Accurate Diffuse Lighting from Spherical Gaussian Lights

Yusuke Tokuyoshi
yusuke.tokuyoshi@amd.com
Advanced Micro Devices, Inc.
Japan

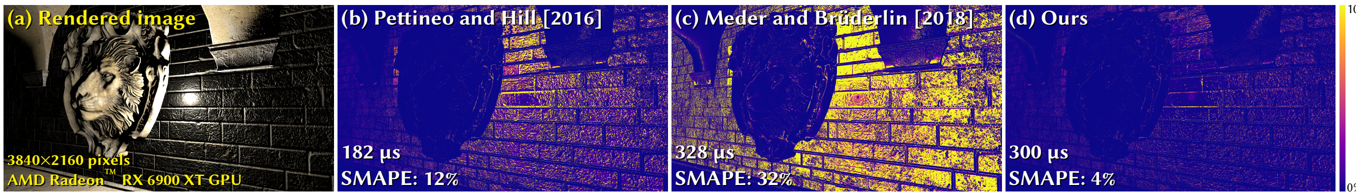

Figure 1: (a) Rendering using our method for a scene lit by a sharp SG light close to a wall. (b–d) Visualizations of symmetry mean absolute percentage errors (SMAPEs). Our approximation (d) produces a smaller error than the previous methods (b–c).

ACM Reference Format:
Yusuke Tokuyoshi. 2022. Accurate Diffuse Lighting from Spherical Gaussian Lights. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters (SIGGRAPH '22 Posters), August 07-11, 2022.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3532719.3543209

1 INTRODUCTION

Spherical Gaussian (SG) lights [Wang et al. 2009] are an efficient approximation for area lights, environment maps, and indirect illumination. It is often used for real-time rendering [Tokuyoshi 2015] and relighting [Zhang et al. 2021]. Although diffuse lighting from an SG light is given by the product integral of the SG and clamped cosine, it does not have a closed-form exact solution. Therefore, efficient approximation is required for real-time applications. Pettineo and Hill [2016] fitted the irradiance from an SG light, but it can produce a significant error for low-frequency SGs. Meder and Brüderlin [2018] introduced an SG approximation for the clamped cosine. Then, since the product of two SGs is an SG, they also introduced an approximation for the hemispherical integral of the SG. However, for a sharp SG light, their approximation error is relatively large especially at grazing angles (Fig. 1). To reduce the error, we present more accurate approximations for the clamped cosine and hemispherical integral than the previous work. Our approximation is simple and easy to implement. By using our method, we are able to improve the quality of real-time SG lighting.

2 APPROXIMATION FOR CLAMPED COSINE

Previous Work. Meder and Brüderlin's SG approximation for clamped cosine is given as $\max(\boldsymbol{\omega}\cdot\mathbf{n},0) \approx (\mu G(\boldsymbol{\omega};\mathbf{n},\lambda)-\alpha)\,H(\boldsymbol{\omega}\cdot\mathbf{n})$, where $\boldsymbol{\omega}\in\mathbb{S}^2$ is a direction of incident light, $\mathbf{n}\in\mathbb{S}^2$ is the surface normal, $G(\boldsymbol{\omega};\mathbf{n},\lambda)=\mathrm{e}^{\lambda(\boldsymbol{\omega}\cdot\mathbf{n})-\lambda}$ is the SG with the sharpness parameter $\lambda\in[0,\infty]$, and $H(\boldsymbol{\omega}\cdot\mathbf{n})$ is the Heaviside function: 1 if $\boldsymbol{\omega}\cdot\mathbf{n}>0$ and 0 if $\boldsymbol{\omega}\cdot\mathbf{n}\le 0$. They obtained constant parameters $\lambda=0.0315$, $\mu=32.708$, and $\alpha=31.703$ using non-linear fitting. However, their parameters are not optimal in practice. In this paper, we derive more near-optimal parameters in an analytic way.

Our Approximation Form. Let $z=\lambda(\boldsymbol{\omega}\cdot\mathbf{n})$, then we can approximate the SG with the following Taylor series:

$$\mathrm{e}^{\lambda}G(\boldsymbol{\omega};\mathbf{n},\lambda)=\mathrm{e}^{z}\approx 1+z+\frac{z^2}{2}+\cdots+\frac{z^m}{m!},$$

where $m\in\mathbb{N}$. The approximation error of this Taylor series converges to zero for $m\to\infty$ or $\lambda\to 0$. By assuming a small λ, we use $m=1$, i.e. first order approximation: $\mathrm{e}^{\lambda}G(\boldsymbol{\omega};\mathbf{n},\lambda)\approx 1+\lambda(\boldsymbol{\omega}\cdot\mathbf{n})$. It yields $\max(\boldsymbol{\omega}\cdot\mathbf{n},0)\approx\frac{1}{\lambda}\left(\mathrm{e}^{\lambda}G(\boldsymbol{\omega};\mathbf{n},\lambda)-1\right)H(\boldsymbol{\omega}\cdot\mathbf{n})$. To improve the accuracy for the spherical integral, we normalize this first order approximation as follows:

$$\max(\boldsymbol{\omega}\cdot\mathbf{n},0)\approx\frac{\left(\mathrm{e}^{\lambda}G(\boldsymbol{\omega};\mathbf{n},\lambda)-1\right)H(\boldsymbol{\omega}\cdot\mathbf{n})\int_{\mathbb{S}^2}\max(\boldsymbol{\omega}\cdot\mathbf{n},0)\mathrm{d}\boldsymbol{\omega}}{\int_{\mathbb{S}^2}\left(\mathrm{e}^{\lambda}G(\boldsymbol{\omega};\mathbf{n},\lambda)-1\right)H(\boldsymbol{\omega}\cdot\mathbf{n})\mathrm{d}\boldsymbol{\omega}}$$

$$=\boxed{\alpha\left(\mathrm{e}^{\lambda}G(\boldsymbol{\omega};\mathbf{n},\lambda)-1\right)H(\boldsymbol{\omega}\cdot\mathbf{n})},$$

where $\alpha=\lambda/(2\mathrm{e}^{\lambda}-2-2\lambda)$. Our approximation error converges to zero for $\lambda\to 0$. However, due to the floating-point representation, a smaller λ introduces a larger numerical error. Therefore, a near-optimal λ is obtained by minimizing the sum of the approximation error and numerical error.

Minimization of the Error Bound. For lighting computation, the numerical error occurs when calculating the product integral of an SG light and clamped cosine. Let $\mathbf{v}\in\mathbb{S}^2$ and $\kappa\in[0,\infty]$ be the axis and sharpness of the SG light, then the product integral of the normalized SG (a.k.a von Mises–Fisher distribution) and clamped cosine is approximated as follows:

$$L(\mathbf{v},\kappa)=\frac{\int_{\mathbb{S}^2}G(\boldsymbol{\omega};\mathbf{v},\kappa)\max(\boldsymbol{\omega}\cdot\mathbf{n},0)\mathrm{d}\boldsymbol{\omega}}{\int_{\mathbb{S}^2}G(\boldsymbol{\omega};\mathbf{v},\kappa)\mathrm{d}\boldsymbol{\omega}}\approx\alpha p-\alpha q,$$

where $p=\mathrm{e}^{\lambda}\frac{\int_{\mathbb{S}^2}G(\boldsymbol{\omega};\mathbf{v},\kappa)G(\boldsymbol{\omega};\mathbf{n},\lambda)H(\boldsymbol{\omega},\mathbf{n})\mathrm{d}\boldsymbol{\omega}}{\int_{\mathbb{S}^2}G(\boldsymbol{\omega};\mathbf{v},\kappa)\mathrm{d}\boldsymbol{\omega}}$ and $q=\frac{\int_{\mathbb{S}^2}G(\boldsymbol{\omega};\mathbf{v},\kappa)H(\boldsymbol{\omega},\mathbf{n})\mathrm{d}\boldsymbol{\omega}}{\int_{\mathbb{S}^2}G(\boldsymbol{\omega};\mathbf{v},\kappa)\mathrm{d}\boldsymbol{\omega}}$ are obtained using our hemispherical integral approximation (Sect. 3).

SIGGRAPH '22 Posters, August 07-11, 2022, Vancouver, BC, Canada
© 2022 Copyright held by the owner/author(s).
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 *Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters (SIGGRAPH '22 Posters), August 07-11, 2022*, https://doi.org/10.1145/3532719.3543209.

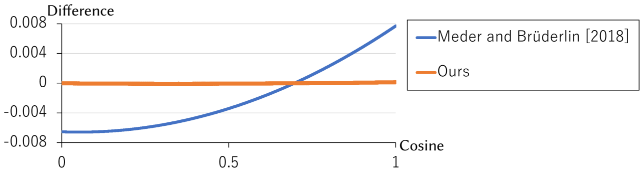

Figure 2: Plots of differences between SG approximations and exact clamped cosine. Our approximation has a significantly smaller error than the previous approximation.

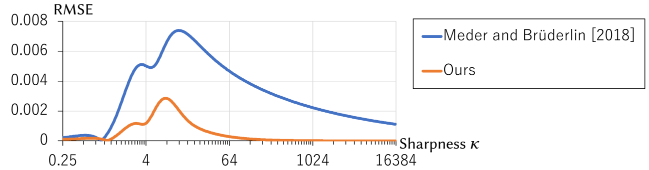

Figure 3: Root mean square errors (RMSEs) for normalized integral approximations. Our method (Eq. 1) has a smaller error than the previous method for large sharpness.

Figure 4: Dynamic indirect illumination (3840×2160 pixels) using virtual SG lights. Although Pettineo and Hill's method (P&H) [2016] is faster than Meder and Brüderlin (M&B) [2018] and our method, it can produce undesirable black splotches. Our method does not produce such artifacts, and it is more accurate and slightly faster than M&B [2018].

Although we can compute the difference of products $\alpha p - \alpha q$ precisely [Kahan 2004], it still produces a numerical error due to the rounding of α, p, and q. For floating point values, the maximum rounding error for α is given as $0.5\epsilon|\alpha|$ where ϵ is the machine epsilon. Therefore, our approximation with the maximum rounding error is written as $L(\mathbf{v}, \kappa) \approx (\alpha \pm 0.5\epsilon\alpha)\,(p - q \pm 0.5\epsilon(p + q))$. Thus, the upper bound of the sum of approximation error and the numerical error is $E_{\mathbf{v},\kappa}(\lambda) = |(\alpha \pm 0.5\epsilon\alpha)\,(p - q \pm 0.5\epsilon(p+q)) - L(\mathbf{v}, \kappa)|$. When $\kappa = \infty$, we obtain the numerical error of the clamped cosine: $L(\mathbf{v}, \infty) = \max(\mathbf{v}\cdot\mathbf{n}, 0)$. For this case, since $p = \mathrm{e}^{\lambda} G(\mathbf{v}; \mathbf{n}, \lambda) H(\mathbf{v}\cdot\mathbf{n})$ and $q = H(\mathbf{v}\cdot\mathbf{n})$, the error is largest at $\mathbf{v} = \mathbf{n}$ as follows:

$$E_{\mathbf{n},\infty}(\lambda) = \frac{(1 + 0.5\epsilon)\,\lambda\left((1 + 0.5\epsilon)\,\mathrm{e}^{\lambda} - 1 + 0.5\epsilon\right)}{2\mathrm{e}^{\lambda} - 2 - \lambda} - 1.$$

Minimizing this error bound for single precision (i.e. $\epsilon = 2^{-23}$), we obtain the near-optimal parameter $\boxed{\lambda = 0.00084560872241480124}$. Fig. 2 shows the error of our approximation and Meder and Brüderlin's approximation for clamped cosine. For closeups of our error, please see the supplemental material.

3 APPROXIMATION FOR HEMISPHERICAL INTEGRAL

The hemispherical integral of an SG is represented with an interpolation between the upper hemispherical integral $A(\kappa)$ and lower hemispherical integral $B(\kappa)$ [Meder and Brüderlin 2018] as follows:

$$\int_{S^2} G(\boldsymbol{\omega}; \mathbf{v}, \kappa) H(\boldsymbol{\omega}\cdot\mathbf{n}) \mathrm{d}\boldsymbol{\omega} = s(\mathbf{v}\cdot\mathbf{n}, \kappa) A(\kappa) + (1 - s(\mathbf{v}\cdot\mathbf{n}, \kappa)) B(\kappa),$$

where $A(\kappa) = 2\pi(1 - \mathrm{e}^{-\kappa})/\kappa$, $B(\kappa) = 2\pi\mathrm{e}^{-\kappa}(1 - \mathrm{e}^{-\kappa})/\kappa$, and $s(\mathbf{v}\cdot\mathbf{n}, \kappa) \in [0, 1]$ is the normalized integral. Since this normalized integral is sigmoid-form, the previous work roughly approximated it using a normalized logistic function similar to Wang et al. [2009]'s visibility approximation. Unlike such a logistic approximation, we use the cumulative distribution function (CDF) of a Gaussian on a plane. This is because an SG asymptotically approaches to a planar Gaussian for $\kappa \rightarrow \infty$. By normalizing this CDF for $\mathbf{v}\cdot\mathbf{n} \in [-1, 1]$, we yield

$$\boxed{s(\mathbf{v}\cdot\mathbf{n}, \kappa) \approx \frac{1}{2} + \frac{\mathrm{erf}\,(t(\kappa)(\mathbf{v}\cdot\mathbf{n}))}{2\mathrm{erf}\,(t(\kappa))}}, \qquad (1)$$

where $t(\kappa)$ is the steepness of the CDF, and it is given by

$$t(\kappa) \approx \kappa\sqrt{\frac{0.5\kappa + 0.65173288269070562}{\kappa^2 + 1.3418280033141288\kappa + 7.2216687798956709}}.$$

We obtained this approximated $t(\kappa)$ by fitting it to numerically computed optimal steepness (please see the supplemental material).

Thanks to this fitted steepness, our approximation error is smaller than the previous logistic approximation in most cases even for small κ (see Fig. 3).

4 RESULTS

Here we show images rendered on an AMD Radeon™ RX 6900 XT GPU in a numerically stable manner (please refer to the supplemental material for implementation details). Fig. 1 shows surfaces lit by a sharp SG light at grazing angles. For this case, our method produces significantly less error than the previous methods. Fig. 4 shows dynamic indirect illumination roughly approximated by two virtual SG lights [Tokuyoshi 2015]. For such low-frequency SGs, Pettineo and Hill's method (P&H) [2016] can produce undesirable artifacts. On the other hand, Meder and Brüderlin's method (M&B) [2018] and our method do not produce such artifacts. Although our method is more expensive than P&H, it robustly performs higher-quality lighting with less computational cost than M&B.

REFERENCES

William Kahan. 2004. On the Cost of Floating-Point Computation Without Extra-Precise Arithmetic. https://people.eecs.berkeley.edu/~wkahan/Qdrtcs.pdf

Julian Meder and Beat Brüderlin. 2018. Hemispherical Gaussians for Accurate Light Integration. In *ICCVG' 18*. 3–15. https://doi.org/10.1007/978-3-030-00692-1_1

Matt Pettineo and Stephen Hill. 2016. SG Series Part 3: Diffuse Lighting From an SG Light Source. https://therealmjp.github.io/posts/sg-series-part-3-diffuse-lighting-from-an-sg-light-source/

Yusuke Tokuyoshi. 2015. Virtual Spherical Gaussian Lights for Real-time Glossy Indirect Illumination. *Comput. Graph. Forum* 34, 7 (2015), 89–98.

Jiaping Wang, Peiran Ren, Minmin Gong, John Snyder, and Baining Guo. 2009. All-Frequency Rendering of Dynamic, Spatially-Varying Reflectance. *ACM Trans. Graph.* 28, 5 (2009), 133:1–133:10. https://doi.org/10.1145/1618452.1618479

Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. 2021. PhySG: Inverse Rendering with Spherical Gaussians for Physics-based Material Editing and Relighting. In *CVPR '20*.

©2022 Advanced Micro Devices, Inc. All rights reserved. AMD, Radeon and combinations thereof are trademarks of Advanced Micro Devices, Inc.
Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.