



AALBORG UNIVERSITY
DENMARK

Aalborg Universitet

Performance Comparison of Techniques for Approximating Image-Based Lighting by Directional Light Sources

Madsen, C; Laursen, R

Published in:
Image Analysis

Publication date:
2007

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Madsen, C., & Laursen, R. (2007). Performance Comparison of Techniques for Approximating Image-Based Lighting by Directional Light Sources. In B. K. Ersbøll, & K. S. Pedersen (Eds.), *Image Analysis: 15th Scandinavian Conference, SCIA 2007, Aalborg, Denmark, June 10-14, 2007* (pp. 888-897). IEEE Computer Society Press.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Performance Comparison of Techniques for Approximating Image-Based Lighting by Directional Light Sources

Claus B. Madsen and Rune E. Laursen

Laboratory of Computer Vision and Media Technology, Aalborg University, Denmark
{cbm,rul}@cvmt.aau.dk
<http://www.cvmt.aau.dk>

Abstract. Image-Based Lighting (IBL) has become a very popular approach in computer graphics. In essence IBL is based on capturing the illumination conditions in a scene in an omni-directional image, called a light probe image. Using the illumination information from such an image virtual objects can be rendered with consistent shading including global illumination effects such as color bleeding.

Rendering with light probe illumination is extremely time consuming. Therefore a range of techniques exist for approximating the incident radiance described in a light probe image by a finite number of directional light sources. We describe two such techniques from the literature and perform a comparative evaluation of them in terms of how well they each approximate the final irradiance. We demonstrate that there is significant difference in the performance of the two techniques.

Key words: Augmented Reality, Image-Based Lighting, median cut, irradiance, real-time rendering, directional light sources

1 Introduction

Image-based approaches have gained widespread popularity in computer graphics, [1]. Image-based techniques have been used for 3D modeling of real scenes (Image-Based Modeling), for rendering from a bank of images with no 3D model whatsoever (Image-Based Rendering), and for modeling the complex illumination conditions in real scenes (Image-Based Lighting).

Image-Based Lighting (IBL) has become a frequently used technique to render a virtual object with illumination conditions that are consistent with those of a real scene given that the scene is distant, [2]. The idea is simply to use a camera to measure the light arriving at some point in the scene, the point at which you want to insert a virtual object. In practice people most often use a polished steel ball, place it somewhere in a scene, and take an image of it with a tele-lens from some distance away. The image thus contains information about how much light arrives at the position of the ball from all possible directions. Figure 3 shows example light probe images re-mapped to the longitude-latitude

format, where the full 360 degrees are represented along the horizontal (longitude) axis, and 180 degrees are represented along the vertical (latitude) axis. For construction and remapping of light probe images we use HDRShop 2.09, [3]. Handling the dynamic range of the light in the scene is done by acquiring the same view at different exposures, gradually lowering the exposure time until no pixels in the image are saturated. These multiple exposure are then fused into a single High Dynamic Range (HDR) floating point image, [4].

The light probe is a map of the incident radiance at the acquisition point. Each pixel in the map corresponds to a certain direction and solid angle, and together all pixels cover the entire sphere around the acquisition point. A light probe can thus also be called a radiance map, or an environment map. With this information virtual objects can be rendered into the scene with scenario consistent illumination e.g., [2, 5, 6]. Light probes can also be used to estimate the reflectance distribution functions of surfaces from images, as demonstrated in [7, 8]. For a review of illumination models in mixed reality see [9].

Actually using light probes for rendering is computationally very heavy. Using image-based lighting for a full global illumination rendering with path tracing is time consuming in order to reduce the noise level in the final rendering, simply because the light probe has to be treated as a spherical area light source enclosing the entire scene. To get a noise free estimate of the irradiance at a certain point requires thousands and thousands of samples of this area source, unless the light probe has very low frequency content.

To combat this problem several approaches have been proposed which take a light probe and attempt to approximate its illumination by a relatively low number of directional light sources. That is, the idea of these approaches is to find directions and the radiances of some number, say 64, directional light sources, such that the combined illumination from these sources approximate the combined illumination from the entire light probe. With such a directional light source approximation to a light probe, Image-Based Lighting using light probes can also be implemented in real-time applications taking into account that each source causes shadows to be cast.

The aim of the present paper is simply to test the performance of such approximation techniques in terms of how well they actually approximate the light probe for a given number of sources. The paper is organized as follows. Section 2 gives an overview of the approach and results in the paper. Section 3 introduces some concepts and terminology. Section 4 describes two completely different approaches to computing a directional light sources approximation. Section 5 then tests these two techniques in terms of their relationships between approximation error and number of sources used. Conclusions and directions for future research are given in section 6.

2 Overview of the idea of this work

Figure 1 shows a light probe together with the result from one of the approximation techniques studied in this paper. In this particular case the technique



Fig. 1. Result from running the Median Cut approximation technique using 8 directional sources on Galileo’s Tomb light probe. The light probe is obtained from [10]. Each rectangular region contains a red dot. This red dot marks the chosen direction for a particular directional source, and all the combined radiance from the region has been transferred to this particular source direction.

has been allocated 8 directional sources which it has then distributed across the light probe longitude-latitude map in an attempt to capture the radiance distribution of the original light probe. Naturally, the accuracy of the approximation depends on the number of sources allocated. The original light probe is simply W times H directional sources, where W is the number of pixels in the longitude direction, and H is the number of pixels in the latitude direction.

We then run any given technique on some light probe image to produce approximations with 2, 4, 8, 16, etc. light sources. Given these sets of approximated sources we compute what the resulting error in irradiance is compared to ground truth, which in this case is the irradiance computed by using the radiance from all pixels in the light probe. Figure 2 shows an example of true and approximated irradiance. In section 5 we test and compare two different approximation techniques on three different light probes.

The irradiance is chosen as the error measure in order to have a compact, quantitative performance measure, which is independent of surface characteristics, and independent on view point. For highly glossy surfaces approximating the light probe by a relatively low number of light sources will obviously lead to visible errors, whereas reflected radiance from diffuse surfaces will be correct if the approximated irradiance is correct.

3 Terminology

Before we proceed with the actual techniques and their performances we need to establish a small theoretical basis. Above we have used the term ”directional light source” a few times. There are correctness problems relating to using radiometric properties in conjunction with point and directional sources as they have no physical extent, [11, 12]. The purpose of the remainder of this section is to establish a physically correct terminology regarding light probe images.




Fig. 2. Left: ground truth irradiance for around 20000 normals distributed evenly on a sphere computed for the Galileo’s Tomb light probe, figure 1. Right: irradiances resulting from running the Median Cut source approximation technique to produce 8 directional sources. On print the difference may be visually subtle, but the average error is actually around 25 percent, and the maximum error is more than 75 percent.

In formal terms the light probe image is a spatially discrete measurement of the continuous function describing the incident radiance (measured in $W/(m^2 \cdot Sr)$), which in turn is a function of the incident direction. Let \mathbf{n} be the normal of a differential area surface, and let $\Omega_{\mathbf{n}}$ be the hemi-sphere defined by this normal. By integrating the incident radiance, $L(\boldsymbol{\omega})$, from the direction $\boldsymbol{\omega}$ over the hemi-sphere the total irradiance, $E(\boldsymbol{\omega})$, can be computed:

$$E(\mathbf{n}) = \int_{\Omega_{\mathbf{n}}} L(\boldsymbol{\omega})(\mathbf{n} \cdot \boldsymbol{\omega})d\boldsymbol{\omega} \quad (1)$$

which then is measured in W/m^2 . The term $d\boldsymbol{\omega}$ signifies the differential solid angle $d\boldsymbol{\omega} = |d\boldsymbol{\omega}|$ in the direction $\frac{d\boldsymbol{\omega}}{|d\boldsymbol{\omega}|}$.

For computational purposes it is beneficial to formulate these matters in terms of standard spherical coordinates. A direction in space is then written as $\boldsymbol{\omega}(\theta, \phi) = [\sin(\theta) \cos(\phi), \sin(\theta) \sin(\phi), \cos(\theta)]$, where θ is the angle the direction vector makes with the coordinate system z -axis (latitude), and ϕ is the angle the projection of the vector on the xy -plane makes with the x -axis. The irradiance from Eq. 1 then becomes:

$$E(\mathbf{n}) = \int \int_{(\theta, \phi) \in \Omega_{\mathbf{n}}} L(\theta, \phi)(\mathbf{n} \cdot \boldsymbol{\omega}(\theta, \phi)) \sin(\theta)d\theta d\phi \quad (2)$$

In this paper we will exclusively use the latitude-longitude mapping (LL mapping) of light probe images. Let the resolution of the LL light probe image be W by H pixels, and let u and v represent pixel coordinates in an image coordinate system with origin in the top left corner of the LL map, and v -axis oriented downwards. Thus the *middle row* in the image corresponds to the

equator of the unit sphere, i.e., corresponds to $\theta = \pi/2$, the top row corresponds to $\theta = 0$ and the bottom row corresponds to $\theta = \pi$. Moreover $\phi = 0$ corresponds to the *leftmost column*. Each light probe pixel, $P(u, v)$, represents the radiance in $W/(m^2 \cdot Sr)$ (if the light probe acquisition is radiometrically calibrated) from the direction given by $\boldsymbol{\omega}(u, v) = \boldsymbol{\omega}(\theta(v), \phi(u))$, where $\theta(v) = v\Delta_\theta$ and $\phi(u) = u\Delta_\phi$, where $\Delta_\theta = \pi/H$ and $\Delta_\phi = 2\pi/W$. The discrete version of Eq. 2 then becomes:

$$E(\mathbf{n}) \approx \sum_u \sum_v P(u, v)(\mathbf{n} \cdot \boldsymbol{\omega}(u, v)) \sin(\theta(v)) \Delta_\theta \Delta_\phi \quad (3)$$

where the summations are subject to the constraint that $(\theta(v), \phi(u)) \in \Omega_{\mathbf{n}}$, i.e., that the combinations of u and v represent pixels inside the region corresponding to the hemi-sphere defined by the surface normal \mathbf{n} .

From Eq. 3 it is evident that if every pixel, $P(u, v)$, in the LL map is scaled with $\Delta_\theta \cdot \Delta_\phi = 2\pi^2/(W \cdot H)$ and weighted by $\sin(\theta(v))$, we get a very simple summation. We therefore produce a new LL map, where each pixel $Q(u, v) = 2\pi^2 P(u, v) \sin(\theta(v))/(W \cdot H)$. The irradiance for a given normal is then simply computed as:

$$E(\mathbf{n}) \approx \sum_u \sum_v Q(u, v)(\mathbf{n} \cdot \boldsymbol{\omega}(u, v)) \quad (4)$$

where the summations again are subject to the constraint that $(\theta(v), \phi(u)) \in \Omega_{\mathbf{n}}$.

To recapitulate in a different way: Each pixel in the LL map acts as a small area light source subtending a solid angle of $A_p = 2\pi^2/(W \cdot H)$ [$Sr/pixel$]. By weighting each pixel by $\sin(\theta(v))$ we achieve "permission" to treat all pixels equally in the sense that we cancel out the effect of the non-uniform sampling density of the LL mapping (poles are severely over-sampled). By subsequently scaling by A_p we convert the solid angle domain from steradians to pixels. I.e., each $Q(u, v) = 2\pi^2 P(u, v) \sin(\theta(v))/(W \cdot H)$ measures the radiance in $W/(m^2 \cdot pixel)$, such that by performing a simple cosine weighted sum of pixels we directly get the irradiance contributed by the pixels involved in the sum (Eq. 4). Another way of putting it is: each pixel $Q(u, v)$ is an area light source contributing $Q(u, v)(\mathbf{n} \cdot \boldsymbol{\omega}(v, u))$ irradiance to the differential area surface with normal \mathbf{n} .

In the remainder of the paper we take the meaning of a directional light source to be a very small area light source (there are normally a lot of pixels in a light probe image). The direction to such a source is taken to be the direction to its center, and it is assumed that for each such source we know its radiance and its area.

4 Light probe approximation techniques

We have found three different approaches to finding a set of directional light sources which approximate a full radiance map in the form of a light probe. Two of the approaches are closely related and operate directly in the radiance space image domain of the light probe, in particular on the longitude-latitude mapping. The last technique is quite different in that it operates in *irradiance* space.

4.1 Median Cut

The LightGen, [13], and the Median Cut, [14], techniques operate directly in the image domain of the light probe. We have chosen to focus on the Median Cut technique since it is better documented than the LightGen technique, which only exists as a plugin to HDRShop. The median cut technique, [14], is conceptually wonderfully simple. The idea is to recursively split the LL map into regions of approximately equal summed radiance. Since the method splits all regions K times the technique produces 2^K sources, i.e., 2, 4, 8, 16, 32, etc. Figure 1 illustrates the result of running the technique on a light probe. The algorithm is as follows:

1. Add the entire light probe image to the region list as a single region.
2. For each region in the list, subdivide along the longest dimension such that its light energy is divided evenly.
3. If the number of iterations is less than K , return to step 2.
4. Place a light source at the centroid of each region, and set the light source radiance to the sum of the pixel values within the region.

The strength of this approach is that it is so straight forward, computationally light and easy to implement. The problems with this approach lies in two issues. The first issue is that it subdivides all regions at each iteration and depending on K there can be a large jump in the number of sources, which may be disadvantageous for real-time rendering with the approximated sources, where one would like as many sources as possible, but at the same time there is a performance limit in the graphics hardware. The second issue relates to step 4, where, for small K , and thereby large regions, a lot of radiance is transferred quite large distances over the sphere without any cosine weighting. This transfer could be done physically correct, but only for a single known surface normal. For arbitrary normals there is no alternative to just transferring the radiance and hope the regions are small enough that it does not constitute a grave error. This is naturally an invalid assumption for very small K .

4.2 Irradiance Optimization

The Irradiance Optimization technique by Madsen et al., [15], is significantly different from the previous one. While the Median Cut technique operates entirely on a pixel level in the light probe image, i.e., operate in radiance space, the Irradiance Optimization method operates in *irradiance* space.

The method is based on first using the original light probe image to compute the ground truth irradiance for a large number (M) of normal directions uniformly distributed across the unit sphere using Eq. 4. These M irradiance values constitute the goal vector in an optimization to estimate the parameters of N directional sources. Each source is defined by five parameters (RGB radiances and two angles for direction).

Given an estimate of these 5 times N parameters it is possible to compute the approximated irradiances for the M normals. Let L_i be the radiance of the i th

source, and let $\omega(\theta_i, \phi_i) = [\sin(\theta_i) \cos(\phi_i), \sin(\theta_i) \sin(\phi_i), \cos(\theta_i)]$ be the direction vector to the i th source. Furthermore, let A be a fixed, small area (in steradians) of each light source to accommodate physical correctness. A disc area light source of 1 degree has an area of $2.392 \cdot 10^{-4}$ steradian. Finally, let \mathbf{n}_k be the k th normal. The irradiance on a differential area surface with normal \mathbf{n}_k is then:

$$E(\mathbf{n}_k) = \sum_{i=1}^N L_i \cdot A \cdot (\mathbf{n}_k \cdot \omega(\theta_i, \phi_i)) \quad (5)$$

By comparing the approximated irradiances to the ground truth irradiances an error vector is obtained, which can be converted to a parameter update vector. The source estimation process is thus an iterative, non-linear optimization process based on Newton’s iterative method, since the Jacobian can be expressed analytically.

The strength of this method is that it inherently produces a configuration of directional source which result in irradiances that are similar to the irradiances achieved by full IBL using the entire light probe. Furthermore, this method can produce source configurations with any number of sources, not just a power of 2. The weakness of the method is that it is significantly more complicated to implement than the Median Cut method, and it is computationally much heavier. The estimation part takes a few minutes in Matlab, but computing the ground truth irradiances can be time consuming. We use HDRShop, [3], to precompute the irradiance map, and it takes on the order of 24 hours for a 1024x512 resolution irradiance map. We tried using the fast diffuse convolution method implemented in HDRShop (executes in seconds), but this method in itself is approximative and has too many artifacts to be used as ground truth.

5 Comparative evaluation

We have tested the two methods (Median Cut and Irradiance Optimization) on a number of qualitatively quite different light probe image shown in figure 3. We ran the methods on the test light probes, and produced directional source approximations with 2, 4, 8, 16, 32, 64, and 128 sources. The Irradiance Optimization technique can produce any number of sources, but was constrained to the source number cases which were also feasible for the Median Cut approach. We were unable to obtain a convergence on a 128 source solution with the Irradiance Optimization technique. This will be discussed later.

The evaluation is based on computing the irradiances resulting from the estimated set of sources for a large number (20480 resulting from a subdivision of an icosahedron) of surface normals evenly distributed on a unit sphere, and comparing them to the ground truth irradiances (computed using HDRShop). Figure 2 showed an example of true and approximated irradiance. For each color channel we then compute the mean and the maximum of the absolute differences between estimated and ground truth irradiances. Figure 4 shows curves representing average and maximum error for each of the two methods as a function of the number of light sources used. The errors are an average over RGB.




Fig. 3. The tested light probes. Top row: Galileo’s tomb in Florence, Italy, and St. Peter’s Cathedral, Rome, Italy. Both acquired from [10]. Bottom: Outdoor daylight scene. Everything below the horizon has been set to zero, so the light probe only accounts for the sun and the sky, not the ground. All probes are available via [10].

5.1 Discussion

The plots of mean and maximum irradiance errors in figure 4 clearly show that both approximation techniques, regardless of what light probe they are applied to, perform consistently better the more light sources the technique is allowed to employ. This simple fact means that in terms of accuracy it will always be better to choose a higher number of sources.

Secondly the experiments show that the Irradiance Optimization technique consistently performs much better than the Median Cut method. This is the case for both average and maximum error. Generally the Median Cut method requires 2 to 3 times as many sources to achieve the same error as the Irradiance Optimization technique. For rendering this is very important from a computational point of view, since it will always be an advantage to use as few sources as possible. Typically the Irradiance Optimization gets below 5 percent error for approximately 5 to 6 sources, whereas the Median Cut technique requires on the order of 20 sources or more.

The Irradiance Optimization technique could not converge when the number of sources comes above some threshold (64). This is due to the fact that when the number of sources grows too high there is too little energy for some sources to latch on to. Early in the iterations the dominant sources become stable, leaving ever smaller amounts of energy to distribute among the rest of the sources. At the same time the sources tend to repel each other when they distribute across the sphere, because each source has a semi-spherical ”footprint” (each source is like shining a torch on a sphere) and sources will be reluctant to overlap footprints too much.




Fig. 4. Mean and max irradiance error in percent for the two tested techniques for three different light probes. Left column: Irradiance Optimization. Right column: Median Cut. Top to bottom row: Galileo, St. Peter's, and the Sky probe, respectively. See figure 3 for the three light probes.

6 Conclusions

We have demonstrated a significant difference in the performance of available techniques to approximate light probes with a set of directional light sources. Tests clearly demonstrate that the Irradiance Optimization technique requires substantially fewer sources to achieve the same error level as the other technique.

In terms of rendering speed it will always be an advantage to have as few light sources as possible and still achieve visibly acceptable performance. If a rendering is solely for visual purposes it may not be crucial whether the irradiance at a point is 1 percent or 5 percent wrong, but renderings can be also used in more radiometrically challenging contexts such as for inverse methods, aim-

ing at estimating surface reflectance parameters from images, [7, 8]. For inverse problems the accuracy in the approximation can be very important.

Acknowledgments

This research is funded by the CoSPE project (26-04-0171) under the Danish Research Agency. This support is gratefully acknowledged.

References

1. Oliveira, M.M.: Image-based modelling and rendering: A survey. *RITA - Revista de Informatica Teorica a Aplicada* **9**(2) (October 2002) 37 – 66 Brazilian journal, but paper is in English.
2. Debevec, P.: Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In: *Proceedings: SIGGRAPH 1998*, Orlando, Florida, USA. (July 1998)
3. Debevec et al., P.: Homepage of HDRShop, www.hdrshop.com
4. Debevec, P., Malik, J.: Recovering high dynamic range radiance maps from photographs. In: *Proceedings: SIGGRAPH 1997*, Los Angeles, CA, USA. (August 1997)
5. Debevec, P.: Tutorial: Image-based lighting. *IEEE Computer Graphics and Applications* (March/April 2002) 26 – 34
6. Gibson, S., Cook, J., Howard, T., Hubbard, R.: Rapid shadow generation in real-world lighting environments. In: *Proceedings: EuroGraphics Symposium on Rendering*, Leuven, Belgium. (June 2003)
7. Yu, Y., Malik, J.: Recovering photometric properties of architectural scenes from photographs. In: *Proceedings: SIGGRAPH 1998*, Orlando, Florida, USA. (July 1998) 207 – 217
8. Yu, Y., Debevec, P., Malik, J., Hawkins, T.: Inverse global illumination: Recovering reflectance models of real scenes from photographs. In: *Proceedings: SIGGRAPH 1999*, Los Angeles, California, USA. (August 1999) 215 – 224
9. Jacobs, K., Loscos, C.: State of the art report on classification of illumination methods for mixed reality. In: *EUROGRAPHICS*, Grenoble, France (September 2004)
10. Debevec, P.: Homepage of Paul Debevec, www.debevec.org/probes
11. Phar, M., Humphreys, G.: *Physically Based Rendering – From Theory to Implementation*. Elsevier (2004)
12. Jensen, H.W.: *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters (2001)
13. Cohen, J.M., Debevec, P.: The LightGen HDRShop plugin. (2001) www.hdrshop.com/main-pages/plugins.html
14. Debevec, P.: A median cut algorithm for light probe sampling. In: *Proceedings: SIGGRAPH 2005*, Los Angeles, California, USA. (August 2005) Poster abstract.
15. Madsen, C.B., Sørensen, M.K.D., Vittrup, M.: Estimating positions and radiances of a small number of light sources for real-time image-based lighting. In: *Proceedings: Annual Conference of the European Association for Computer Graphics, EUROGRAPHICS 2003*, Granada, Spain. (September 2003) 37 – 44