Interactive Backprojected Soft Shadows with an Occlusion Camera Shadow Map (sap_0295)

Qi Mo*
University of Iowa

Chris Wyman†
University of Iowa

1 Introduction

Soft shadows are crucial for computer generated imagery, as they increase the level of realism and improve spacial perception. Generating soft shadows is inherently difficult because it involves evaluating the visibility between an area light source and every point in the scene. One way to accurately compute this visibility is to backproject occluder geometry onto the light for every point in the scene. However, searching for relevant edges is costly so backprojection scales poorly with increased scene complexity. We present an interactive soft shadow algorithm that draws insight from the depth discontinuity occlusion camera model (DDOC) [Popescu and Aliaga 2006], a non-pinhole camera that bends light rays near discontinuity edges to sample normally hidden surfaces. By carefully choosing the size of DDOC distortion regions based on light and occluder geometry, we generate a map that stores the closest silhouette edges for each pixel. This enables us to find the most relevant edge to backproject using a single map lookup and additional edges with a small neighborhood search. Our method combines the accuracy of object-space backprojection methods with the efficiency of an image-based approach.

2 Backprojection Shadows using a DDOC

Previously, we showed the applicability of non-pinhole cameras to the interactive rendering of soft shadows [Mo et al. 2007]. We observed that the basic problem for rendering soft shadows, computing light visibility at partially-occluded surfaces, is similar to the image-based rendering problem of rendering a scene from novel viewpoints based on a single reference image. In particular, the DDOC proves useful for this application, as it aims to capture all visible and "nearly-visible" geometry from a viewpoint by bending light rays around silhouettes. By rendering a shadow map with a DDOC instead of a pinhole camera, we can thus store a single image containing fully illuminated geometry as well as partially occluded surfaces (i.e., in the penumbra).

In this work, we observe that silhouette edge information is vital both to creation of the DDOC as well as soft shadow techniques based upon backprojection [Drettakis and Fiume 1994]. Instead of using the DDOC to explicitly distort geometry (as in Mo et al.), here we simply use it to find relevant silhouette edges. We then backproject these edges onto the light source to compute shadow intensity. As identifying the relevant silhouette edges for a given pixel is the most costly component of backprojection, this technique allows backprojected soft shadows at interactive rates. This differs fundamentally from our previous technique, which iteratively applied a distortion (determined by the DDOC shadow map) for each pixel to approximate its shadow intensity.

This new approach has numerous advantages: because we use explicit edge information instead of discretized visibility, most shadow map aliasing is avoided; our method requires less memory that the original occlusion camera soft shadows; should we conservatively identify too many silhouettes, backprojection automatically handles extras without artifacts; and it allows backprojection within an image-space framework, without introducing the cracking artifacts of Guennebaud et al. [2006].

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

Figure 1: *Examples of soft shadows rendered interactively with our backprojection technique.*

The main limitation is that a single shadow map texel may only store a singular silhouette edge, which leads to difficulty identifying all edges for complexly shadowed regions. Currently, we perform an image-space search to find other overlapping edges. Even so, this search is more efficient than with a standard shadow map (e.g., Guennebaud et al.), because edge information is spread among multiple texels instead of remaining concentrated at the edges.

3 Results and Conclusions

Figure 1 demonstrates our results. The accompanying video shows the interactivity and realistic soft shadows attainable even with dynamic lighting and geometry. This work eliminates many of the limitations of our previous work, as we store explicit edge equations for backprojecting silhouettes instead of relying on discretized distortion parameters to compute shadow intensity. This allows larger light sources, while maintaining the efficiency of an image-based approach. Future plans include addressing remaining artifacts where multiple edges project to a single texel and demonstrating results with arbitrarily shaped and textured lights.

References

DRETTAKIS, G., AND FIUME, E. 1994. A fast shadow algorithm for area light sources using backprojection. In *Proceedings of SIGGRAPH*, 223–230.

GUENNEBAUD, G., BARTHE, L., AND PAULIN, M. 2006. Real-time soft shadow mapping by backprojection. In *Eurographics Symposium on Rendering*, 227–234.

MO, Q., POPESCU, V., AND WYMAN, C. 2007. The soft shadow occlusion camera. Tech. Rep. UICS-07-02, University of Iowa. *(Submitted to Pacific Graphics)*.

POPESCU, V., AND ALIAGA, D. 2006. The depth discontinuity occlusion camera. In *Proc. Symp. Interactive 3D Graphics and Games*, 139–143.