Accurate Light Source Acquisition and Rendering

Michael Goesele, Xavier Granier, Wolfgang Heidrich, Hans-Peter Seidel

► To cite this version:

Michael Goesele, Xavier Granier, Wolfgang Heidrich, Hans-Peter Seidel. Accurate Light Source Acquisition and Rendering. ACM Transactions on Graphics, 2003, 22 (3), pp.621 - 630. 10.1145/882262.882316 . hal-00308294v2

HAL Id: hal-00308294

<https://hal.science/hal-00308294v2>

Submitted on 27 Oct 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Accurate Light Source Acquisition and Rendering

Michael Goesele¹⁾, Xavier Granier²⁾, Wolfgang Heidrich²⁾, Hans-Peter Seidel¹⁾

1) MPI Informatik

2) The University of British Columbia

Figure 1: Stages of light source measurement and rendering (from left to right): a) photo of flashlight – b) flashlight in measurement setup – c) 2D reconstruction of the measured data in a virtual plane – d) measured light used in a global illumination simulation.

Abstract

Realistic image synthesis requires both complex and realistic models of real-world light sources and efficient rendering algorithms to deal with them. In this paper, we describe a processing pipeline for dealing with complex light sources from acquisition to global illumination rendering. We carefully design optical filters to guarantee high precision measurements of real-world light sources. We discuss two practically feasible setups that allow us to measure light sources with different characteristics. Finally, we introduce an efficient importance sampling algorithm for our representation that can be used, for example, in conjunction with Photon Maps.

CR Categories: I.3.3 [COMPUTER GRAPHICS]: Picture/Image Generation—Digitizing and Scanning, I.3.3 [COMPUTER GRAPHICS]: Three-Dimensional Graphics and Realism—Raytracing, I.4.1 [IMAGE PROCESSING AND COMPUTER VISION]: Digitization and Image Capture—Radiometry, I.4.1 [IMAGE PROCESSING AND COMPUTER VISION]: Digitization and Image Capture—Sampling.

Keywords: Modeling – Light Sources; Modeling – Object Scanning/Acquisition; Modeling – Physically-based Modelling; Rendering – Global Illumination; Rendering – Image-based Rendering; Rendering – Raytracing

1 Overview and Related Work

Convincing and realistic light source models are an important prerequisite for photorealistic image synthesis. Given the exact de-

sign of a luminaire (including geometry, materials, and emission characteristics of the illuminant), one can apply standard global illumination computations to incorporate complex luminaires into synthetic scenes. This approach, however, is impractical in most real-world scenarios since the precise design data represents an intellectual property of significant value, and is unlikely to be shared for rendering purposes. But even if such design data is made available, it does not account for manufacturing tolerances. However, relatively minor variations of parameters like the glass thickness in halogen light bulbs can have a significant impact on the illumination patterns generated. In addition, rendering based on the geometric model can be computationally very demanding.

As an alternative, one can incorporate measured light data from complex luminaires directly into the rendering system. Such data is provided by many companies in the form of goniometric diagrams [Verbeck and Greenberg 1984], which represent measurements of the far field of a luminaire (i.e. directional information of the emitted light from a point light source). Unfortunately, the far field is only a faithful approximation of the emitted light when the light source is sufficiently far away from the object to be illuminated – as a rule of thumb at least five times the maximum luminaire dimension [Ashdown 1995]. Even improved models using spatial intensity maps or complex emitting geometry [Rykowski and Wooley 1997] are still too weak to represent the near field of a complex light source.

On the other hand, a light field [Gortler et al. 1996; Levoy and Hanrahan 1996] completely represents near and far field illumination of a light source, and can thus be used to represent the light source without knowing its geometry and internal structure.

1.1 Near Field Photometry

In his work on *near field photometry* [Ashdown 1993; Ashdown 1995], Ashdown has presented methods for measuring the light field of luminaires in both its near field and its far field components. A digital camera mounted on a robot arm observes a light source from many positions on a surrounding (hemi-)sphere. Both Rykowski and Wooley [1997], and Jenkins and Mönch [2000] employ a similar setup to acquire the light field of a luminaire, while Siegel and Stock [1996] replaced the camera lens with a pinhole.

On the rendering side, Heidrich et al. [1998] have described relatively efficient algorithms based on a similar representation they

called a *canned light source*. While their algorithms are faster than a brute-force rendering approach, they are still quadratic in the light field resolution. However, the similarity between the representations used by Ashdown and by Heidrich et al. illustrates that measurement and rendering can work on the same data structures, thereby allowing for an efficient processing pipeline from acquisition to image synthesis.

In near field photometry, a number of cameras are pointed at the luminaire to be measured (or, more practically, a single camera is moved around), and then the irradiance incident to the film plane is recorded using an imaging sensor (e.g. a CCD chip). The camera positions correspond to a sampling of some virtual sampling surface $\mathcal{S}$. In practice, the light source may produce arbitrarily high spatial frequencies on $\mathcal{S}$, in which case we have just introduced aliasing by not applying a low-pass filter before the sampling step.

This issue has been raised by Halle [1994] and was noted by Levoy and Hanrahan [1996]. They both propose to use the finite aperture of the camera lens as a low-pass filter by choosing the aperture size equal to the size of a sample on the camera plane. The shape of the filter kernel is thus a side effect of the lens system rather than a deliberate choice based on sampling theory.

1.2 Contributions

In this paper, we describe a measurement system that significantly improves on near field photometry by projecting the light field emitted by the light source into a finite basis *before* sampling. This is done using a simple optical system. The shape and support of the basis functions are specifically designed for a particular sampling scheme. Based on the resulting measurements we can *exactly* reconstruct the least-squares approximation of the true light field in our basis. Alternatively, we can reconstruct with a more efficient, shift-invariant filter, and obtain a close approximation to the least-squares solution that is suitable for hardware accelerated rendering with the approach proposed by Heidrich et al. [1998].

In comparison to previous acquisition approaches, our optical setup lets us move in very close to the surface of the light source. In particular, the distance between the light and the aperture of our system is smaller than the distance between the aperture and the imaging plane. This allows us to capture a much wider field of view in a single photograph.

On the rendering side, we introduce a solution to perform importance sampling of our light field representation. This allows us to integrate the acquired light sources into global illumination methods like the Photon Map Algorithm [Jensen 2001b]. In contrast to the work by Heidrich et al., this algorithm has constant time complexity, i.e. it is independent of the light field resolution.

The remainder of this paper is structured as follows: first, we present our conceptual approach to filtered light source measurement (Section 2) before we discuss the theoretical underpinnings of our method in Section 3. We then describe the physical setup we use for measurement in Section 4, and our rendering algorithm in Section 5. Finally, we present results based on both measurements and simulation (Section 6), and conclude with some ideas for future work (Section 7).

2 Basic Approach

The conceptual setup for our approach to filtered light source measurement is depicted in Figure 2. Light rays are emitted from the light source and hit a filter in a plane $\mathcal{S}$. This plane is opaque except for the finite support area of the filter which can be moved to a number of discrete positions on a regular grid. The filter is a semi-transparent film, similar to a slide, containing the 2D image of a basis function $\Phi'_{ij}(u, v)$. The light falling through this filter hits a

Figure 2: Cross section through the conceptual acquisition setup.

second plane $\mathcal{M}$, on which we are able to measure the irradiance with a high spatial resolution.

When comparing this setup with the work on near field photometry [Ashdown 1993; Siegel and Stock 1996], $\mathcal{S}$ roughly corresponds to the plane on which the pinholes or finite apertures of the cameras are located, while $\mathcal{M}$ corresponds to the imaging sensor of the camera. The fundamental difference in our approach is that we can select an arbitrary filter $\Phi'_{ij}(u, v)$ which projects the light field emitted by the luminaire into an arbitrarily chosen function space. This results in an optical low-pass filtering of the light field *before* sampling, and thus helps avoiding aliasing artifacts.

In general, we will select a standard reconstruction filter $\Phi_{ij}(u, v)$ for rendering, and its biorthogonal basis $\Phi'_{ij}(u, v)$ for measurement. This will make it necessary to perform measurements with filters that contain negative values, and can therefore not be implemented optically as an opacity. In this case, we split the filters into a positive and a negative component:

$$\Phi'_{ij(+)}(u, v) := \begin{cases} \Phi'_{ij}(u, v) & \Phi'_{ij}(u, v) \geq 0 \\ 0 & \text{else} \end{cases}$$

$$\Phi'_{ij(-)}(u, v) := \begin{cases} 0 & \Phi'_{ij}(u, v) \geq 0 \\ -\Phi'_{ij}(u, v) & \text{else} \end{cases}$$

We then perform separate measurements for both parts, and subtract the results in software. Similarly, we can rescale basis functions with values > 1 , and correct for it in software.

In principle, one also has to perform low-pass filtering on the measurement plane $\mathcal{M}$. This, however, is optically very difficult or even impossible to implement if the filter kernels overlap. Fortunately, it is technically easy to achieve a very high measurement resolution on $\mathcal{M}$, so that we can consider the measured data as point samples, and implement an arbitrary filter kernel while downsampling to the final resolution.

Note that the filter kernels for $\mathcal{S}$ and $\mathcal{M}$ can be chosen independently as required by the reconstruction process. In this paper, we are using a piecewise quadratic basis to construct $\Phi'_{ij}(u, v)$ (see Section 4.2), and a box filter for $\Phi'_{kl}(s, t)$. The latter works best with our importance sampling algorithm (Section 5 and Appendix B).

3 Theoretical Framework

Before we discuss the theory behind the proposed method in detail, we first introduce the mathematical notation used throughout this document. This notation is summarized in Table 1 and Figure 2.

Symbol	Meaning
$\Psi_{ijkl}(u, v, s, t)$	basis function for approximating the light field
Φ_i	1D basis used for reconstruction
$\Phi_{ijkl}(u, v, s, t)$	4D tensor product basis for reconstruction
Φ'_i	biorthogonal 1D basis used for measurement
$\Phi'_{ijkl}(u, v, s, t)$	4D tensor product basis for measurement
$\mathcal{M}$	surface on which the irradiance is measured (<i>measurement surface</i> , or (s, t) -plane)
$\mathcal{S}$	surface on which the optical filters are placed (<i>sampling surface</i> , or (u, v) -plane)
$L(u, v, s, t)$	radiance passing through (u, v) on $\mathcal{S}$ and (s, t) on $\mathcal{M}$
$\tilde{L}(u, v, s, t)$	projection of $L(u, v, s, t)$ into basis $\{\Psi_{ijkl}(u, v, s, t)\}$
$L_{mn}(u, v, s, t)$	radiance $\Phi'_{mn}(u, v) \cdot L(u, v, s, t)$ projected through <i>one</i> filter $\Phi'_{mn}(u, v)$
$E_{mn}(s, t)$	irradiance caused by $L_{mn}(u, v, s, t)$ on the measurement surface $\mathcal{M}$

Table 1: Notation (overview) – see also Figure 2.

For the measurement, we assume that the light field emitted by the light source is well represented by a projection into a basis $\{\Psi_{ijkl}(u, v, s, t)\}_{ijkl \in \mathbb{Z}}$:

$$L(u, v, s, t) \approx \tilde{L}(u, v, s, t) := \sum_{i,j,k,l} \Psi_{ijkl}(u, v, s, t) \cdot L_{ijkl}. \quad (1)$$

We assume that Ψ_{ijkl} has local support, and i, j, k , and l roughly correspond to translations in u, v, s , and t , respectively. Note, however, that the translated basis functions will not in all cases have the same shape, i.e. $\Psi_{i'j'k'l'}(u, v, s, t)$ may not be an exact copy of $\Psi_{ijkl}(u, v, s, t)$ in general.

We also define two additional sets of basis functions, one for *measuring* and one for *reconstruction*. For reconstruction we use a 1D basis $\{\Phi_i\}_{i \in \mathbb{Z}}$ with the property $\Phi_i(x) = \Phi(x + i)$. The 4D reconstruction basis is then given as the tensor product basis

$$\Phi_{ijkl}(u, v, s, t) := \Phi_{ij}(u, v) \cdot \Phi_{kl}(s, t) \quad (2)$$

$$= \Phi_i(u) \cdot \Phi_j(v) \cdot \Phi_k(s) \cdot \Phi_l(t). \quad (3)$$

For measurement, we use the *biorthogonal* (or *dual*) $\{\Phi'_i(x)\}_{i \in \mathbb{Z}}$ of the reconstruction basis with

$$\int_{-\infty}^{\infty} \Phi'_i(x) \cdot \Phi_{i'}(x) dx = \begin{cases} 1 & \text{if } i = i' \\ 0 & \text{else} \end{cases} \quad (4)$$

and again we use a tensor-product construction for the 4D basis.

3.1 Measured Irradiance

Our approach is based on measuring the irradiance $E_{mn}(s, t)$ in the (s, t) -plane that is caused by the incident radiance $L_{mn}(u, v, s, t) = \Phi'_{mn}(u, v) \cdot L(u, v, s, t)$. In Section 4.1 we will discuss two physical setups for performing this kind of measurement. The result of such a measurement is

$$\begin{aligned} E_{mn}(s, t) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{\cos^2 \theta}{R^2} \cdot \Phi'_{mn}(u, v) \cdot L(u, v, s, t) du dv \\ &\approx \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{\cos^2 \theta}{R^2} \cdot \Phi'_{mn}(u, v) \cdot \tilde{L}(u, v, s, t) du dv \\ &= \sum_{i,j,k,l} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{\cos^2 \theta}{R^2} \cdot \Phi'_{mn}(u, v) \\ &\quad \cdot \Psi_{ijkl}(u, v, s, t) \cdot L_{ijkl} du dv. \end{aligned} \quad (5)$$

The geometric term $\cos^2 \theta / R^2$ is composed of the distance R of the point on the (u, v) -plane from the point on the (s, t) plane, as

well as the cosine of the angle θ between the plane normals and the vector connecting the two points. Note that this term also accounts for any differences in the parameterizations on the two planes (i.e. different grid spacings).

3.2 Exact Reconstruction

We now describe an exact reconstruction algorithm given the measurements E_{mn} . To this end, we first define what the relationship between the basis functions Ψ_{ijkl} and the reconstruction and measurement bases should be. We define

$$\Psi_{ijkl}(u, v, s, t) := \frac{R^2}{\cos^2 \theta} \cdot \Phi_{ij}(u, v) \cdot \Phi_{kl}(s, t). \quad (6)$$

Inserting this definition into Equation 5, and using the biorthogonality relationship (Equation 4) yields

$$E_{mn}(s, t) = \sum_k \sum_l \Phi_{kl}(s, t) \cdot L_{mnkl}. \quad (7)$$

To determine which reconstruction filter to use, we now rewrite Equation 1 using Equations 6 and 7:

$$\begin{aligned} \tilde{L}(u, v, s, t) &= \sum_{m,n,k,l} \Psi_{mnkl}(u, v, s, t) \cdot L_{mnkl} \\ &= \sum_{m,n,k,l} \frac{R^2}{\cos^2 \theta} \cdot \Phi_{mn}(u, v) \cdot \Phi_{kl}(s, t) \cdot L_{mnkl} \\ &= \sum_{m,n} \frac{R^2}{\cos^2 \theta} \cdot \Phi_{mn}(u, v) \cdot E_{mn}(s, t) \end{aligned} \quad (8)$$

This indicates that we can *exactly* reconstruct $\tilde{L}$, the projection of L into the basis $\{\Psi_{ijkl}\}$ by using the reconstruction filter $R^2 / \cos^2 \theta \cdot \Phi_{mn}(u, v)$.

Note that this reconstruction filter does contain a shift-variant component in form of the geometric term. Interestingly, this geometric term will cancel out for the Photon Map particle emission algorithm we discuss in Section 5. However, for other rendering algorithms such as the ones proposed by Heidrich et al. [Heidrich et al. 1998], the reconstruction step should ideally only involve a convolution with a shift-invariant filter kernel. In Appendix A we discuss an approximate reconstruction scheme that satisfies this property.

4 Physical Setup and Measurements

In this section we explain how the theoretical framework can be transformed into a practical acquisition method. We propose two different hardware setups and discuss their advantages and disadvantages. Finally, we describe the filter kernels used in our measurements.

4.1 Possible Implementations

The first setup we propose – Setup A – is based on a standard photographic camera where the lens is removed and replaced by the filter Φ'_{ij} (see left side of Figure 3). As with conventional light fields, the camera has to be positioned at the vertices of a regular grid within the $\mathcal{S}$ plane and images can be captured. Although this setup is compact and easy to handle, it does have some drawbacks. The fixed and often quite small size of the imaging sensor (e.g. a CCD chip) imposes a strong limitation on the size of the filter kernel and the incident angles of illumination covered by the measurement. In addition, the distance between the filter plane and the imaging plane determines the focal length of the camera (similar to a pinhole camera). A large distance leads to a tele-lens effect and further limits the useful incident angles.

Figure 3: Possible setups for the acquisition. Setup A (left): The camera lens is replaced by a filter and the light field is directly projected onto the imaging sensor inside the camera. Setup B (right): The filter $\Phi'_{ij}(u, v)$ projects an image on a diffuse reflector which is captured by a camera.

These problems can be avoided with Setup B. Here, the filter Φ'_{ij} is used to project an image of the light source onto a diffusely reflecting screen. If a high quality diffuse (Lambertian) reflector material is available, the filter size and viewing angles are (almost) only limited by the size of the reflector. On the other hand, the larger setup makes it potentially harder to use, and errors due to indirect light reflecting from parts of the equipment such as the reflection screen or the environment tend to increase. Furthermore, the calibration effort increases since the intrinsic and extrinsic parameters of the camera system relative to the projected image need to be recovered.

4.2 Filter Design

The theoretical framework and the implementations detailed above allow for the use of arbitrary filters Φ_i and Φ'_i with finite support.

It is therefore possible to design filters adapted to the specific properties of the light source at hand, for example taking the structure of the light source or the intended use of the data into account.

Figure 4: A quadratic reconstruction basis and one dual. Left: function Φ_i (green) and a dual Φ'_i (red) used in this paper. Center and right: 2D filter images used for the acquisition (split into positive and negative components and normalized).

For our measurements we use a piecewise quadratic basis function $\Phi_{i, \text{bellq}}$

$$\Phi_{i, \text{bellq}} = \begin{cases} 1 - 2x^2 & |x| \leq \frac{1}{2} \\ 2(|x| - 1)^2 & \frac{1}{2} < |x| \leq 1 \\ 0 & \text{else} \end{cases} \quad (9)$$

Although it would also be possible to use the classical quadri-linear basis, we choose this one due to its C^1 continuity that helps preventing Mach bands. Figure 4 shows a plot of this function and its dual.

Note that some parts of the dual basis are negative. Using this function as a filter will result in negative coefficients. This will occur in places with strong intensity gradients in the original light field. There are two ways of dealing with this situation. The first one is to clamp the value of the coefficients after measurement, guaranteeing that each query in the light field will be positive. However, this approach will remove fine details in the light field and contradicts our filtering approach. We therefore choose to keep the negative coefficients in the data structures, and to only clamp negative results after reconstruction. With this approach, the positive parts of the reconstructed light field remain unaltered.

5 Rendering

In this section, we discuss the integration of our measured light sources into a particle based global illumination system. We chose the Photon Map algorithm [Jensen 1996; Jensen 2001a; Jensen 2001b] for rendering, as it is one of the most common and simple to use methods available. We present a solution to emit particles from our measured data, and we also describe some adjustments to the classical Photon Map algorithm. Note that our importance sampling approach does not only apply to Photon Maps, but to any algorithm that requires shooting from the light source such as particle tracing [Walter et al. 1997] or bidirectional path-tracing [Lafortune and Willems 1993].

5.1 Importance Sampling

For the scene description, we attach the light field to a virtual geometric object, representing an impostor for the area from which light is emitted. This can be, for example, the glass surface of a car head light, or simply a planar polygon representing one of the light field planes. This geometry helps in positioning the light source in the 3D scene and can also be used to aid the sampling in our rendering algorithm.

To emit a particle from our light source representation, we could then use the standard method of selecting a random point on the impostor geometry and then a random direction. While this approach might be sufficient for a relatively uniform light source, it can be inefficient when high spatial and directional frequencies are present.

Figure 5: A light field cell for PDF computation formed by a single pixel on the measurement plane $\mathcal{M}$ and by $1/4$ of a filter on the filter plane $\mathcal{S}$. This corresponds to the region on $\mathcal{S}$ where $\Phi_i(u)$, $\Phi_{i+1}(u)$, $\Phi_j(v)$, and $\Phi_{j+1}(v)$ overlap.

We therefore develop an importance sampling strategy for our light field representation.

The particles that we want to emit have to estimate the total energy E coming from the light field:

$$E = \int_{u,v,s,t} \tilde{L}(u,v,s,t) \frac{\cos^2 \theta}{R^2} \quad (10)$$

In order to achieve a constant energy for each photon, the probability density function (PDF) of the photon emission has to be:

$$p(u,v,s,t) = \frac{1}{E} \tilde{L}(u,v,s,t) \frac{\cos^2 \theta}{R^2} \quad (11)$$

To generate random photons distributed according to this PDF, we use the standard approach of generating uniformly distributed random points in 4D, and passing them through the inverted cumulative density function (CDF, [Pitman 1992]). If we have pre-inverted and tabulated the CDF, then generating a new random photon according to this importance is possible in constant time, i.e. independent of the light field resolution.

One complication in our case is, however, that we have the PDF represented in terms of some complicated and overlapping basis functions instead of the usual point sampled (i.e. tabulated) representation. We therefore do not directly generate and invert the CDF from the basis representation, but rather choose a two-steps approach. We first compute a 4D table of constant probabilities for selecting a 4D cell from the light field (see Figure 5). A single 4D cell corresponds to all the directions coming from one pixel on the measurement plane to $1/4$ of a filter on the filter plane (corresponding to the minimum overlap between two neighboring filters in our basis). The probability on this cell can be computed by some simple summations over the light field coefficients. The details are described in Appendix B. Then, as we now have a constant probability for each cell, we can invert this table more easily than the continuous representation of the PDF. Note that all these steps are performed only once in a preprocessing stage. The resulting table is part of the light source description that we save on disk.

5.2 Direct Illumination

As pointed out in the previous section, the light field is attached to a virtual base geometry. The irradiance at a point $\mathbf{x}$ can easily be computed by sampling this geometry. But as in the case of indirect illumination, this approach can be quite inefficient depending on the directional variation of the light source.

Our approach is therefore to separate the Global Map of the traditional Photon Map algorithm into two: the Direct Map (i.e. the direct impacts from the lights sources), and a new Global Map (i.e. impacts after one diffuse reflection). With this separation, we have two possibilities for computing the direct illumination: we can either choose to reconstruct it straight from the Direct Map, or we can use the Direct Map for generating an importance to guide a standard Monte Carlo lighting integration.

The former method promises to be quite efficient for highly directional light sources. The quality of this reconstruction can also be increased by storing a larger photon number in the Direct Map, while keeping a low number for the Global Map. This approach involves either limiting the integration to regions where some direct photons are present in the neighborhood, or building an importance sampling table from the nearest particles in the Direct Map (similar to the approach used for the Global Map by Jensen [1995]).

The idea of identifying the direct impacts was already introduced in [Keller and Wald 2000]. However, instead of using two different maps, they use the classical Global Map and extend the photon data structure by storing a light source identifier. While this approach is well suited for determining which light source the particle is coming from, it is less efficient for finding all direct impacts, which is what we require here.

In general, the separation of the Global Map into two maps gives us a lot more flexibility in the reconstruction process. As pointed out above, the user can choose between different methods for both the direct and the indirect illumination. Note that this algorithm is independent of our light source representation and can also be used in other Photon Map implementations. This is also true for the following section.

5.3 Other Photon Map Improvements

Our implementation uses a slightly different algorithm to find the N_p closest photons in one of the two maps. We find our approach to be better in determining which photons actually contribute to the irradiance estimate.

In the following, $\vec{d}_i$ will denote the impact direction of a photon p_i , and $\mathbf{x}_i$ will denote its position. We then search for the particles closest to a position $\mathbf{x}$ with a normal $\vec{n}$. During the search, we only consider the particles that have a negative dot product ($\vec{d}_i \cdot \vec{n}$), meaning that the particle arrives from above the surface. This guarantees that only the incident photons on one side of a wall, for example, will contribute to the reconstructed illumination.

Then, with an approach similar to the one described by Lastra et al. [2002], the density is estimated on the tangent plane. The particles p_i are projected into the tangent plane at $\mathbf{x}$, and the Euclidean distance of that projection from $\mathbf{x}$ is used as the distance metric. The distance is thus defined as:

$$\delta(\mathbf{x}, p_i)^2 = \alpha^2 - 2\alpha(\mathbf{x} - \mathbf{x}_i \cdot \vec{d}_i) + \|\mathbf{x} - \mathbf{x}_i\|^2$$

$$\alpha = \frac{(\mathbf{x} - \mathbf{x}_i \cdot \vec{n})}{(\vec{d}_i \cdot \vec{n})}$$

This allows us to increase the accuracy of the reconstruction, especially for non-planar surfaces and corners.

6 Results

We first present results from a simulation of the complete process, from acquisition (with Setup A) to rendering as a validation of our approach. We then discuss our measurement setups and show renderings of acquired data sets.

6.1 Simulations

Before building the real setup, we validate our approach by simulating the acquisition of a strongly directional light field such as a slide projector. It has an aperture of 50 mm and a focal length of 59 mm. We use an inverted SIGGRAPH 2003 logo as a slide, to create strong directional variations (see Figure 6). Simulations are done using Setup A with a filter size of $5.068 \text{ mm} \times 5.068 \text{ mm}$ and the quadratic basis. The simulated camera has an imaging sensor with a size of $27.15 \text{ mm} \times 18.1 \text{ mm}$ and a resolution of 150×100 pixels. The distance between the filter and the imaging sensor is 39 mm, and the distance between the slide projector and the camera is 20 mm. We simulate the acquisition of a 41×41 array of images. We store the data in a signed version of the RGBE format [Ward 1991] to have a more efficient usage of the memory. The resulting simulation uses about 96 MB of memory.

Figure 6: Projections from the simulated light field of a slide projector. The first image (left-up) is a projection on a plane in focus, the second one (left-down) is a projection on a plane closer to the slide projector.

Based on this data, we then computed a sequence of projections on a plane moving from the lens of the slide projector to the projector’s focus plane, to show the directional changes (see Figure 6). The in-focus image shows that we manage to capture the whole shape of the light source, even if the small details are not present which is expected for this setup geometry and the optical filtering. We also used this data to compute a global illumination solution that validates the particle emission approach (see Figure 7). This scene has about 15,000 primitives, and it took about 2 minutes to compute the inverted probability table (size $32 \times 32 \times 128 \times 128$), 4 seconds for emitting 40,000 photons, and 52 minutes to render one image.¹

6.2 Measurements

In order to measure real light sources we built systems for each of the two setups proposed in Section 4.1. The filters were printed with 1016 dpi resolution on high density film using a calibrated LightJet 2080 Digital Film Recorder. This produces slides with a dynamic range of more than $1:50,000$, which are then mounted into suitable frames. The light sources to be measured – we used a bike light as well as several different flashlight configurations – are mounted on manually operated translation stages so that they can be shifted in front of the camera with high precision and repeatability (see Figure 8). The images were captured with a Kodak DCS 560

¹ All timings are for a 1GHz Pentium III Workstation.

Figure 7: Global illumination solution for the simulated acquisition data.

Figure 8: Left: The Mini Mag-Lite and the bike light mounted to the translation stage. Right: The negative part of the filter illuminated by the bike light. Both images show parts of Setup B.

digital camera in our measurement lab which is covered with dark material to avoid interference with the measurements [Goesele et al. 2000].

For Setup A, a custom-built filter holder was attached to the camera instead of a lens, so that the filters can easily be exchanged without affecting the geometry of the setup (see Figure 9, left). High-dynamic range acquisition techniques [Debevec and Malik 1997; Robertson et al. 1999] were used to acquire the images.

For Setup B, we used a $1.2 \text{ m} \times 1.2 \text{ m}$ flat board that was spray-painted with ultra matte paint as a projection screen. A black board with an integrated filter holder was used to mount the filter and shield the light source from the projection screen. To avoid occlusion problems we mounted the camera above the whole setup and equipped it with a 24 mm shift lens allowing us to capture the whole projection screen. The camera parameters were recovered using standard camera calibration techniques and the images were rectified and registered to the filter setup. Figure 9 (right) shows an overview of Setup B. A 5×5 median filter was used to suppress dark current noise resulting from longer exposure times in this setup; remaining noise appears as background illumination in the light field.

6.2.1 Measured Data Sets

We acquired several data sets with these setups. Table 2 summarizes the acquisition geometry, Table 3 gives an overview over the number of views captured and the filters used. All acquired images were downsampled to achieve a resolution of 300×300 pixels on the measurement plane $\mathcal{M}$. Acquisition times are on the order of

Figure 9: Left: Setup A. The filter holder mounted to the camera in front of the lens of a light source. Right: Setup B. The projection screen is observed by the camera located above the setup. The black board contains the filter holder. The light sources can be moved on the table parallel to the filter and the projection screen.

1–2 min per view due to the manual work involved in moving the translation stage.

Setup	Distance $\mathcal{M} - \mathcal{S}$	Size of $\mathcal{M}$	Filter Width (Dual Basis)
Setup A	55 mm	18 mm $\times$ 27 mm	10 mm (small)
Setup B	750 mm	1.2 m $\times$ 1.2 m	20 mm (medium) 30 mm (large)

Table 2: Geometry of the acquisition setups.

Data Set	Setup	# Views	Filter Width (Dual Basis)
<i>Mag-Lite</i>	A	9 $\times$ 9	10 mm (small)
<i>Mag-Lite Medium</i>	B	7 $\times$ 7	20 mm (medium)
<i>Mag-Lite Large</i>	B	5 $\times$ 5	30 mm (large)
<i>Bike Light</i>	B	9 $\times$ 7	30 mm (large)
<i>LED Flashlight</i>	B	5 $\times$ 5	30 mm (large)

Table 3: Parameters for the acquired data sets. Note that the width of the reconstruction filter is 1/2 of the width of the dual filter, the filter spacing on the filter plane $\mathcal{S}$ is 1/4 of the width of the dual filter.

Figure 10: *Mag-Lite* acquisition with Setup A: Left, a projection on a distant (3 m) plane, Right: a global illumination solution (Importance table construction - 20 s, Particle emission - 40 s, Rendering 20 min). The acquired data set was converted to grayscale.

Only the *Mag-Lite* data set was captured with Setup A. Due to geometry constraints imposed by the dimensions of the camera (see Table 2), we were only able to capture a small part of the light field.

Figure 10 shows on the left the light source pointed at a projection screen and on the right onto the crypt. The captured part corresponds roughly to the central dark region in the projection of the real light source (see top of Figure 11). The low resolution is due to the comparably large filter size. Using a smaller filter would increase both the resolution and the number of views required. In this configuration, Setup A serves as a proof-of-concept but can probably not be used to acquire a complete model of a real light source.

Figure 11: Top: photograph of the Mag-Lite illuminating a plane at 75 cm distance. Bottom left: reconstruction of the same setup from data measured with a filter of 30 mm width. Bottom right: reconstruction from data measured with a filter of 20 mm width.

With Setup B, we can generally capture a much larger portion of the light field of the light source due to the 67° viewing angle of the configuration. Additionally, a much larger filter kernel can be used for the filter plane while still improving the resolution of the data set.

In Figure 11 we show the light pattern that a Mini Mag-Lite flash light causes on a plane at 75 cm distance for a specific focus setting on that flashlight. The top image is a photograph of the flashlight pointed at a diffuse surface. The bottom row shows two synthetic images that are reconstructed from measurements. The left image was generated from a dataset where a dual filter width of 30 mm has been used, while a dual filter of 20 mm width was used for the right image. As expected, the larger filter size results in a slightly more blurry image on the left. In comparing the reconstructed images with the photograph, we can see that there are some tone-mapping issues that result in slightly different colors. However, the main features of the light cone, such as the black ring in the center, are represented very well in both reconstructions.

In Figure 12 we show similar images for the light field emitted by a bike light (the light source itself is depicted in the top left of the figure). The top right of the figure again shows a photograph of the light source illuminating a plane at 75 cm distance. The image on the lower right is a reconstruction for the same distance, while the image on the lower left shows the reconstruction at a distance of only 18.7 cm.

The narrow vertical stripes are an interesting feature of the bike light that should be preserved in a model. They are mainly emitted from the upper and lower edges of the luminaire and are reproduced

Figure 12: Acquisition on the *Bike Light* model. The upper right image shows a plane illuminated by the real light source, at the distance of 75 cm. The lower right image shows the same plane illuminated by the acquired data, while the lower left image shows the result for a plane at a quarter of that distance.

very well in the renderings. Furthermore, due to the effective size of the luminaire (5 cm×2.5 cm) and its strongly varying emission characteristics, the near field is quite extended and inhomogeneous. In the rendered image of the *Bike Light* model projected at different distances, variations of the near field are clearly visible. Figure 14 shows a global illumination solution based on this dataset. Note both the distinctive patterns on the wall and the stripe patterns on floor and ceiling.

Finally, we have performed the same test for a flashlight with 3 LEDs instead of a light bulb (see left of Figure 13 - the fourth, smaller LED is a low battery light, and is not used in normal operation). Although the actual flashlight is identical to the one from Figure 11 except for the light bulb and the reflector, the cone of light produced by this setup is quite different from the one before. In particular, the illumination is much more uniform, although roughly triangular due to the LED layout. This illustrates the importance of using actual measurements from real world light sources in image synthesis, since even small changes in the reflectors can cause major differences in appearance.

Figure 13: Photograph and reconstruction of the *LED Flashlight*.

Figure 14: Exploring the cloister with the *Bike Light*. Importance table construction - 37 s, Particle emission - 12 s, Rendering 30 min.

7 Conclusions

In this paper we presented a novel approach for accurate light acquisition. A model is measured using a simple optical system to project the real light source into a finite element basis. This projection allows a correct optical pre-filtering of the acquired light field. We introduced a theoretical framework for this approach and two possible setups for real acquisitions. The tests we performed to validate this approach both on a simulated process and on the two physical setups, have shown that our setup can faithfully represent rather complex near and far field information with a relatively small number of measurements. The fact that we perform optical low-pass filtering avoids aliasing artifacts such as missing directional peaks. In our evaluations we found that the second setup, although larger and harder to calibrate, is more practical for real-world light sources, since it covers a larger field-of-view.

All our acquisitions were done manually. As only linear movements are involved, the entire measurement can be speed up easily by using a motor driven translation stage. This will lead to an entirely automatic process, that would allow us to further increase the resolution of the measurements.

We also presented an efficient algorithm to integrate our new light representation into particle-based global illumination systems. Using a pre-inverted cumulative density function, this method allows a constant time importance sampling to emit those particles. Combined with the acquisition step, we therefore have a complete pipeline to bring real light into virtual worlds.

Acknowledgements

We would like to thank Ian Ashdown and the anonymous Siggraph reviewers for valuable comments on earlier versions of the paper. Thanks to REVES/INRIA Sophia-Antipolis for allowing us to use the GIS Global Illumination platform for development in this project, and to UBC's Institute of Applied Mathematics for the access to their PC-Cluster for the measurement simulation. Special thanks to the MPI technical staff for building most of the measurement apparatus and to Oliver Ashoff who helped to perform the measurements. This work was funded by the PIMS Post-doctoral Fellowship program, the BC Advanced Systems Institute, and by the DFG Schwerpunktprogramm V3D2.

References

- ASHDOWN, I. 1993. Near-Field Photometry: A New Approach. *Journal of the Illuminating Engineering Society* 22, 1 (Winter), 163–180.
- ASHDOWN, I. 1995. Near-Field Photometry: Measuring and Modeling Complex 3-D Light Sources. In *ACM SIGGRAPH 95 Course Notes - Realistic Input for Realistic Images*, ACM, 1–15.
- DEBEVEC, P., AND MALIK, J. 1997. Recovering High Dynamic Range Radiance Maps from Photographs. In *Computer Graphics Proceedings (ACM SIGGRAPH 97)*, ACM, 369–378.
- GOESELE, M., HEIDRICH, W., LENSCH, H., AND SEIDEL, H.-P. 2000. Building a Photo Studio for Measurement Purposes. In *Proc. of the 5th VMV Conference*, 231–238.
- GORTLER, S. J., GRZESZCZUK, R., SZELINSKI, R., AND COHEN, M. F. 1996. The Lumigraph. In *Computer Graphics Proceedings (ACM SIGGRAPH 96)*, ACM, 43–54.
- HALLE, M. W. 1994. Holographic stereograms as discrete imaging systems. In *Practical Holography VIII*, SPIE, vol. 2176, 73–84.
- HEIDRICH, W., AND GOESELE, M. 2001. Image-based measurement of light sources with correct filtering. Tech. Rep. TR-2001-08, Department of Computer Science, The University of British Columbia.
- HEIDRICH, W., KAUTZ, J., SLUSALLEK, P., AND SEIDEL, H.-P. 1998. Canned lightsources. In *Rendering Techniques '98*, Eurographics, 293–300.
- JENKINS, D. R., AND MÖNCH, H. 2000. Source Imaging Goniometer Method of Light Source Characterization for Accurate Projection System Design. In *Proc. of SID Conference*, 862–865.
- JENSEN, H. W. 1995. Importance Driven Path Tracing Using the Photon Map. In *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, Springer-Verlag, New York, NY, P. M. Hanrahan and W. Purgathofer, Eds., Eurographics, 326–335.
- JENSEN, H. W. 1996. Global Illumination Using Photon Maps. In *Rendering Techniques '96 (Proceedings of the Seventh Eurographics Workshop on Rendering)*, Springer-Verlag/Wien, New York, NY, Eurographics, 21–30.
- JENSEN, H. W. 2001. A practical guide to global illumination using photon mapping. In *ACM SIGGRAPH 2001 Course Notes CD-ROM*, ACM SIGGRAPH, ACM. Course 38.
- JENSEN, H. W. 2001. *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters, Natick, MA.
- KELLER, A., AND WALD, I. 2000. Efficient importance sampling techniques for the photon map. In *Proceedings of Vision Modelling and Visualization 2000*, 271–279.
- LAFORTUNE, E. P., AND WILLEMS, Y. D. 1993. Bi-directional Path Tracing. In *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics '93)*, 145–153.
- LAстра, M., UREA, C., REVELLES, J., AND MONTES, R. 2002. A density estimation technique for radiosity. In *Ibero-American Symposium in Computer Graphics (SIACG'2002)*.
- LEVOY, M., AND HANRAHAN, P. 1996. Light field rendering. In *Computer Graphics Proceedings (ACM SIGGRAPH 96)*, ACM, 31–42.
- PITMAN, J. 1992. *Probability*. Springer.
- ROBERTSON, M. A., BORMAN, S., AND STEVENSON, R. L. 1999. Dynamic Range Improvement Through Multiple Exposures. In *Proc. of the Int. Conf. on Image Processing (ICIP'99)*, IEEE, 159–163.
- RYKOWSKI, R. F., AND WOOLEY, C. B. 1997. Source Modeling for Illumination Design. In *Lens Design, Illumination, and Optomechanical Modeling*, SPIE, vol. 3130, 204–208.
- SIEGEL, M. W., AND STOCK, R. D. 1996. A General Near-Zone Light Source Model and its Application to Computer Automated Reflector Design. *SPIE Optical Engineering* 35, 9 (September), 2661–2679.
- VERBECK, C. P., AND GREENBERG, D. P. 1984. A comprehensive light source description for computer graphics. *IEEE Computer Graphics & Applications* 4, 7 (July), 66–75.
- WALTER, B., HUBBARD, P. M., SHIRLEY, P., AND GREENBERG, D. P. 1997. Global illumination using local linear density estimation. *ACM Transactions on Graphics* 16, 3 (July), 217–259.
- WARD, G. 1991. *Graphics Gems II*. Morgan Kaufman Publishers Inc., San Francisco, CA, USA, ch. Real Pixels, 80–83.

A An Approximate Shift-Invariant Reconstruction

For an approximate shift-invariant reconstruction, we start with the measured irradiance from Section 3.1, but we define the relationship between the Ψ_{ijkl} and the measurement and reconstruction bases differently than we did for the exact reconstruction (Section 3.2). Here, we define $\Psi_{ijkl}(u, v, s, t) := \Phi_{ij}(u, v) \cdot \Phi_{kl}(s, t)$. From this, we get

$$E_{mn}(s, t) = \sum_{i,j,k,l} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{\cos^2 \theta}{R^2} \cdot \Phi'_{mn}(u, v) \cdot \Phi_{ij}(u, v) \cdot \Phi_{kl}(s, t) \cdot L_{ijkl} du dv \quad (12)$$

This is still a shift variant filter, since the geometric term depends on the integration variables u and v , and can therefore not be moved outside the integral. However, if the distance d between the (u, v) -plane and the (s, t) -plane is large compared to the support of $\Phi_{ij}(u, v)$, and if θ is small, then this term is well approximated by one constant per point on the (s, t) -plane:

$$\frac{\cos^2 \theta(u, v, s, t)}{R(u, v, s, t)^2} \approx \frac{\cos^2 \theta_{kl}(s, t)}{R_{kl}(s, t)^2}. \quad (13)$$

This yields an approximation of the measured irradiance:

$$E_{mn}(s, t) \approx \sum_{k,l} \frac{\cos^2 \theta_{kl}(s, t)}{R_{kl}(s, t)^2} \cdot \Phi_{kl}(s, t) \cdot L_{mnkl}. \quad (14)$$

Since both the approximate geometric term and $\Phi_{kl}(s, t)$ are known, it is in principle possible to compute

$$E'_{mn}(s, t) \approx \sum_{k,l} \Phi_{kl}(s, t) \cdot L_{mnkl} \quad (15)$$

by de-convolution. In practice, this is only feasible for basis functions $\Phi_{kl}(s, t)$ with a small support. We do not expect this to be a major problem, however, since the practical measurement setups discussed in Section 4.1 have a very high resolution on the (s, t) -plane, so that a bilinear or even box filter can be used. Again, we apply the definition of Ψ_{ijkl} to determine the appropriate reconstruction filter:

$$\begin{aligned} \tilde{L}(u, v, s, t) &= \sum_{m,n,k,l} \Psi_{mnkl}(u, v, s, t) \cdot L_{mnkl} \\ &= \sum_{m,n,k,l} \Phi_{mn}(u, v) \cdot \Phi_{kl}(s, t) \cdot L_{mnkl} \\ &\approx \sum_m \sum_n \Phi_{mn}(u, v) \cdot E'_{mn}(s, t). \end{aligned} \quad (16)$$

Thus, an *approximate* reconstruction of $\tilde{L}$, the projection of L into the basis $\{\Psi_{ijkl}\}$ is obtained using the shift-invariant reconstruction filter $\Phi_{mn}(u, v)$. For example, if Φ is a hat function,

the reconstruction process is reduced to quadri-linear interpolation, which is frequently used for light field rendering algorithms.

The quality of this approximation depends on the error introduced by assuming the geometric term constant over the support of basis function $\Phi_{ij}(u, v)$ in Equation 13. An analysis of this error [Heidrich and Goesele 2001] shows that it is below 10% if the filter is no larger than 10% of the distance between $\mathcal{M}$ and $\mathcal{S}$. That error is reduced to about 1% for filter sizes no larger than 2% of that distance.

B Importance Sampling Details

This appendix details the computation of the 4D probability table introduced in Section 5.1.

To simplify the following equations and the creation of the 4D table, we make the following assumptions: (a) - the configuration is centered; (b) - the overlap between two consecutive basis functions on the sampling plane $\mathcal{S}$ correspond to half of its size for width and height; (c) - the width and height shift (s_w and s_h , in pixel number) between two measurements are integers; (d) - the basis function Φ_i is symmetric; (e) - the basis function on the measurement plane $\mathcal{M}$ is the box filter; (f) - the support of the original reconstruction basis function Φ is $[-1, 1]$ and its integral is 1 on this support. With this assumption, the integrals of the reconstruction basis function are (if their dimension are $w_S \times h_S$):

$$\int \Phi_i(u) = \frac{w_S}{2}, \quad \int \Phi_j(v) = \frac{h_S}{2}$$

With this we can compute the total energy emitted by the light field (see Equation 10):

$$E = \frac{w_S \times h_S \times w_M \times h_M}{4} \sum_{i,j,k,l} L_{ijkl}$$

Here, w_M and h_M represent the dimensions of a pixel on the measurement plane.

The 1D-intervals defining a cell can be described by the following equations ($\mathbb{I}_{W,n}^S$ and $\mathbb{I}_{H,o}^S$ on the sampling plane, and $\mathbb{I}_{W,n,p}^M$ and $\mathbb{I}_{H,o,q}^M$ on the measurement plane - these depend on the selected intervals on the sampling plane):

$$\begin{aligned} \mathbb{I}_{W,n}^S &= (-(W_S + 1)/2 + [n, n + 1]) \times w_S \quad \forall n \in [0..W_S] \\ \mathbb{I}_{H,o}^S &= (-(H_S + 1)/2 + [o, o + 1]) \times h_S \quad \forall o \in [0..H_S] \\ \mathbb{I}_{W,n,p}^M &= \frac{(2n - W_S - 1)w_S/2 - w_M W_M}{2} + [p, p + 1] \times w_M \\ &\quad \forall p \in [0..W_M + s_w - 1] \\ \mathbb{I}_{H,o,q}^M &= \frac{(2o - H_S - 1)h_S/2 - h_M H_M}{2} + [q, q + 1] \times h_M \\ &\quad \forall q \in [0..H_M + s_h - 1] \end{aligned}$$

Note that $W_S \times H_S$ represents the resolution on the sampling plane (width×height), and $W_M \times H_M$ on the measurement plane. A set of light directions $\mathbb{I}_{nopq}^S = \mathbb{I}_{W,n}^S \times \mathbb{I}_{H,o}^S \times \mathbb{I}_{W,n,p}^M \times \mathbb{I}_{H,o,q}^M$ corresponds to all the rays passing through a rectangular patch $\mathbb{I}_{W,n}^S \times \mathbb{I}_{H,o}^S$ on the sampling plane with the dimension $w_S/2 \times h_S/2$ from an existing pixel in the dataset $\mathbb{I}_{W,n,p}^M \times \mathbb{I}_{H,o,q}^M$.

On these intervals, we define the events $U_n = (u \in \mathbb{I}_{W,n}^S)$, $V_o = (v \in \mathbb{I}_{H,o}^S)$, $S_{np} = (s \in \mathbb{I}_{W,n,p}^M)$ and $T_{oq} = (t \in \mathbb{I}_{H,o,q}^M)$. We can then compute the probabilities the conditional probabilities of these events:

$$\begin{aligned} P(U_n \wedge V_o \wedge S_p \wedge T_q) &= \frac{1}{4} \frac{\sum_{i=n-1}^n \sum_{j=o-1}^o L_{ijk_{inp}l_{joq}}}{\sum_{i,j,k,l} L_{ijkl}} \\ P(T_q / (U_n \wedge V_o \wedge S_p)) &= \frac{\sum_{i=n-1}^n \sum_{j=o-1}^o L_{ijk_{inp}l_{joq}}}{\sum_{i=n-1}^n \sum_{j=o-1}^o \sum_l L_{ijk_{inp}l}} \\ P(S_p / (U_n \wedge V_o)) &= \frac{\sum_{i=n-1}^n \sum_{j=o-1}^o \sum_l L_{ijk_{inp}l}}{\sum_{i=n-1}^n \sum_{j=o-1}^o \sum_{k,l} L_{ijkl}} \\ P(V_o / U_n) &= \frac{1}{2} \frac{\sum_{i=n-1}^n \sum_{j=o-1}^o \sum_{k,l} L_{ijkl}}{\sum_{i=n-1}^n \sum_{j,k,l} L_{ijkl}} \\ P(U_n) &= \frac{1}{2} \frac{\sum_{i=n-1}^n \sum_{j,k,l} L_{ijkl}}{\sum_{i,j,k,l} L_{ijkl}} \end{aligned}$$

with $k_{inp} = p - (i - n + 1)s_w$ and $l_{joq} = q - (j - o + 1)s_h$.

The first probability concerns the choice of one cell, while the others represent the conditional probabilities of selecting one interval only.