

Real-time Soft Shadow Mapping using the DDOC

Qi Mo*
University of Iowa

Chris Wyman†
University of Iowa




Figure 1: Scenes rendered using our soft shadow mapping algorithm

Soft shadows provide important information to viewers about spatial relationships between objects. Generating accurate soft shadows in real time has been a focus of research in computer graphics. Shadow mapping is a standard way of generating hard shadows in real time and several methods including Smoothies [Chan and Durand 2003] and penumbra maps [Wyman and Hansen 2003] extend the standard shadow map algorithm to handle soft shadows. A major problem with these methods is that they use the hard shadow regions given by the standard shadow map as an approximation of the umbra regions; therefore, they only capture the outer penumbras, resulting in shadows with larger umbrae that appear darker than accurate soft shadows. We propose a real time soft shadow algorithm that extends the standard shadow map algorithm but captures both inner and outer penumbra, leading to more accurate soft shadows.

Our method takes advantage of the depth discontinuity occlusion camera model (DDOC) proposed in [Popescu and Aliaga 2006]. DDOC is a non-pinhole camera model with bent light rays that distort the scene at depth discontinuities and sample normally hidden surfaces. Although pinhole camera models are dominant in computer graphics, such cameras do not give information about the inner penumbra regions because they are hidden in the shadow map. On the contrary, the DDOC is able to capture samples that are hidden (thus will be missing in a standard shadow map) but are close to the depth discontinuities (thus are partially visible from an area light source), without adding more layers to the shadow map. In other words, the DDOC can effectively sample more than one layers of depth information in the scene while maintaining the simplicity and efficiency of a single layer shadow map.

The original DDOC algorithm builds a distortion map that provides fine-grain control for the distortion at each pixel in the scene. We modified their algorithm for soft shadow generation: our method builds a similar map that stores the potential penumbra regions and the corresponding light intensities within such regions. The main steps of our algorithms are as follows:

1. Generate a standard shadow map from the light's view
2. Compute 1-bit depth discontinuity map from the shadow map
3. For each non zero pixel in the depth discontinuity map, splat it in the distortion map, storing the light intensities of pixels

within the splat and the near and far z values that create the depth discontinuity

4. Render the scene from the viewer's point of view, using a shader to calculate soft shadows by looking into both the shadow map and the distortion map

Figure [1] gives images that are generated by our prototype soft shadow map algorithm. Currently we achieve framerates between 30 and 60 fps on a NVIDIA GeForce 7800 graphics card. The algorithm captures both the inner and outer penumbra, so when the size of the light source increases, the umbra shrinks as it should be. There are still several artifacts that we must eliminate from our method: the first occurs when more than two layers of depth are present, pixel in the soft shadow map can be covered by two conflicting splats. Since we have not implemented the DDOC's splat adjustment step, we choose the splat with center closest to the pixel for now, which causes artifacts where points occluded by multiple layers of occluders can get shaded by the wrong splat. The second artifact comes from splatting individual pixels in the depth discontinuity map; the direction of depth discontinuity can change abruptly and result in non-smooth light intensity between neighboring splats. We are considering eliminating the second type of artifact by extracting the silhouette edges in object space instead of using image space edge detection.

Acknowledgements We would like to thank Professor Voicu Popescu for his helpful discussions with us on the DDOC.

References

- CHAN, E., AND DURAND, F. 2003. Rendering fake soft shadows with smoothies. In *Proceedings of the Eurographics Symposium on Rendering*, 208–218.
- POPESCU, V., AND ALIAGA, D. 2006. The depth discontinuity occlusion camera. In *Proceedings of the symposium on Interactive 3D graphics and games*, 139–143.
- WYMAN, C., AND HANSEN, C. 2003. Penumbra maps: Approximate soft shadows in real-time. In *Proceedings of the Eurographics Symposium on Rendering*, 202–207.

*e-mail: qmo@cs.uiowa.edu

†e-mail: cwyman@cs.uiowa.edu