# Shadow and Environment Maps

Lecture #8: Monday, 28 September 2009
Lecturer: Ravi Ramamoorthi
Scribe: Aaron Hong

## 1 Introduction

Shadows and complex lighting are an essential part of realistic rendering, however they can be computationally expensive (such as path tracing shadows). Therefore image-based methods, *Shadow* and *Environment* Mapping (See figure 1), are used to accelerate and approximate these effects in interactive real-time rendering.

Shadow mapping is an old technique [Williams 78], but very widely used and developed with many extensions and implementations (Renderman). Environment mapping is also a very old technique [Blinn and Newell 76], but has only seen more use and development in recent times due to previous hardware limitations.

## 2 Shadow Maps

There are many real-time shadow techniques. Common techniques are usually problematic and do not work for all situations. Projected planar shadows use a projective transform to render the casting object as the shadow, but this technique only works well

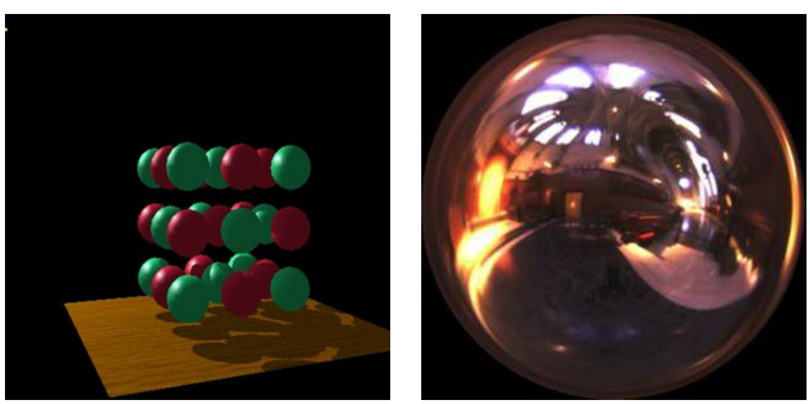

(a) Shadow Map (b) Environment Map

Figure 1: Image-based Rendering Techniques

for flat surfaces. In the shadow volume technique, the light shoots rays through the vertices of the casting object to construct geometry, which volume is used to test for shadowing. Light maps are unsuitable for dynamic shadows. Finally, in general it is just hard to get everything to shadow everything right. See figure 2.

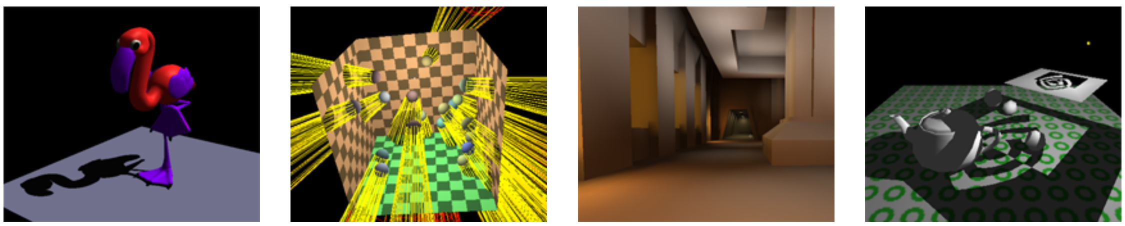

(a) Projected Planar (b) Shadow Volumes (c) Light Maps (d) Hybrid

Figure 2: Common Real-time Shadow Techniques

However, the most common software rendering technique, shadow mapping, works effectively in many situation. Developed by Lance Williams in 1978, along with other completely image-space algorithms, the technique precomputes images to avoid retracing shadow rays for every frame. The basic idea is to take an image of what the light sees in the scene and then compare it to what the eye sees by projecting the depth back to the light. The places where it is seen by both the light and the eye we know is lit. Otherwise, it is in the shadow. See figure 3.

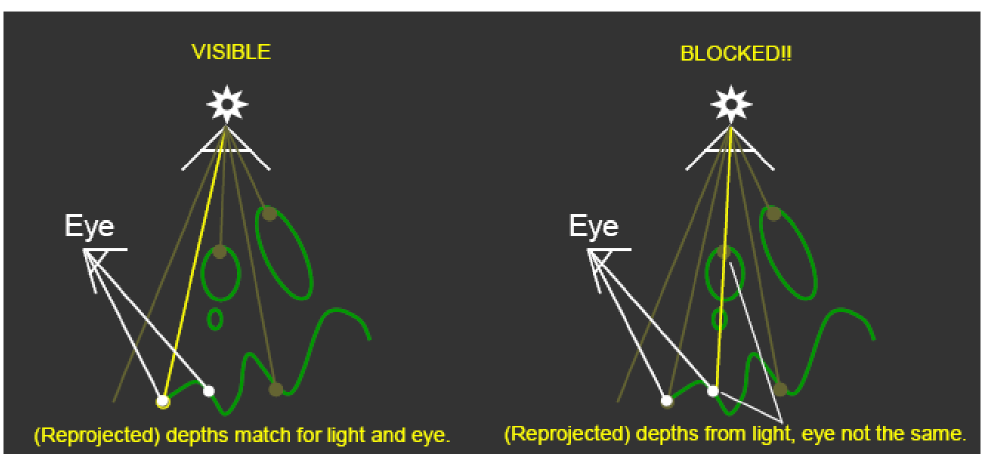


Figure 3: Shadow Mapping

The saved image from the light is called the shadow map, which is essentially a bitmap where depth corresponds to a gray scale value. See figure 4.

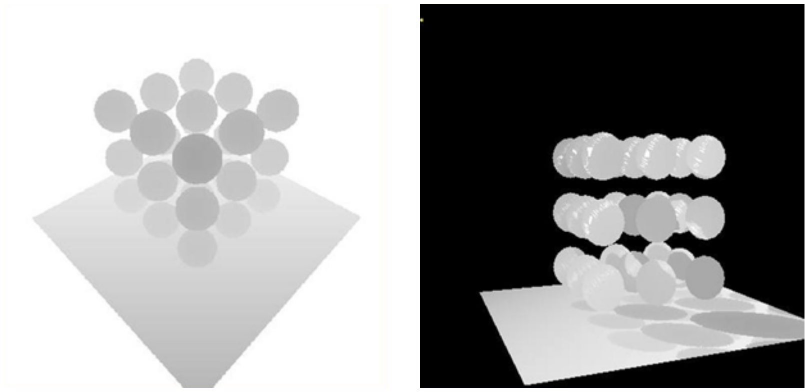

(a) Depth Buffer from Light (b) Projecting the depth onto the eye's view

Figure 4: Visualizing Shadow Mapping

## 2.1 Map Filtering

Normal texture filtering just averages the color components, however that would be inadequate for our uses in shadow mapping. If we average the color components of a depth map we would instead get depths of things that never existed. Essentially we will be changing the position of the objects in the scene and potentially creating new surfaces.

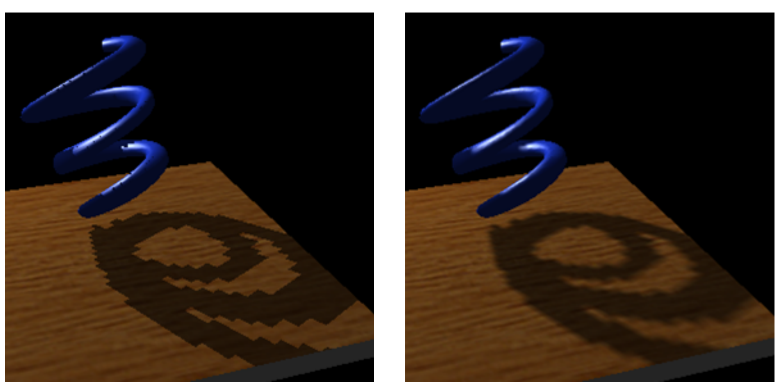

(a) GL_NEAREST: blocky (b) GL_LINEAR: antialiased

Figure 5: Low shadow map resolution used to heighten filtering artifacts.

Instead another method is used called "percentage closer" filtering [Reeves 87], where the hardware takes multiple samples and then averages the results of the comparisons to get a percentage seen value. This provides anti-aliasing at shadow map edges and can be used to fake soft shadows. See figure 5.

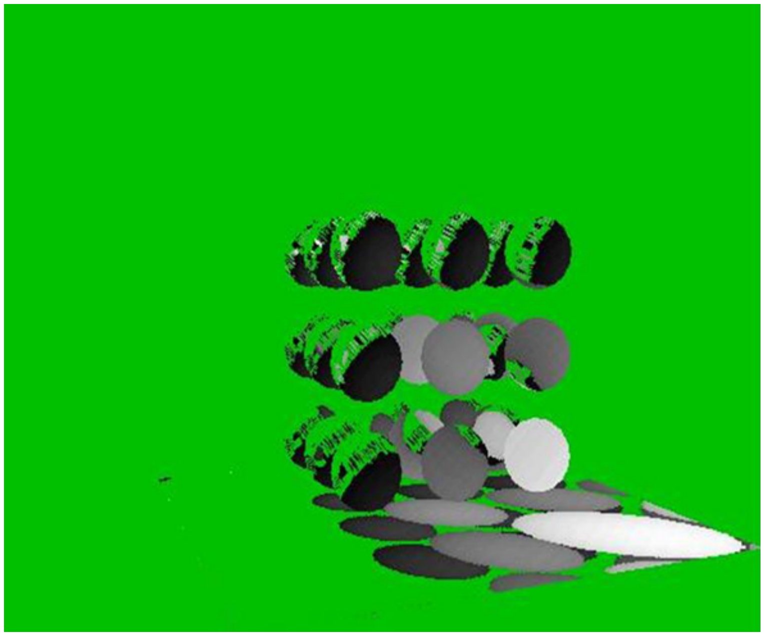

Figure 6: Green is where the light planar distance and the light depth map are approximately equal. Note the aliasing due to the tolerance.

### 2.2 Limitations

Although shadow maps are one of the most viable real-time shadow techniques, it still has it's limitations. One, it is generally only applicable to point lights and produces hard shadows. The quality of the shadow also depends on the resolution of the shadow map. Extreme angles or low resolutions reveal aliasing effects. Finally, there is equality comparisons of floating point depth values which means there are issues of scale, bias, and tolerance (See figure 6). There however have been many extension including adaptive sampling.

## 3 Environment Maps

Environment maps are a way of creating and using complex lighting through the use of images. The idea was conceived as early as 1976 when Blinn and Newell demonstrated reflection maps (See figure 7a). Horn contributed the reflectance maps in 1977. Miller and Hoffman are other major contributors, who developed irradiance and Phong reflection maps (See figure 7c). There has been abundant development in recent years, including work from Paul Debevec, who has contributed his famous high dynamic range environment maps including the Eucalyptus Grove and Grace Cathedral.

Since environment maps are by nature an image, it also makes sense to collect environment maps from the real world through photographs and different techniques (See figure 7b). We will see how we can sample an environment map in a few lectures from now. First we will talk about irradiance environment maps, which be used to light diffuse surfaces.

(a) Reflection Map

(b) Ornament used to capture map.

(c) Composited image with 3D element using a environment map

Figure 7: Environment Map

## 3.1 Irradiance Environment Map

Some assumptions we will make for lighting using the environment map is that the surface of the object is diffuse. Also the illumination is distant, so that the lighting can generalize to all points on the surface of an object. Finally, there is no shadowing or interreflection. Hence, irradiance is a function of the surface normal. The diffuse reflection equation is just the reflectance $\rho$ times the irradiance $E$ equals the radiosity $B$ (See figure 8).

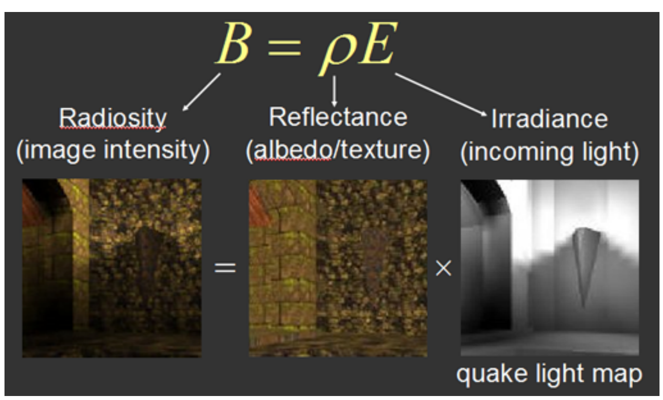


Figure 8: Diffuse Equation

There is another property of diffuse surfaces we can exploit in order to do real-time rendering, which is a diffuse surface acts like a low pass filter. We will use an analytic expression for the irradiance in terms of spherical harmonic coefficients (See figure 9) of the lighting to approximate diffuse surfaces [Ramamoorthi and Hanrahan 01].

**Definition 1.** *Analytic Irradiance Formula*

$$E_{lm} = A_l L_{lm} \tag{1}$$

$$A_l = 2\pi \frac{(-1)^{\frac{1}{2}-1}}{(l+2)(l-1)} \left[ \frac{l!}{2^l (\frac{1}{2}!)^2} \right] \tag{2}$$

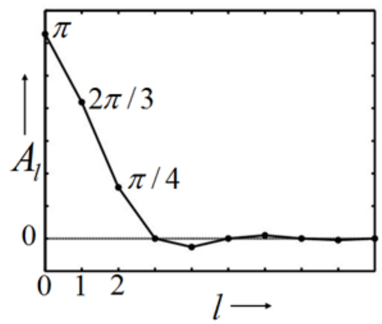


Figure 9: Graph of Irradiance Approximation

Since it acts like a low pass filter, we only need to use the lowest-frequency modes of illumination, which corresponds to using only 9 spherical harmonic coefficients. For any illumination, on average the error is less than 3% [Basri and Jacobs 01]. See figure 10.

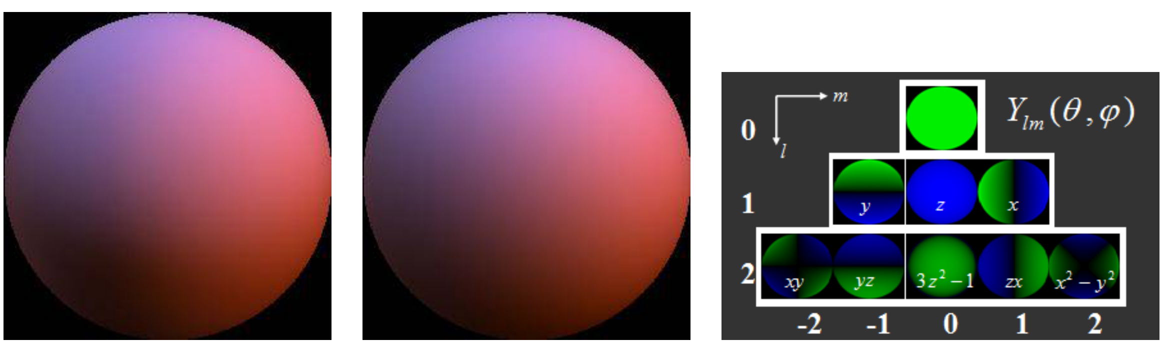


(a) Exact image (b) Order 2 with 9 Terms (c) Spherical harmonics

Figure 10: 9 Parameter Approximation

This can be written as a simple procedural method, which makes it easy for rendering either in software or hardware because it only requires matrix-vector multiplication and a dot product. Due to it's simplicity, it is widely implemented and used in games (AMPED for Microsoft Xbox) and movies (Pixar and Framestore CFC).

**Definition 2.** *Irradiance Approximation in Matrix Form*

$$E(n) = n^t M n \tag{3}$$

# 4 Conclusion

Both Shadow Mapping and Environment Mapping are popular methods used for real-time rendering. However, each have their own limitations and do not account for all visual effects. A combination of both methods give many of the realistic effects we want, but ultimately is a non-trivial problem. For more information about rendering shadows from sampling environment maps see Ben-Artzi and Ramamoorthi 2004. There have been many extensions to do more advance effects, but they still exist today in there venerable forms.

# References

[1] Lance Williams, *Casting Curved Shadows on Curved Surfaces*. Proc. Siggraph '78, Computer Graphics, pp. 270-274, August 1978.

[2] James F. Blinn, Martin E. Newell, *Texture and reflection in computer generated images*. Communications of the ACM, v.19 n.10, p.542-547, Oct. 1976

[3] Reeves, W.T., Salesin, D. H., and Cook, R. I., *Rendering Antialiased Shadows with Depth Maps*. Proc. Siggraph '87, Computer Graphics, pages 283-291, July 1987.

[4] Ramamoorthi, Ravi and Hanrahan, Pat, *A signal-processing framework for reflection*. ACM Trans. Graph., v.23 n.4, p.1004-1042, 2004

[5] Basri, R. and Jacobs, D., *Photometric stereo with general, unknown lighting*. IEEE International Conference on Computer Vision, Kauai, II:374-381. 2004