Directional Occlusion via Multi-Irradiance Mapping

Martin Misiak
TH Köln
martin.misiak@smail.th-koeln.de

Arnulph Fuhrmann
TH Köln
arnulph.fuhrmann@th-koeln.de

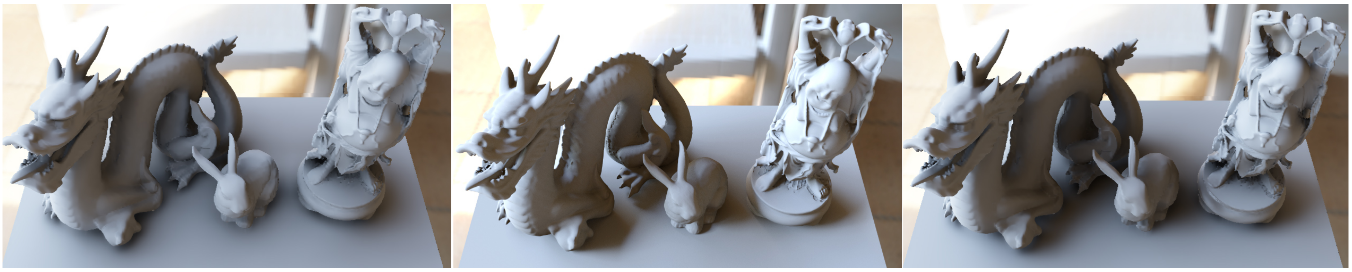

Figure 1: Left: Irradiance map with bent normal + AO (8.0 ms). **Middle:** Raytraced ground truth. **Right:** Our approach (7.2 ms).

ABSTRACT

We present a new, physically plausible, real-time approach to compute directional occlusion for dynamic objects, lit with image based lighting. For this, we partition the hemisphere into multiple sectors and pre-convolve these into separate irradiance maps. At runtime the contributions of each sector are then individually occluded and summed together.

CCS CONCEPTS

• **Computing methodologies → Rendering**;

KEYWORDS

directional occlusion, real-time rendering, irradiance mapping

ACM Reference format:
Martin Misiak and Arnulph Fuhrmann. 2017. Directional Occlusion via Multi-Irradiance Mapping. In *Proceedings of SIGGRAPH '17 Posters, Los Angeles, CA, USA, July 30 - August 03, 2017,* 2 pages.
DOI: 10.1145/3102163.3102217

1 INTRODUCTION

The shading of objects under image based lighting in real-time is very challenging, as it requires to solve the triple product integral of incident radiance, material BRDF and visibility function for each shaded point. The diffuse material contribution is often handled separately, as its BRDF is view-independent, and radiance contributions from the whole hemisphere need to be considered. A widespread approach is to omit the visibility term and pre-convolve, for each surface normal, the arriving irradiance into a look up texture, the irradiance map [Miller and Hoffman 1984].

To make up for the missing occlusion, this technique is often used in conjunction with ambient occlusion (AO), which is defined as the integral of the visibility function over the hemisphere and is solely dependent on the scene's geometric information. While it strengthens the geometric perception of objects, the decoupling of visibility from the incident radiance makes AO produce the most plausible results only in homogeneous environments.

Landis [2002] proposed to compute the average unoccluded direction (the bent normal) alongside the AO factor and use it, instead of the surface normal, as a lookup into the irradiance map. Klehm et al. [2011] implemented a real-time version of this technique using a screen space ambient occlusion (SSAO) approach and extended it via a bent cone, whose aperture represents the variance of the computed bent normal. Based on this aperture, a lookup into a correspondingly pre-filtered environment map is made, to reduce the contribution of occluded directions. With increasing difference between normal and bent cone however, the error in incident radiance increases. Additionally, having only a single unoccluded direction, the bent normal approach requires blockers to form a simple horizon.

Ritschel et al. [2009] also extended SSAO with directional occlusion, by evaluating the triple product integral for N discrete directions. While this evaluation allows for multiple directions to contribute to the incident radiance, the point sampling scheme only works for very low frequency environments.

Crassin et al. [2011] introduced voxel cone tracing, which allows to compute smooth, object space AO with only a few samples.

Herholz et al. [2013] achieve directional occlusion by evaluating the triple product integral in the spherical harmonics (SH) domain. Due to the use of SH, the results do not scale well to higher frequencies and are prone to ringing artifacts when used with high intensity environment maps.

2 OUR APPROACH

Our approach to directional occlusion builds upon the idea of pre-convolving the irradiance map. However instead of filtering the entire hemisphere to a single irradiance lookup for normal n, we

---

SIGGRAPH '17 Posters, Los Angeles, CA, USA
978-1-4503-5015-0/17/07…$15.00
DOI: 10.1145/3102163.3102217

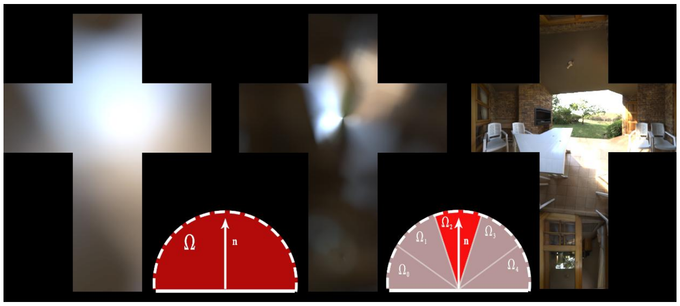

Figure 2: Left: Irradiance map. **Middle:** A Multi-Irradiance map from our approach corresponding to one hemispherical sector. **Right:** Radiance map.

divide the computations into N equally sized sectors.

$$E(n) = \sum_{i}^{N} E_i(n) = \sum_{i}^{N} \int_{\Omega_i} L_e(l_i)\langle n \cdot l_i \rangle dl_i \tag{1}$$

$$E(n) \approx \sum_{i}^{N} V_i(\omega_i) E_i(n) \tag{2}$$

Each sector Ω_i is pre-convolved and stored into a separate texture E_i (Figure 2 middle). At runtime, the irradiance values of each sector are looked up using the surface normal n. For unoccluded surfaces, the sum of these values corresponds to a lookup in the classical irradiance map. However, due to the partitioning of the hemisphere, we are able to occlude each contribution individually, resulting in a physically plausible directional occlusion of the incident irradiance (Equation 2). For this, we use voxel cone tracing, as it allows us to efficiently approximate the occlusion within each sector.

Our approach can suffer from shading transition artifacts when paired with very high frequency environment maps. To counteract this, two choices are available. Either the input radiance map is prefiltered in order to reduce higher frequencies, or the problematic light sources can be handled separately. We choose the later and implemented an area light extraction algorithm as described by Annen et al. [2008], which removes the problematic light sources before the convolutions are performed. Their contribution is then added explicitly at runtime. To ensure proper occlusion, a cone trace is performed towards each light. This enables us to also capture high frequency shadows (Figure 3 right).

3 RESULTS

The timings were done on a system with the following specifications: Intel i5 2500, Geforce 980Ti, 8 GB Ram. The voxelization of the scene was done at a resolution of 256^3 into a 3D texture and 15 cones were used to approximate the occlusion inside the hemisphere. Since we do not target very high frequency occlusion, we determine the incident irradiance at 1/4 resolution and upsample it.

The test scene (Figure 1) consists of 267K polygons and is rendered at 1920x1080. We benchmarked three setups. The first uses a standard irradiance map with occlusion provided through AO (7.0 ms). The second computes a bent normal alongside the AO factor, which is used as a lookup into the irradiance map (8.0 ms). And lastly our approach (7.2 ms). Although we perform N-1 more texture reads, this overhead is quite small when compared to the cost of tracing cones. Interestingly our solution is faster than the bent normal approach and we speculate that this difference is due to an suboptimal texture caching behavior, as the bent normal does not change in a coherent pattern.

Figure 3: Scene lit by our approach under different illuminations.

4 CONCLUSION

We presented a novel, physically plausible approach to incorporate directional occlusion into scenes lit with image based lighting. It supports fully dynamic objects and through the number of hemispherical segments an easy quality/performance adjustment can be made, as the approach can converge to ground truth. We also implemented bent normal irradiance mapping using cone traced AO and compared it to our solution. Multi-Irradiance mapping requires N times more storage, but has a negligible performance overhead, while providing superior results.

5 ACKNOWLEDGMENTS

We would like to thank Greg Zaal from hdrihaven.com for providing environment maps, as well as Stanford University for their 3D models. Additionally, we thank Philipp Schoemacker and Tobias Bayer for insightful discussions. This work was partially funded by the German Federal Ministry for Economic Affairs and Energy under the grant number KF2793703SS3 as part of the ZIM program.

REFERENCES

Thomas Annen, Zhao Dong, Tom Mertens, Philippe Bekaert, Hans-Peter Seidel, and Jan Kautz. 2008. Real-time, All-frequency Shadows in Dynamic Scenes. *ACM Trans. Graph.* 27, 3, Article 34 (Aug. 2008), 8 pages. `DOI`:http://dx.doi.org/10.1145/1360612.1360633

Cyril Crassin, Fabrice Neyret, Miguel Sainz, Simon Green, and Elmar Eisemann. 2011. Interactive Indirect Illumination Using Voxel-based Cone Tracing: An Insight. In *ACM SIGGRAPH 2011 Talks (SIGGRAPH '11)*. ACM, New York, NY, USA, Article 20, 1 pages. `DOI`:http://dx.doi.org/10.1145/2037826.2037853

Sebastian Herholz, Jens-Uwe Hahn, and Andreas Schilling. 2013. Dual space directional occlusion. *The Visual Computer* 29, 9 (2013), 917–926. `DOI`:http://dx.doi.org/10.1007/s00371-013-0856-7

Oliver Klehm, Tobias Ritschel, Elmar Eisemann, and Hans-Peter Seidel. 2011. Bent Normals and Cones in Screen-space. In *Vision, Modeling, and Visualization (2011)*, Peter Eisert, Joachim Hornegger, and Konrad Polthier (Eds.). The Eurographics Association. `DOI`:http://dx.doi.org/10.2312/PE/VMV/VMV11/177-182

Hayden Landis. 2002. RenderMan in Production. In ACM SIGGRAPH 2002 Course 16. (2002).

Gene S. Miller and C. Robert Hoffman. 1984. Illumination and Reflection Maps: Simulated Objects In Simulated and Real Environments. SIGGRAPH 84 Advanced Computer Graphics Animation seminar notes. (1984).

Tobias Ritschel, Thorsten Grosch, and Hans-Peter Seidel. 2009. Approximating Dynamic Global Illumination in Image Space. In *Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games (I3D '09)*. ACM, New York, NY, USA, 75–82. `DOI`:http://dx.doi.org/10.1145/1507149.1507161