

Directional Lightmap Encoding Insights

Peter-Pike Sloan
Activision Publishing
ppsloan@activision.com

Ari Silvennoinen
Activision Publishing
Ari.Silvennoinen@
activision.com




Figure 1: Left is AHD interpolation, right is new encoding. There is a static light near the ground shadowed by a pillar, and a light bouncing off the back wall that has higher albedo. Interpolating AHD, you can see artifacts in the transition from lit to shadowed. Using the new encoding, the lighting interpolates properly. The images are looking at lighting only, no albedo. ©Activision Publishing, Inc.

ABSTRACT

Lightmaps that respond to normal mapping are commonly used in video games. This short paper describes a novel parameterization of a standard lightmap encoding, Ambient Highlight Direction (AHD) — a model for directional irradiance consisting of ambient and directional light — that eliminates common interpolation artifacts. It also describes a technique for fitting the AHD model to lighting represented as spherical harmonics, where the unknown model parameters are solved in the null space of the constraint that irradiance is preserved.

CCS CONCEPTS

• Computing methodologies → Rendering;

KEYWORDS

lightmaps, global illumination, spherical harmonics

ACM Reference Format:

Peter-Pike Sloan and Ari Silvennoinen. 2018. Directional Lightmap Encoding Insights. In *SIGGRAPH Asia 2018 Technical Briefs (SA '18 Technical Briefs)*, December 4–7, 2018, Tokyo, Japan. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3283254.3283281>

SA '18 Technical Briefs, December 4–7, 2018, Tokyo, Japan

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *SIGGRAPH Asia 2018 Technical Briefs (SA '18 Technical Briefs)*, December 4–7, 2018, Tokyo, Japan. <https://doi.org/10.1145/3283254.3283281>.

1 INTRODUCTION

Since global illumination is often too expensive to compute in real-time, it is common for video games to precompute some light transport paths and store them on surfaces for static objects and in volumes for dynamic objects and effects. Lightmaps represent lighting response stored in textures and are often directional [Chen 2008; McTaggart 2004]: the data stored at each surface location can be evaluated using a higher frequency normal map, to more accurately represent the surface response. This is commonly used for diffuse transport, but glossy reflectance can also be handled. Different representations are used: spherical harmonics, or subsets of spherical harmonics restricted to the hemisphere, spherical Gaussians [Neubelt and Pettineo 2015] and simple models of incident radiance. In this paper we focus on a popular model, Ambient Highlight Direction, or AHD for short. It encodes a single High Dynamic Range (HDR) ambient color, a highlight direction and HDR highlight color for the directional light. While various forms of spherical harmonics often have higher quality, they require significantly more coefficients, particularly in HDR lighting environments. Unfortunately, AHD is a non-linear representation, so simply interpolating the parameters can result in visual artifacts. We present two alternate representations that explicitly encode irradiance and interpolate it linearly, which is correct, pushing the non-linearities outside of this constraint. We show why the artifacts occur when naively interpolating AHD parameters and how our parameterization addresses them.

It is common to use spherical harmonics internally when pre-computing lighting data, even if the final format does not use them. Fitting AHD parameters from SH is relatively straightforward, but

we have found it important to incorporate two constraints when doing so:

- The irradiance should be exactly captured, so that when no normal maps are used the resulting lighting is what one would expect.
- The two HDR colors should be strictly non-negative, doing something reasonable when one goes to zero.

2 AMBIENT AND HIGHLIGHT DIRECTION

The AHD model represents lighting as a combination of *ambient color* C_a and a *directional light* with highlight direction \mathbf{d} and highlight color C_d . This was used to represent spherical functions for dynamic objects in Quake III [id Software 1999], but it seems to have been independently used for lightmaps in multiple video games [Iwanicki 2013; Lazarov 2011]. Using the AHD model, irradiance values $I(\mathbf{n})$ can be reconstructed in any arbitrary direction \mathbf{n} – given usually by a detail normal map – by

$$I(\mathbf{n}) = C_a + \max(0, \mathbf{n} \cdot \mathbf{d})C_d. \quad (1)$$

The AHD representation of the illumination, (C_a, \mathbf{d}, C_d) , is compact and has a strong response to normals, often exaggerated compared to ground truth, that is preferred by the artists.

Unlike spherical harmonics or the variant of spherical Gaussians used in games [Neubelt and Pettineo 2015], AHD is a non-linear model. The mixture of two AHD values is not representable as a single AHD unless the directions are identical. Blending AHD-tuples is typically done by treating the color channel components as linear and computing a normalized and (luminance) weighted average direction for the highlight. Unfortunately, this makes leveraging the texture filtering units in modern GPUs challenging and can cause artifacts when AHD values are reconstructed via linear texture filtering. The directions are often stored in tangent space in a parameterization that only represents normalized vectors, or renormalized after interpolation. If the highlight direction \mathbf{d} changes quickly between texel samples, linear interpolation of the AHD values results in overshooting the irradiance, even without any detail normal maps (Figure 2). Some prior work [Iwanicki 2013] would filter the highlight direction, but this results in a loss of contrast in the normal maps. These interpolation problem are even more pronounced when trying to use it for volumetric models of lighting [Hooker 2016].

2.1 Alternate Parametrization

To alleviate the issues with linear interpolation with the standard AHD model, we propose a new parametrization based on the observation that we can require the reconstructed irradiance to be exact in the simplest case, i.e., without any detail normal maps by a simple change of variables. The irradiance $I_z = I(\mathbf{z})$ in the local z -direction $\mathbf{z} = [0, 0, 1]^T$ is given by

$$I_z = C_a + \max(0, \mathbf{z} \cdot \mathbf{d})C_d \quad (2)$$

$$= C_a + d_z C_d. \quad (3)$$

Now, we can reparametrize the AHD representation using the following mapping:

$$(C_a, \mathbf{d}, C_d) \mapsto (C_a, \mathbf{d}, I_z). \quad (4)$$

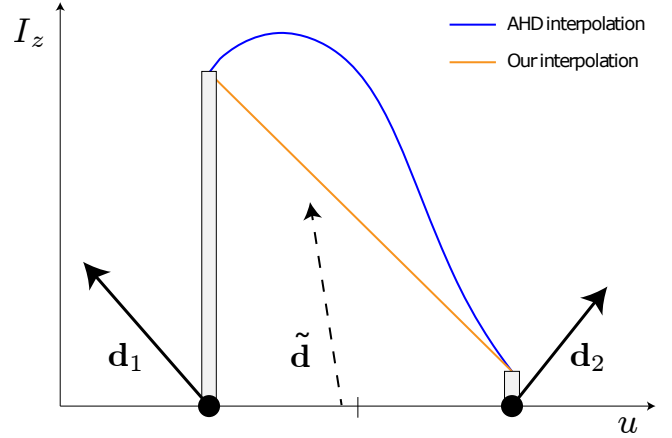


Figure 2: Simple 1D example. The texel on the left has a directional light of unit intensity pointing to the left and on the right one of zero intensity pointing to the right. While irradiance should be linear, it overshoots when the highlight direction is encoded in tangent space and interpolated. The black arrows represent the highlight directions for the two texels, the vertical axis is irradiance in local Z , the horizontal axis interpolates between the two texels.

This remapping makes the irradiance in the local z -direction I_z linear by construction and it is invertible at texel centers. The highlight color can now be reconstructed *after* interpolation by

$$C_d = \frac{I_z - C_a}{d_z}, \quad (5)$$

which gives us a AHD-representation which can be used for shading as usual. It's useful to consider what happens when interpolating between strongly tilted directions using the new parametrization: the local z -component d_z of the interpolated highlight direction \mathbf{d} increases, causing the highlight color intensity C_d to be reduced, which eliminates the typical interpolation artifacts (Figure 2).

2.1.1 Ambient Color Compression. The new parametrization is defined for RGB-triples. For further compression, we can drop the RGB ambient color C_a and replace it with a scalar luminance ratio of the ambient and local irradiance colors, giving us the following parametrization

$$(C_a, \mathbf{d}, C_d) \mapsto (L(C_a)/L(I_z), \mathbf{d}, I_z), \quad (6)$$

where L is the luminance function. In particular, this representation guarantees that the RGB irradiance value in the local z -direction is reconstructed without error even without storing the ambient color. The ambient color is reconstructed after interpolating the luminance ratio and local z -irradiance by

$$C_a = \frac{L(C_a)}{L(C_d)} I_z, \quad (7)$$

and the highlight color C_d is given by Equation (5). Since the luminance ratio is normalized to $[0, 1]$, it can be stored using a low dynamic range fixed point texture format. This can also be done with RGB values, requiring only a single HDR color to be encoded.

3 FITTING AHD DATA FROM SPHERICAL HARMONICS

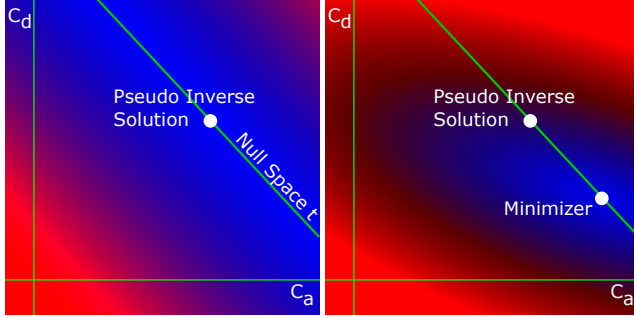


Figure 3: Axes are C_a and C_d , error is relative to I_z , the diagonal green line is the space of solutions to the constraint equation for a specific I_z , with the direction being the null-space of the I_z constraint. One point is the solution from the pseudo-inverse, which is the location on the line closest to the origin. The image on the left is error relative to I_z and on the right integrated squared error relative to the SH representation. Blue is low error, red is high. The minimizer is a point on the line of solutions that minimizes error over the rest of the hemisphere.

The problem of fitting AHD values to SH irradiance functions has been previously considered by Sloan [2008]. Similarly to previous methods, we use the optimal linear SH direction [Sloan et al. 2005] for the highlight direction, so the remaining task is to find the unknown ambient and highlight direction colors, C_a and C_d respectively. In contrast to previous methods, we solve for the colors, one channel at a time, under the constraint that irradiance in the local z-direction should be preserved. This is expressed by the following constraint equation:

$$\begin{bmatrix} 1 & d_z \end{bmatrix} \begin{bmatrix} C_a \\ C_d \end{bmatrix} = I_z. \quad (8)$$

The equation is underdetermined having an infinite number of solutions. See Figure 3 for an example of the solution space. The point on the line of solutions that satisfy the constraint for a particular I_z value closest to the origin is easily computed using the Moore-Penrose pseudoinverse of the constraint equation:

$$\begin{bmatrix} C_a \\ C_d \end{bmatrix} = \frac{I_z}{d_z^2 + 1} \begin{bmatrix} 1 \\ d_z \end{bmatrix}. \quad (9)$$

We parametrize the full solution space by considering the null space of the constraint equation, which is given by the line

$$t \begin{bmatrix} -d_z \\ 1 \end{bmatrix}. \quad (10)$$

Adding the null space solutions to the minimum norm least-squares solution gives us the form of the general solution, parametrized by t (c.f., Figure 3):

$$\frac{I_z}{d_z^2 + 1} \begin{bmatrix} 1 \\ d_z \end{bmatrix} + t \begin{bmatrix} -d_z \\ 1 \end{bmatrix}. \quad (11)$$

Finally, to fit the color values to the target lighting over the hemisphere, we perform optimization in the null space of the constraint using the objective function

$$E(t) = \int_{\Omega^+} \left(I(\omega, t) - L(\omega) \right)^2 d\omega, \quad (12)$$

where $L(\omega)$ is the target light environment and $I(\omega, t)$ is the AHD irradiance, where the color values are parametrized by t and Equation (11). Minimizing the objective function $E(t)$ is a 1D optimization problem that we solve in the spherical harmonic domain. In addition, we enforce non-negativity by restricting the feasible region of the search space to the positive quadrant of the solution space, given by the interval $[t_{\min}, t_{\max}]$ in terms of our parametrization. The final color values are given by plugging the minimizer $t_{\min} \leq t^* \leq t_{\max}$ of the objective function $E(t)$ to Equation (11).

4 CONCLUSIONS AND FUTURE WORK

The AHD lightmap encoding is compact and generates good results compared to other hemispherical parameterizations [Iwanicki and Sloan 2017], but it can suffer from interpolation artifacts, particularly when direct lighting is also baked into the lightmaps. We present an alternate parametrization that makes irradiance in local Z linear, eliminating these artifacts. This parameterization results in less HDR data and is easy to LOD if normal maps are not used in the distance. We also discuss how this invariance can be used to fit AHD from spherical harmonic lighting data, fixing some issues from simple unconstrained least squares [Sloan 2008].

ACKNOWLEDGMENTS

We have had discussed these ideas with various lighters and engineers: Seo Atsushi, Dave Blizzard, Hao Chen, Danny Chan, Shane Daley, Michal Drobot, Adrien Dubouchet, Michal Iwanicki, Stephanus Indra, Dimitar Lazarov and Luka Rommel.

REFERENCES

- Hao Chen. 2008. Lighting and Materials of Halo3. In *Game Developers Conference*.
- JT Hooker. 2016. Volumetric Global Illumination at Treyarch. In *SIGGRAPH 2016 Course: Advances in Real-Time Rendering in 3D Graphics and Games*.
- id Software. 1999. Quake III Arena. <https://github.com/id-Software/Quake-III-Arena>
- Michal Iwanicki. 2013. Lighting Technology of the Last of Us. In *ACM SIGGRAPH 2013 Talks*.
- Michal Iwanicki and Peter-Pike Sloan. 2017. Precomputed lighting in Call of Duty: Infinite Warfare. In *SIGGRAPH 2017 Course: Advances in Real-Time Rendering in Games*.
- Dimitar Lazarov. 2011. Physically-based lighting in Call of Duty: Black Ops. In *SIGGRAPH 2011 Course: Advances in Real-Time Rendering in Games*.
- Gary McTaggart. 2004. Half-Life 2 source shading. In *Game Developers Conference*.
- David Neubelt and Matt Pettineo. 2015. Advanced Lighting R&D at Ready At Dawn Studios. In *SIGGRAPH 2015 Course: Physically Based Shading in Theory and Practice*.
- Peter-Pike Sloan. 2008. Stupid Spherical Harmonics (SH) Tricks. In *Game Developers Conference*.
- Peter-Pike Sloan, Ben Luna, and John Snyder. 2005. Local, Deformable Precomputed Radiance Transfer. *ACM Trans. Graph.* 24, 3 (July 2005), 9.