Irradiance Rigs

Hong Yuan
UMass Amherst

Derek Nowrouzezahrai
University of Toronto

Peter-Pike Sloan
Disney Interactive Studios

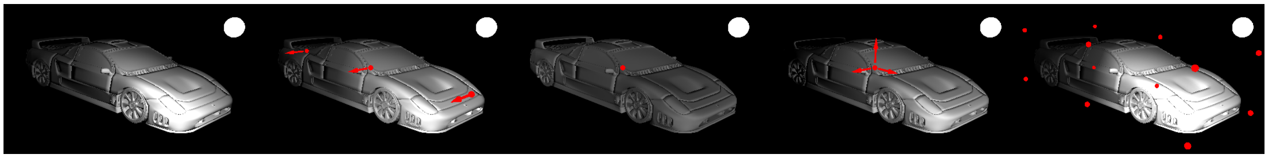

Figure 1: *Mixed rigs capture more accurate spatial variation than traditional rigs and nearly match ground truth quality.*

Abstract

When precomputed lighting is generated for static scene elements, the incident illumination on dynamic objects must be computed in a manner that is efficient and that faithfully captures the near- and far-field variation of the environment's illumination. Depending on the relative size of dynamic objects, as well as the number of lights in the scene, previous approaches fail to adequately sample the incident lighting and/or fail to scale. We present a principled, error-driven approach for dynamically transitioning between near- and far-field lighting. A more accurate model for sampling near-field lighting for disk sources is introduced, as well as far-field sampling and interpolation schemes tailored to each dynamic object. Lastly, we apply a flexible reflectance model to the computed illumination.

1 Irradiance Rigs

We are interested in rigs that reconstruct Spherical Harmonic (SH) irradiance coefficients at points p in a bounded volume around a dynamic object. The spatially-varying irradiance coefficients $\mathbf{l}_i$ can be computed every frame using an abstract model function $M(p, \mathbf{c})$, where $\mathbf{c}$ is a vector of model parameters for a given object.

The simplest type of rig for a dynamic object is a spatially-constant SH function. Projecting point or spherical light sources into SH is simple [Sloan 2008], and this rig's only parameter is the SH projection of the light sources as observed from the center of the object, $\mathbf{l}_c$. Concretely, this rig model is $\mathbf{l}_i = M(p, \mathbf{l}_c) = \mathbf{l}_c$.

The simplest *spatially varying* rig augments the "centroid-lighting" rig with a linear model of how lighting changes in space based on its gradient: $\mathbf{g}_x$, $\mathbf{g}_y$, $\mathbf{g}_z$. And so the model for this rig is $\mathbf{l}_i = M(p, \{\mathbf{l}_c, \mathbf{g}_x, \mathbf{g}_y, \mathbf{g}_z\}) = \mathbf{l}_c + p_x\mathbf{g}_x + p_y\mathbf{g}_y + p_z\mathbf{g}_z$. The gradient can be evaluated efficiently for spherical lights and computed numerically for other types of sources.

Another common rig fixes a lattice around an object, where lighting is sampled at each lattice point c_j and tri-linearly interpolated to reconstruct irradiance on the surface of the object. Given an arbitrary reconstruction kernel $b(p)$, this model is $\mathbf{l}_i = M(p, \mathbf{c}) = \sum_j b_j(p - c_j)\mathbf{c}_j$, where $\mathbf{c}_j$ is the SH irradiance at point c_j. This can be evaluated efficiently in hardware with volume textures.

For objects with different spatial scales along each axis, it can often be advantageous to employ a rig with different models for each dimension. In Figure 1, the car's surface irradiance varies at different rates across its length, height and width. We apply a three-sample lattice model for the length, a constant model for the height, and a gradient model for the width. This *mixed rig* model's parameter vector is composed of only six SH vectors, yet generates results comparable to a $8\times4\times2 = 64$-sample lattice rig and outperforms a 12-sample rig. We compare to traditional (e.g., constant [Smedberg and Wright 2009]) rigs and a ground-truth analytic evaluation.

Error driven criteria are used to transition lights between analytic evaluation and evaluation in the rig.

We have experimented with several other approaches, most notably, coupling radial basis functions with linear polynomials, combining functions and gradients in a lattice, as well as a form of hermite interpolation; however, all of these approaches performed poorly from a quality-per-unit-performance metric.

2 Disk Light Source Model

When using an irradiance rig, light sources near the character are evaluated directly. Unfortunately, the common approximation for disk light sources [Wallace et al. 1989] has extremely high error in this situation. We alternatively tabulate the response to a disk light source in SH and use this data to evaluate irradiance at runtime. Two small textures and a simple shader suffice to perform this computation efficiently, and we model the problem in a coordinate frame that eliminates three degrees of freedom in the SH response (whereas direct tabulation would require an unwieldily 4D table.)

3 Wrap Lighting

Wrap reflectance models are a commonly used approach for convincing character shading. We generalize Valve's model [Mitchell et al. 2006] while retaining its important properties. Our new parameterized model, $f(\theta|a) = ((\cos\theta + a)/(1 + a))^{1+a}$, has quadratic per-band SH scaling coefficients

$$\mathbf{f} = \left[\frac{2(1+a)}{2+a}, \frac{4(1+a)}{(2+a)(3+a)}, \frac{2(1+a)\left(3-2a+a^2\right)}{(2+a)(3+a)(4+a)} \right] .$$

This model can decay faster than a diffuse BRDF while retaining its first two derivatives, and reduces to a clamped cosine lobe at $a = 0$.

References

MITCHELL, J., MCTAGGART, G., AND GREEN, C. 2006. Shading in Valve's Source Engine. In *Advanced Real-Time Rendering in 3D Graphics and Games, SIGGRAPH '06*.

SLOAN, P.-P. 2008. Stupid Spherical Harmonics (SH) Tricks. In *GDC '08*.

SMEDBERG, N., AND WRIGHT, D. 2009. Rendering Techniques in Gears of War 2. In *GDC '09*.

WALLACE, J. R., ELMQUIST, K. A., AND HAINES, E. A. 1989. A Ray Tracing Algorithm for Progressive Radiosity. In *Proceedings of SIGGRAPH '89*.

SIGGRAPH 2010, Los Angeles, California, July 25 – 29, 2010.
ISBN 978-1-4503-0210-4/10/0007