBL enables the production of highly realistic renderings using real world lighting environments exhibiting complex spatial variations with preserved parallax. The proposed ILF techniques have mainly

---

\*e-mail:Joakim.Low@itn.liu.se

been directed towards off-line rendering or limited to environments with single light sources. The main difficulties in ILF rendering are the large data size of captured illumination and efficient random access during evaluation of the rendering integral at each point in the scene.

To extend the general ILF technique towards realtime applications, such as computer games, preview rendering for scene modeling and lighting setup software, there is a need for efficient data structures and novel rendering algorithms. Therefore this paper proposes a technique where a sequence of light probe images, captured at different locations in the scene, is used to generate a 3D grid of illumination information describing the spatial and angular variations in the captured scene lighting. At each point in the grid, the illumination is projected onto a *spherical harmonic* (SH) basis which can then be efficiently used to evaluate the interaction between the scene lighting and diffuse and glossy materials. Although the basis can be chosen arbitrarily, we use spherical harmonics since their properties are well explored in computer graphics applications. This allows a straightforward integration of ILF lighting into most rendering frameworks.

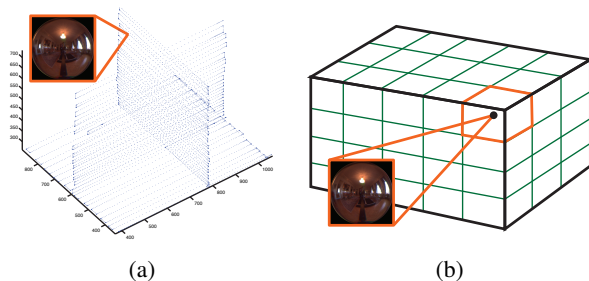
## 2 Related work

The methods used for resampling and realtime rendering using spatially varying real world illumination data sets relates to work in several different computer graphics areas, here divided into: *Image Based Lighting*, *Incident Light Fields* and *Realtime rendering*. This section gives an overview of the work upon which the proposed technique builds.

**Image Based Lighting** - *Image Based Lighting* (IBL) was introduced by Debevec [1998] and is a technique for creating photo-realistic renderings of synthetic objects. The objective of IBL is to measure and recreate the *plenoptic function*,  $P(\mathbf{x}, \vec{\omega}, \lambda, t)$ , introduced by Adelson and Bergen [1991]. This function describes the illumination incident at any point,  $\mathbf{x}$ , from any direction,  $\vec{\omega}$ , of any wavelength,  $\lambda$ , at any time,  $t$ . By letting time be constant, and considering three bands of wavelengths separately ( $r, g, b$ ), the plenoptic function is reduced to a 5D form,  $P(\mathbf{x}, \vec{\omega})$ .

In IBL, real world scene illumination is measured by capturing omni-directional *High Dynamic Range* (HDR) images. This is usually done using wide angle optics such as fisheye-lenses or by capturing the environment through the reflection in a mirror sphere. This kind of angular radiance measurement captured at a single position in space is called a *light probe* or *radiance map*. For an overview of HDR imaging and IBL see Reinhard et al. [2006]. Each pixel,  $l$ , in the light probe image corresponds to a radiance measurement,  $\hat{L}(\vec{\omega}_q)$ , incident at the capture location from a direction,  $\vec{\omega}_q$ . This information is used to compute the light incident upon synthetic objects during a rendering, giving them the appearance of being present in the real world scene. Under the assumption that any geometry in the environment captured by the light probe is far from the synthetic scene, each radiance sample,  $\hat{L}(\vec{\omega}_q)$ , can be considered a valid radiance measurement for the direction  $\vec{\omega}_q$  of any point in the synthetic scene.

**Incident Light Fields** - Traditional IBL techniques use illumination,  $L(\vec{\omega})$ , captured at a single point only. This means that it cannot capture or recreate effects such as cast shadows, shafts of light or other spatial variations that are important for the visual appearance of the scene. By combining IBL with *light field* imaging, introduced in computer graphics by Levoy and Hanrahan [1996] and Gortler et al [1996], Unger et al. [2003] presented a technique for



**Figure 2:** An example ILF data set (a). The cloud of points, in this case roughly layed out in three planes, displays the position of each captured light probe image. Each pixel in each image represents an individual radiance sample. This data set is resampled to a regular 3D grid of local radiance maps (b) which are used for spherical harmonic projection.

capturing and rendering with spatially varying real world illumination. They called such 4D radiance data sets *Incident Light Fields* (ILF). Based on similar ideas, Masselus et al. [2003] later captured reflectance fields in a way such that relighting of real objects with 4D ILFs could be performed. Capture and representation of ILFs was further improved by Unger et al. [2008], showing how to build efficient and adaptive data structures taking advantage of the 4D nature of ILFs, without significant loss of quality, as well as how to extract and separate the light sources of the captured scene. Ihrke et al. [2008] proposed an ILF capture technique using a static HDR video camera and a handheld planar mirror. They also proposed a method for resampling the unstructured measurements of the light field into a directional representation, creating an orthographic view of the light field for each direction.

**Realtime rendering** - The benefit of using real world illumination has motivated research and development of techniques for realtime IBL rendering. Ramamoorthi and Hanrahan [2001a] projected the captured illumination,  $L(\vec{\omega})$ , onto *Spherical Harmonics*. They showed that diffuse transfer could be accurately simulated by representing the environment illumination with a nine component spherical harmonic. Sloan et al. [2002] later extended the use of spherical harmonics for *Precomputed Radiance Transfer* (PRT) of glossy materials, performing expensive computations as a precomputation step.

Many rendering applications need to represent a spatial variance of the illumination in a scene, at the same time requiring the representation to be fast enough to produce realtime framerates. Greger et al. [1998] presented a method to represent spatially variant illumination called *Irradiance Volumes*, storing irradiance maps in adaptive 3D grids in the scene. Oat [2006] extended this technique using spherical harmonics to represent the irradiance maps, yielding an extremely compact representation of the spatially variant illumination. Using spherical harmonic grids to represent volumetric data has found its way into other applications as well. Moon et al. [2008] used the technique to represent scattered radiance in hair to speed up physically realistic hair renderings.

## 3 ILF representation for realtime rendering

In this work an ILF is captured as a sequence of *light probes*, omnidirectional radiance maps, measured at a large set of different locations in the scene. The key idea for achieving realtime rendering performance is to resample the captured sequences to a regular

3D grid of radiance maps. The purpose of resampling is to generate a data structure well suited for efficient lookup and lighting calculations during rendering.

The sampled light probes, each consisting of a large set of radiance samples, are captured in a spatial volume assumed to contain no occluding geometry, effectively sampling the 4D ILF of the volume. Representing this inherently 4D data set by a 5D structure is motivated by efficiency requirements. The data set often consists of billions of radiance samples captured at irregularly spaced points, meaning the memory requirements are in the order of Gb, resulting in time consuming random access of the radiance samples. The computation of material response to incident illumination during rendering therefore presents a significant problem. Building upon the ideas of irradiance volumes [Greger et al. 1998] and the extension using spherical harmonics [Oat 2006], we solve the problem by projecting each radiance map of the resampled 3D grid onto spherical harmonic basis functions. This representation has a small memory footprint and allows for efficient rendering of diffuse and glossy objects inside the ILF volume.

The input to our algorithm is a large set of light probes along with tracking information describing where each probe was captured. An example of a light probe distribution is given in Figure 2. A pixel,  $q$ , in image  $p$  of the sequence corresponds to a radiance measurement,  $\hat{L}(\mathbf{x}_p, \vec{\omega}_q)$ , captured at a position,  $\mathbf{x}_p$ , observed in the direction  $\vec{\omega}_q$ . The objective of the algorithm is to resample the full set of radiance measurements,  $\hat{L}(\mathbf{x}_p, \vec{\omega}_q)$ , to a regular grid of local radiance maps,  $L(\mathbf{x}_{p'}, \vec{\omega})$ . The local radiance map at each grid point,  $\mathbf{x}_{p'}$ , is then projected to spherical harmonics. Because of the method used for the projection, the entire radiance map of each grid point never needs to be stored explicitly, reducing the total memory footprint during resampling considerably. The algorithm consists of four steps:

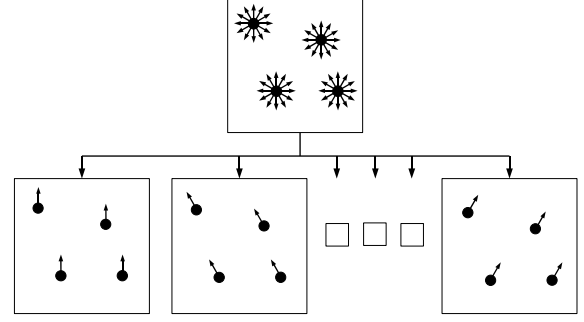
**Sorting** In order to reduce the demand on memory during the resampling process, all radiance samples,  $\hat{L}(\mathbf{x}_p, \vec{\omega}_q)$ , are sorted into buckets based on their direction,  $\vec{\omega}_q$ , as shown in Figure 3. In other words, a bucket is created for all samples extracted from the same corresponding pixel in the light probe images. Sorting all samples on direction will allow us to work through them in a sequential manner, one direction at a time, without exhausting main memory or causing large amounts of memory shuffling for large data sets. Using the unsorted samples would either cause a lot of memory shuffling, or make it necessary to keep a full local radiance map at each grid point, requiring huge amounts of memory.

**Radiance sample projection** For each bucket, all corresponding radiance samples are projected through the volume, adding a contribution to the local radiance maps,  $L(\mathbf{x}_{p'}, \vec{\omega})$ , at the grid points,  $\mathbf{x}_{p'}$ .


**Repair** Depending on the spatial distribution of the samples from a given bucket, there may be grid points that do not get a contribution for the corresponding direction, leaving the radiance maps incomplete. If this is the case, a repair step is run, adding directional information to empty grid points from nearby grid points.

**Spherical harmonic projection** When each grid point has received a sample for the current direction, the information is used to update the computation of the spherical harmonic projection, before proceeding to the next bucket of samples.

Below the three last steps of the algorithm are described in detail. Note that the last step, spherical harmonic projection, is explained first, in order to explain why there is no need for storing the complete resampled radiance maps.



**Figure 3:** Each light probe represents a set of radiance samples, as illustrated in the top image. A radiance sample is denoted as  $\hat{L}(\mathbf{x}_p, \vec{\omega}_q)$ , and corresponds to a position  $\mathbf{x}_p$ , where the light probe was captured, and a direction  $\vec{\omega}_q$ , which corresponds to a specific pixel in the light probe. The radiance samples are sorted into buckets based on their direction, to allow for efficient access.



**Figure 4:** Illustration of how one of the radiance samples,  $\hat{L}(\mathbf{x}_p, \vec{\omega}_q)$ , in a bucket is used to add radiance to the local radiance maps,  $L(\mathbf{x}_{p'}, \vec{\omega})$ , at the grid points,  $\mathbf{x}_{p'}$ , for the direction  $\vec{\omega}_q$ . The same procedure is repeated for all radiance samples in the bucket, which contains all samples having direction  $\vec{\omega}_q$ , one from each sampled light probe. The radiance samples are projected in both directions along  $\vec{\omega}_q$  from their capture point,  $\mathbf{x}$ .

### 3.1 Spherical harmonic projection

In this work real-valued spherical harmonics are used. Definitions and details of the properties of the spherical harmonics basis functions can be found in, for example, [Ramamoorthi and Hanrahan 2004; Green 2003]. The projection of a function,  $f(\vec{\omega})$ , onto the set of spherical harmonic basis functions,  $Y_{lm}(\vec{\omega})$ ;  $l \geq 0, -l \leq m \leq l$ , is expressed as an integral of the product between  $f(\vec{\omega})$  and the individual basis functions over the sphere. For a given local radiance map,  $L(\mathbf{x}_{p'}, \vec{\omega})$ , the coefficients,  $c_{lm}^L$ , of the corresponding spherical harmonic expansion are given by

$$c_{lm}^L = \int_{S^2} L(\mathbf{x}_{p'}, \vec{\omega}) Y_{lm}(\vec{\omega}) d\vec{\omega} \quad (1)$$

yielding the expansion

$$L(\mathbf{x}_{p'}, \vec{\omega}) \approx \sum_{l=0}^n \sum_{m=-l}^l c_{lm}^L Y_{lm}(\vec{\omega}) \quad (2)$$

for the  $n$ -degree spherical harmonic approximation of the map. The integral (1) is evaluated numerically using

$$c_{lm}^L = \sum_{j=1}^N w(\vec{\omega}_j) \hat{L}(\mathbf{x}_{p'}, \vec{\omega}_j) Y_{lm}(\vec{\omega}_j) \quad (3)$$

where  $w(\vec{\omega}_j)$  is a weighting function that is determined by the sample distribution,  $N$  is the number of samples and  $\hat{L}(\mathbf{x}_{p'}, \vec{\omega}_j)$  is

the sample of  $L(\mathbf{x}_{p'}, \vec{\omega})$  in the direction  $\vec{\omega}_j$ . Since mirror sphere mapping is used, all radiance samples correspond to equal sized solid angles in the reflected environment, giving a constant weight  $w(\vec{\omega}_j) = \frac{4\pi}{N}$ .

The sum (3) does not require the complete radiance map to be available at any time, all that is needed is to keep track of the partial sum and update it when samples become available. This means the projection to spherical harmonics can be performed for all grid points in parallel, keeping only the partial sum of (3) for each spherical harmonic coefficient at each grid point in memory at once. After the radiance sample projection and repair steps, radiance samples for the current direction will be available for all grid points. These samples are used to update the current partial sums, and are then thrown away, resulting in very low demand on memory.


### 3.2 Radiance sample projection

In order to perform the spherical harmonic projection, samples of the local radiance map,  $L(\mathbf{x}_{p'}, \vec{\omega})$ , are needed at each grid point,  $\mathbf{x}_{p'}$ , in the volume. To obtain these samples, each of the original captured radiance samples,  $\hat{L}(\mathbf{x}_p, \vec{\omega}_q)$ , is projected through the volume along the rays,  $\vec{R}_{pq}(s) = \mathbf{x}_p + s \cdot \vec{\omega}_q$ , as shown in Figure 4. Under the assumption that the volume has no occluding geometry, the rays carrying the radiance will travel unoccluded through the entire volume, from entry point to exit point. Therefore the sample can be projected in both directions along the ray. For each bucket, the projection of the full corresponding sample set allows for reconstruction of the radiance in the local radiance maps for the current direction,  $\vec{\omega}_q$ . The reconstruction is performed by applying a Gaussian filter kernel at each grid point,  $\mathbf{x}_{p'}$ . All rays,  $\vec{R}_{pq}(s)$ , that pass through the support of the kernel at a grid point will contribute to the local radiance map,  $L(\mathbf{x}_{p'}, \vec{\omega})$ , in the direction  $\vec{\omega}_q$ .


### 3.3 Repair

When all radiance samples for a bucket have been projected there may be points in the grid missing a contribution for the corresponding direction. Since there will be no information about radiance in neighbouring directions in the local radiance maps, the radiance samples for the incomplete grid points must, instead, be created using radiance information for the current direction in nearby grid points. To this end, a mip-mapping technique is used. The number of buckets (radiance sample directions) are high, for the examples used in this paper numbering in the tens of thousands. A repair step may need to be performed for every bucket, and mip-mapping provides a computationally efficient approach to fill out missing data.

Having grid points missing information, the mip-map technique starts by building a grid half the resolution of the original one, using a box filter on the information available. Grid points not having information will not contribute in any way to the values of the level below. The method then tries to sample the new lower resolution grid to reconstruct samples for the top level grid. If this reconstruction is not possible for one or more of the grid points in the top level, a new level of the mip-map is built, and this procedure is repeated until all grid points have been assigned a sample. The technique is illustrated in Figures 5 and 6.



**Figure 5:** A 2D illustration of the mip-mapping technique used to fill out missing data in the grid. Each level is created using the information available in the level above.



**Figure 6:** Using the mip-mapping technique to fill out missing data in the grid, a trilinear interpolation of the highest level of the mipmap having information about the neighbourhood of the empty grid point is used.

## 4 Rendering using a 3D spherical harmonic grid

After the spherical harmonic grid has been completed, it is ready to be used for lighting. To compute the outgoing radiance from a surface point the *rendering equation* [Kajiya 1986], with the emission term omitted, is used:

$$B(\mathbf{x}, \vec{\omega}_o) = \int_{S^2} L(\mathbf{x}, \vec{\omega}_i) \rho(\mathbf{x}, \vec{\omega}_i, \vec{\omega}_o) d\vec{\omega}_i \quad (4)$$


where  $S^2$  is the unit sphere in  $\mathbb{R}^3$ ,  $B(\mathbf{x}, \vec{\omega}_o)$  is the outgoing radiance from surface point  $\mathbf{x}$  in direction  $\vec{\omega}_o$ ,  $L(\mathbf{x}, \vec{\omega}_i)$  is radiance incident to the surface point from direction  $\vec{\omega}_i$ , and  $\rho(\mathbf{x}, \vec{\omega}_i, \vec{\omega}_o)$  is the transfer function, describing how incident radiance is reflected and distributed to outgoing radiance. The coefficients of the spherical harmonic approximation to  $\rho(\mathbf{x}, \vec{\omega}_i, \vec{\omega}_o)$  will be computed for the variable  $\vec{\omega}_i$ , that is the weighting of incident radiance will be represented using spherical harmonics. For this reason, the coefficients need to be recomputed every time  $\vec{\omega}_o$  or  $\mathbf{x}$  changes. Due to orthonormality of the spherical harmonic basis functions, the integral (4), when computed using the spherical harmonic approximations of  $L(\mathbf{x}, \vec{\omega}_i)$  and  $\rho(\mathbf{x}, \vec{\omega}_i, \vec{\omega}_o)$ , reduces to a scalar product between coefficient vectors:

$$\int_{S^2} L(\mathbf{x}, \vec{\omega}_i) \rho(\mathbf{x}, \vec{\omega}_i, \vec{\omega}_o) d\vec{\omega}_i \approx \sum_{l=0}^n \sum_{m=-l}^l c_{lm}^L(\mathbf{x}) c_{lm}^{\rho}(\mathbf{x}, \vec{\omega}_o) \quad (5)$$

The coefficients,  $c_{lm}^L(\mathbf{x})$ , which depend on spatial position, represent incident radiance at point  $\mathbf{x}$  and are computed using trilinear interpolation in the spherical harmonic grid representation of the incident light field.  $c_{lm}^{\rho}(\mathbf{x}, \vec{\omega}_o)$  still needs to be determined, and depends on the material properties of the geometry.

Diffuse transfer has no dependence on  $\vec{\omega}_o$  and is computed using  $\rho_D(\mathbf{x}, \vec{\omega}_i, \vec{\omega}_o) = \max(0, \vec{\omega}_i \cdot \vec{n}(\mathbf{x}))$ , where  $\vec{n}(\mathbf{x})$  is the surface normal at point  $\mathbf{x}$ . This function can be well estimated using only a nine





**Figure 7:** Translation stage with motion tracking. This stage is used to capture HDR light probe sequences within a volume of up to  $1.5\text{m} \times 1.5\text{m} \times 1.5\text{m}$  [Unger et al. 2008]

coefficient (degree 2) spherical harmonic approximation, which is fairly simple to compute [Ramamoorthi and Hanrahan 2001a; Ramamoorthi and Hanrahan 2001b].


For glossy transfer, a Phong transfer function is used:

$$\rho_p(\mathbf{x}, \vec{\omega}_i, \vec{\omega}_o) = (\max(0, \vec{\omega}_i \cdot \vec{r}(\vec{\omega}_o, \vec{n}(\mathbf{x})))^s)$$

where  $\vec{r}(\vec{\omega}_o, \vec{n}(\mathbf{x}))$  is the reflection vector of  $\vec{\omega}_o$  with respect to the surface normal  $\vec{n}(\mathbf{x})$ , and  $s$  is the Phong exponent determining the shininess of the surface. The spherical harmonic coefficients of a Phong lobe can be computed as a number of *Zonal Harmonics*, a subset of the spherical harmonic basis functions which are circularly symmetric with respect to one axis. The analytical expression for the Phong lobe coefficients is given by Ramamoorthi and Hanrahan [2004]. The coefficients are computed in a local space, having the reflection vector  $\vec{r}$  as the vector around which we have circular symmetry, and need to be rotated to global space before (5) can be computed. Rotation of zonal harmonics is easier than rotating a full set of spherical harmonic basis functions, and is described by Sloan et al. [2005]. The result of the rotation will be a full set of spherical harmonic coefficients, representing the Phong lobe in global space, ready to use for computation of (5). For low-frequency glossy transfer, a 25 coefficient (degree 4) spherical harmonic approximation to the glossy transfer function yields satisfying results [Sloan et al. 2002], and for the purposes here, a 25 coefficient approximation to a Phong model of low shininess ( $s \leq 6$ ) is used.

## 5 Results and discussion

To validate the method, we have used two different lighting setups, as illustrated in Figure 8. The first scene, illustrated in Figure 8(a), contains two principal light sources: an area light in the ceiling and a projector spotlight projecting a high frequency pattern onto the scene (the four colored squares). The second scene, Figure 8(b), contains two main sources of illumination: one red spotlight that is blocked to illuminate roughly half of the scene and a projector creating the cone of light in the middle of the scene. It also contains a small amount of ambient illumination due to stray light.



**Figure 8:** Illustration of the lighting setup for the two sampled test scenes. The grey planes are drawn to clarify the shape of the illumination, and are not present during the actual sampling of the scenes. The first scene (a) contains an area light source in the ceiling and a light projector with a color filter to produce a colored rectangular pattern. The second scene (b) contains a red spotlight blocked to illuminate only half of the scene and a projector to create a sharp cone of white light.

### 5.1 Illumination capture

The HDR light probe sequences were captured using the technique described by Unger et al. [2008]. They capture light probe image sequences using an HDR video camera (a Ranger C55 from Sick IVP) pointed at a mirror sphere (Figure 7). The camera and sphere are mounted onto a  $1.5\text{m} \times 1.5\text{m} \times 1.5\text{m}$   $xyz$ -translation stage equipped with opto-mechanical tracking units exhibiting sub-millimeter accuracy. The translation stage keeps the orientation of each light probe in the sequence fixed, and thus maintains a common angular quantization for all images. The image of the mirror sphere exhibits a dead angle of  $4.2^\circ$  for directions directly behind the sphere. For the spherical harmonic projection described in this paper, the light incident from this region is assumed to be of intensity zero.

The illumination in the first scene was captured using about twenty-six thousand HDR light probe images, and the second scene using about twenty-nine thousand light probes. To reduce the computational costs and memory footprint, the sequences were adaptively sub-sampled to contain 4397 and 8454 images respectively. This sub-sampling was performed by comparing the energy in the absolute difference between adjacent light probe images. A certain image is considered redundant, and can be removed from the sequence if it is similar to the previous and the next image in the sequence. To ensure the sampling density does not drop below some minimum, the capture positions of the images are also compared to make sure an image is only removed if it is close to its neighbours. A similar reduction technique was described in Unger et al. [2008].

### 5.2 Sorting and resampling

The radiance samples corresponding to each image in the sequences are sorted on their angle,  $\vec{\omega}_q$ , and stored in angular buckets, as described in Section 3. During this process we also, as suggested by Unger et al. [2007], compensate for the non-single viewpoint inherent in the mirror sphere optical setup. An analysis of the optical characteristics of a mirror sphere can also be found in [Swaminathan et al. 2006]. The sorting and resampling were implemented in C++ and performed as single threaded processes. Sorting was run on a 3.00GHz Intel Pentium D workstation, while the resampling was performed on a 2.83GHz Intel Core 2 Quad workstation.

The sorting took about 45 minutes for the 4397 light probe images from the first test scene and about 1 hour and 45 minutes for the 8454 images from the second test scene. The resolution of the light probes images were  $250 \times 250$  pixels. The resampling (the ray-marching and spherical harmonics projection steps) took about 4

	#Probes	SH grid resolution	#SH coeff.	SH grid memory
Test Scene 1	4397	$64^3$	25	75 MB
Test Scene 2	8454	$32^3$	25	9.4 MB

**Table 1:** Number of probes and SH grid resolution for the two test scenes. SH grid memory footprint takes into account three color channels and storing the data as 32-bit floats. On GPU the coefficients are stored as 16-bit floats, effectively cutting the size of the data in half.

hours for the first test scene using a grid resolution of  $64^3$ , and about 1 hour for the second test scene using a grid resolution of  $32^3$ . The filter kernel used when adding radiance contributions to grid points was given a fixed support with a diameter of 3 grid cells. The number of light probe images, grid resolutions and memory footprints for the two scenes are displayed in Table 1.

It should be noted that the sorting step can be avoided completely. Currently, the output HDR stream from the camera units is stored as sequences of light probe images. The output stream could instead be stored in angularly sorted fashion directly, and the multiple viewpoint correction could be performed on the fly during the ray-marching step.

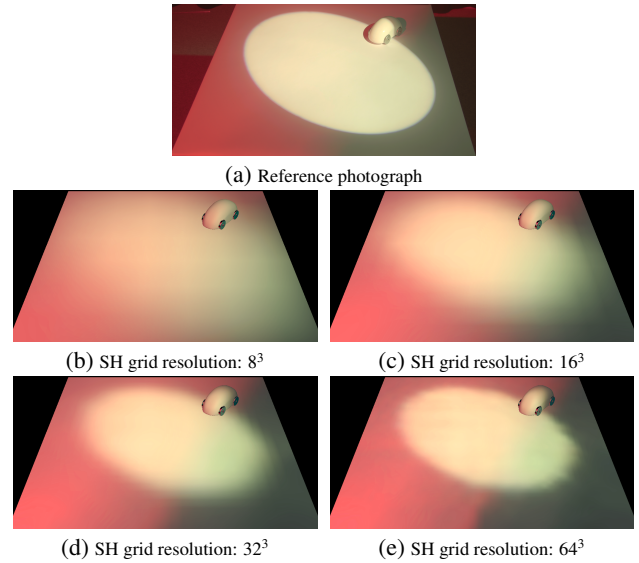
Figure 9 displays the impact of the spatial resolution of the spherical harmonic grid on the rendered result. For comparison, 9(a) displays a reference photograph of an  $80\text{cm} \times 80\text{cm}$  table and a toy car placed in the lighting setup of the second test scene. For the renderings in 9(b,c,d,e) the grid resolutions are set to  $8^3$ ,  $16^3$ ,  $32^3$  and  $64^3$  voxels. In comparison to the physical scene, in which the ILF was captured within a volume of  $1\text{m} \times 1\text{m} \times 1\text{m}$ , these resolutions correspond to a voxel size of 12.5cm, 6.25cm, 3.125cm and 1.5625cm respectively. 9(b,c) shows that the spatial variations are preserved, though heavily blurred, in the low resolutions,  $8^3$  and  $16^3$ . This spatial blurring is expected since the grid clearly undersamples the underlying ILF data.

The highest resolution grid  $64^3$  displayed in 9(e) shows sample artifacts. This is due to the fact that the spatial sample density of the sub-sampled light probe sequence was matched to the  $32^3$  grid displayed in 9(d). Although the  $64^3$  grid produces a sharper result it cannot recreate the under-sampled features in the ILF. This means that if a part of the ILF having high contrasting features is undersampled, unwanted local blur may appear due to the smoothing effects of the mip-mapping technique used to fill out missing data. To reduce this effect, other repair methods could be tested, possibly techniques based on bilateral filtering.

Figure 11 displays the impact of the grid resolution in a scene with more geometry. It should be noted that the angular resolution, that is the degree of the spherical harmonics, is kept the same for the different spatial resolutions, and that the blur is due to the spatial under-sampling in the grid.

### 5.3 Rendering

For realtime rendering using the ILF spherical harmonic grids, a GLSL implementation supporting both diffuse and glossy (Phong) materials was run on a GeForce 8800 GTX. The spherical harmonic coefficients were stored as 16-bit floating point 3D textures on the graphics card and sampled using the built-in trilinear interpolation functionality of the GPU. During rendering, the spherical harmonic coefficients of the Phong lobe are recomputed for every fragment on the GPU, including rotation of the Phong lobe to global space.



**Figure 9:** Comparison between a photograph (a) taken from the second test lighting setup with a table and a toy car, and four synthetic images (b)-(e), with successively increasing spherical harmonic grid resolution.


The generated ILF spherical harmonic grids were used as the source of illumination for the renderings shown in Figures 1, 9, 10 and 11.

Figure 1 shows the Stanford bunny and dragon rendered in realtime, lit using the spherical harmonic grid representations of the ILFs of the two test scenes. 1(a) used a spherical harmonic grid, of resolution  $64^3$ , built from the illumination captured in the first test scene (see Figure 8(a)). 1(b,c) used a spherical harmonic grid, of resolution  $32^3$ , built from the illumination captured in the second test scene (see Figure 8(b)). The plane has a diffuse material in all three images, while the bunny and dragon geometries have both diffuse *and* glossy material. The image resolution was set to  $1024 \times 768$  pixels and the framerates averaged between 150 and 200 fps for the setups shown in the images.


Figure 10 shows six renderings, 10(a,b,c) using the  $64^3$  grid built from the illumination captured in the first test scene and 10(d,e,f) using the  $32^3$  grid from the second test scene. The image resolution in each rendering was set to  $1024 \times 768$  pixels. For the renderings using only diffuse materials, 10(a,b,d,e), the framerate averaged approximately 350 fps, while using a glossy material for the dragon geometry, 10(c,f), resulted in an average framerate of approximately 150 fps.

The lower performance of glossy rendering is due to the relative complexity of the shader performing both Phong coefficient computation and rotation. Exploring the use of precomputed coefficients stored in lookup textures would provide interesting comparisons, and could yield shaders producing higher framerates.

Using spherical harmonics of low degree limits the angular resolution of the lighting, which could be improved not only by increasing the degree of the spherical harmonics, but possibly also by using other sets of basis functions such as wavelets [Ng et al. 2003]. Using spherical harmonics for HDR lighting may yield problems with ringing (Gibbs phenomenon), which can be reduced by windowing [Sloan et al. 2002]. A higher degree of spherical harmonic approximations of transfer functions may also be desirable if there is a higher demand on numerical accuracy, see for example [Mahajan et al. 2008] for a discussion on the accuracy of different spherical harmonic approximations to the Phong transfer function.



**Figure 10:** Results using spherical harmonic grid representations of the incident light fields for the two test scenes used in this paper. Images (a)-(c) are rendered using the spherical harmonic grid of the first test scene, (d)-(f) are rendered using the grid of the second test scene. (a) and (d) shows a flat surface for reference, (b) and (e) shows geometry with diffuse material, and (c) and (f) shows geometry with glossy material. The resolution of the grid for the first test scene is  $64^3$  and for the second test scene  $32^3$ .



**Figure 11:** Images rendered using spherical harmonic grid representations of ILF in second test scene having successively increasing resolution, starting at a low resolution of  $8^3$  in (a),  $16^3$  in (b) and finally a resolution of  $32^3$  in (c).

Using grids to store the spherical harmonics creates higher demands on memory as the resolution increases. Other data structures for the spherical harmonic volume data should be tested to use memory more efficiently. One obvious choice would be to use adaptive grids or octree structures, similar to [Greger et al. 1998] or [Oat 2006]. Care should be taken not to lose the ability to take advantage of the processing power of modern GPUs.

## 6 Conclusion and Future Work

This paper presented an efficient approach to resampling a sequence of light probes to a 3D grid of local radiance maps. It also introduced the idea of using spherical harmonic grids for the application of representing incident light fields, which enables realtime rendering of an ILF on the GPU.

Some discussion and future work has already been presented in the previous section, but, to summarize, more work should be done to explore how sample density and distribution affects the results and

limits the spatial and angular resolution of the spherical harmonic grid. Other structures for storing the spherical harmonics should be tested, since memory will be a crucial factor for larger scenes.

## Acknowledgements

We gratefully acknowledge Matthew Cooper for proofreading the text, Stefan Gustavson for helpful input and discussions, and the anonymous reviewers for valuable comments and suggestions. The project was supported by the Swedish Foundation for Strategic Research through the Strategic Research Center MOVIII.

## References

ADELSON, E. H., AND BERGEN, J. R. 1991. *Computational Models of Visual Processing*. MIT Press, Cambridge, Mass., ch. 1. The Plenoptic Function and the Elements of Early Vision.

- DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH '98: SIGGRAPH 98 Papers*, ACM, New York, NY, USA, 189–198.
- GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. 1996. The lumigraph. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 43–54.
- GREEN, R. 2003. Spherical harmonic lighting: The gritty details. *Proceedings of the Game Developers Conference* (March).
- GREGER, G., SHIRLEY, P., HUBBARD, P. M., AND GREENBERG, D. P. 1998. The irradiance volume. *IEEE Comput. Graph. Appl.* 18, 2, 32–43.
- IHRKE, I., STICH, T., GOTTSCHLICH, H., MAGNOR, M., AND SEIDEL, H.-P. 2008. Fast incident light field acquisition and rendering. *Journal of WSCG (WSCG'08)* 16, 1-3, 25–32.
- 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.
- LEVOY, M., AND HANRAHAN, P. 1996. Light field rendering. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 31–42.
- MAHAJAN, D., TSENG, Y.-T., AND RAMAMOORTHI, R. 2008. An analysis of the in-out brdf factorization for view-dependent relighting. *Computer Graphics Forum* 27, 4 (June), 1137–1145.
- MASSELUS, V., PEERS, P., DUTRÉ, P., AND WILLEMS, Y. D. 2003. Relighting with 4d incident light fields. *ACM Trans. Graph.* 22, 3, 613–620.
- MOON, J. T., WALTER, B., AND MARSCHNER, S. 2008. Efficient multiple scattering in hair using spherical harmonics. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, ACM, New York, NY, USA, 1–7.
- NG, R., RAMAMOORTHI, R., AND HANRAHAN, P. 2003. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph.* 22, 3, 376–381.
- OAT, C. 2006. Irradiance volumes for real-time rendering. In *ShaderX5: Advanced Rendering Techniques*, W. Engel, Ed. Charles River Media, Cambridge, MA, ch. TBD, TBD.
- RAMAMOORTHI, 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.
- RAMAMOORTHI, R., AND HANRAHAN, P. 2001. On the relationship between Radiance and Irradiance: Determining the illumination from images of a convex Lambertian object. *Journal of the Optical Society of America* 18, 10 (Oct), 2448–2459.
- RAMAMOORTHI, R., AND HANRAHAN, P. 2004. A signal-processing framework for reflection. *ACM Trans. Graph.* 23, 4, 1004–1042.
- REINHARD, E., WARD, G., PATTANAİK, S., AND DEBEVEC, P. 2006. *High Dynamic Range Imaging – Acquisition, Display and Image-Based Lighting*. Morgan Kaufmann, San Francisco, CA.
- SATO, I., SATO, Y., AND IKEUCHI, K. 1999. Acquiring a radiance distribution to superimpose virtual objects onto a real scene. *IEEE Transactions on Visualization and Computer Graphics* 5, 1 (January-March), 1–12.
- 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.
- SLOAN, P.-P., LUNA, B., AND SNYDER, J. 2005. Local, deformable precomputed radiance transfer. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, ACM, New York, NY, USA, 1216–1224.
- SWAMINATHAN, R., GROSSBERG, M. D., AND NAYAR, S. K. 2006. Non-Single Viewpoint Catadioptric Cameras: Geometry and Analysis. *International Journal of Computer Vision* 66, 3 (Mar), 211–229.
- UNGER, J., WENGER, A., HAWKINS, T., GARDNER, A., AND DEBEVEC, P. 2003. Capturing and rendering with incident light fields. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 141–149.
- UNGER, J., GUSTAVSON, S., AND YNNERMAN, A. 2007. Spatially varying image based lighting by light probe sequences: Capture, processing and rendering. *Visual Computer* 23, 7, 453–465.
- UNGER, J., GUSTAVSON, S., LARSSON, P., AND YNNERMAN, A. 2008. Free form incident light fields. *Computer Graphics Forum (Proceedings EGSR 2008)* 27, 4, 1293–1301.