

Semi-Dynamic Light Maps

Bekir Öztürk
Middle East Technical University
bekir.ozturk@metu.edu.tr

Ahmet Oğuz Akyüz
Middle East Technical University
akyuz@ceng.metu.edu.tr



Figure 1: Our approach allows modifications to the scene during run-time such as hiding and showing objects and changing the color and intensity of any number of light sources.

ABSTRACT

Light mapping is an important optimization technique, in which the lighting of a scene is precomputed into a texture during a stage known as light baking. However, the primary drawback of this technique is that lights and objects must be static. Our work relaxes several important requirements of light mapping, such as the requirement of the color, intensity, and on-off state of lights as well as the presence or absence of shadow casting objects to be static.

CCS CONCEPTS

•Computing methodologies →Rendering;

KEYWORDS

Light maps, shadows, light baking

ACM Reference format:

Bekir Öztürk and Ahmet Oğuz Akyüz. 2017. Semi-Dynamic Light Maps. In *Proceedings of SIGGRAPH '17 Posters, Los Angeles, CA, USA, July 30 - August 03, 2017*, 2 pages.
DOI: 10.1145/3102163.3102202

1 INTRODUCTION

Real-time detailed lighting calculations is one of the most time consuming tasks for GPUs in forward rendering. Real-time rendering of high number of lights and objects is not possible in many situations, especially in mobile devices. One of the techniques that allow

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGGRAPH '17 Posters, Los Angeles, CA, USA

© 2017 Copyright held by the owner/author(s). 978-1-4503-5015-0/17/07...\$15.00
DOI: 10.1145/3102163.3102202

high quality lighting involves pre-calculating the incoming light and storing it in a texture map. During run-time, the final illumination value of each pixel is fetched from this texture map, freeing the renderer from doing lighting computations. This technique is known as light mapping [Rasmussen et al. 2010].

Contrary to real-time shadow rendering methods, light map generation is an offline process and can be done using highly realistic and computationally intensive methods. This allows features such as indirect illumination and soft shadows to be simulated even in low power GPUs. Generation of light maps can be done with almost any view-independent global illumination method such as path tracing [Kajiya 1986] or photon mapping [Jensen 1996].

The primary drawback of this technique is that lights and objects must be static as lighting computations are done offline [Luksch et al. 2013]. In this paper, we offer a technique to relax this constraint by allowing modifications to the color, intensity, and on-off state of the lights as well as the presence or absence of shadow casting objects (Figure 1).

2 APPROACH

When using light maps, object visibility cannot be changed during run-time, as the contribution of an object to the light map is typically unknown. By the same principle, light intensity for each color channel and whether the light is on or off cannot be altered.

In our method we solve this problem in the following way. For each light source, we first store its contribution to the scene assuming that the light intensity is (1, 1, 1) and all shadow casters are disabled. Then, each shadow caster is enabled one by one and the light map is recomputed. The changed regions with respect to the original light map are copied to a texture atlas. As a result, for each light source, we obtain a base light map and a texture atlas, which contains the shadow effect of all objects. The remapping for each




Figure 2: Overall workflow of our algorithm.

object between these two textures is stored as auxiliary data (this only takes up 20-30 bytes per object).

During run-time three types of modifications are possible: (1) Turning the light on and off, (2) turning the objects on and off, and (3) changing the per channel intensity of the light. If a light is turned off, we simply ignore the base light map of that light source. If the

light is on, we add its value from the base light map but subtract the regions that it cannot make a contribution by looking-up into the texture atlas based on which objects are enabled. If the light intensity is changed, we simply modulate this net result by the actual intensity of the light source. This process is explained in Figure 2.

In practice, not every light in the scene is required to be modifiable. As long as no modifiable objects are illuminated by them, it is unnecessary to calculate a separate contribution map for these lights. To prevent producing redundant textures, we create a single light map called the global base light map (GBLM). The only requirement for GBLM is that it does not illuminate objects that can cast shadows and can be turned off. The final light map value of each pixel is computed by using the following formula:

$$L(u, v) = B(u, v) + \sum_{l=1}^L I_l \max\left(0, C_l(u, v) - \sum_{m=1}^M S_{m \rightarrow l}(u, v)\right)$$

$L(u, v)$: Final illumination value for a point in light map

$B(u, v)$: Contribution of global base light map

$C_l(u, v)$: Contribution of base light map of light l

$S_{m \rightarrow l}(u, v)$: Effect of object m on light l

$I_l(u, v)$: Per channel intensity of light l

(1)

As in standard light mapping, this illumination value can be combined with any other texture to obtain the final object color for each pixel.

3 CONCLUSIONS AND FUTURE WORK

In this work, by storing the effects of each object and light to the illumination of the scene, we removed some of the important restrictions of light maps. While our approach requires more texture memory and extra texture look-ups compared to the standard light mapping, in our experiments we observed real-time performance. As future work, we are planning to investigate how this technique can be extended to fully dynamic objects and light sources.

REFERENCES

- Henrik Wann Jensen. 1996. Global Illumination Using Photon Maps. In *Proceedings of the Eurographics Workshop on Rendering Techniques '96*. Springer-Verlag, London, UK, UK, 21–30. <http://0-dl.acm.org.library.metu.edu.tr/citation.cfm?id=275458>. 275461
- James T. Kajiya. 1986. The Rendering Equation. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '86)*. ACM, New York, NY, USA, 143–150. DOI: <http://dx.doi.org/10.1145/15922.15902>
- Christian Luksch, Robert F. Tobler, Ralf Habel, Michael Schwärzler, and Michael Wimmer. 2013. Fast Light-map Computation with Virtual Polygon Lights. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D '13)*. ACM, New York, NY, USA, 87–94. DOI: <http://dx.doi.org/10.1145/2448196.2448210>
- Jim Rasmusson, Jacob Ström, Per Wennersten, Michael Doggett, and Tomas Akenine-Möller. 2010. Texture Compression of Light Maps Using Smooth Profile Functions. In *Proceedings of the Conference on High Performance Graphics (HPG '10)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 143–152. <http://0-dl.acm.org.library.metu.edu.tr/citation.cfm?id=1921479.1921501>