Imperfect Voxelized Shadow Volumes

Chris Wyman*
NVIDIA

Zeng Dai†
University of Iowa

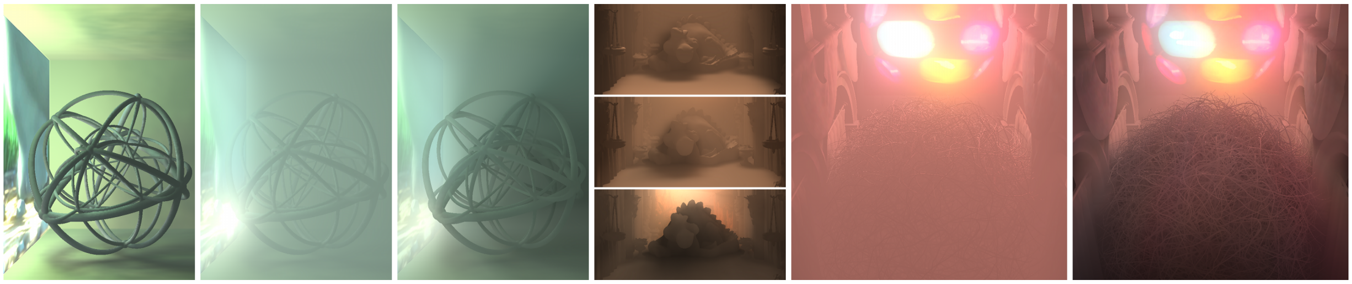

Figure 1: *(Left) A video-illuminated polytope with no media, unshadowed homogeneous media, shadowed media. (Center) A dragon in the Crytek Sponza with a moving light. (Right) A hairball inside the Sponza, with and without volumetric shadows. All run interactively.*

Voxelized shadow volumes (VSVs) [Wyman 2011] are a discretized view-dependent shadow volume representation, but are limited to point or directional lights. We extend them, allowing dynamic volumetric visibility from area lights using *imperfect shadow volumes*. As with imperfect shadow maps [Ritschel et al. 2008], area lights can use coarser spacial sampling without significantly degrading quality. Combining coarser resolution with a parallel shadow volume construction enables interactive rendering of dynamic volumetric shadows from area lights in homogeneous single-scattering media, at around 4x the cost of hard volumetric shadows.

Keywords: shadows, area lights, participating media, voxelization

1 Introduction

Voxelized shadow volumes rely on a spacial discretization based on an angular, epipolar sampling. This sampling's key advantage is a structure where grid axes lie parallel to view and light rays, allowing cache-aligned visibility lookups along camera rays and a parallel scan along light rays to quickly compute dense shadow samples. A VSV retains volumetric visibility inherent in geometric shadow volumes, without many of the costs (exorbitant fill rate consumption and need to identify silhouette edges).

Unfortunately epipolar sampling, by design, depends heavily on the epipole connecting the camera and point light, complicating extensions to area lighting. We take a common approach, sampling an area light as numerous point lights (e.g., Heckbert and Herf [1997]). But this scales linearly with light samples; sampling dense enough for good quality can require over a second per frame.

But computing a naive VSV for each light sample is overkill. We explored the VSV resolution needed when sampling VPLs from an area source. Given averaging from accumulating multiple lights, we found sampling visibility 64x coarser than needed for a single point source feasible (e.g., 128^3 instead of 512^3 or larger).

2 Imperfect Shadow Volumes

In the context of shadow maps, Ritschel et al. [2008] observed not only coarse visibility, but *partially incorrect* visibility could generate realistic shadows from area lights. Given human perception is tuned more to detecting surfaces rather than thin media, we hypothesized artifacts in volumetric visibility would be still less visible.

To create our ”imperfect VSVs,” we used imperfect shadow maps to populate epipolar space with occluders (separately for each sample on the area light) and performed a scan to extrude occlusions away from the light (to form shadow volumes). As these imperfect volumes are low resolution, we can fit hundreds of them in a single render target. This allows a single parallel scan over one buffer, greatly reducing GPU overhead compared to hundreds of scans.

Interpolation. To reduce aliasing artifacts arising from course samples, we introduce an interpolation scheme in epipolar space. Expanding the Riemann sum used to integrate scattering, terms cancel, allowing us to avoid costly trilinear lookups at each step. Instead we need to interpolate at shadow boundaries, and due to VSV view dependency this can occur in screen space. We found a single bilinear interpolation suffices (for each VPL), and due to our epipolar representation this requires exactly two texture lookups.

Results. We implemented imperfect VSVs in OpenGL, and compare with a naive, brute force application of VSVs on a per-VPL basis. The table below compares performance (from Figure 1, left) when naively applying VSVs for 256 VPLs with both the naive method and our new imperfect VSV algorithm at a coarser, more appropriate sampling density for area lights.

	Voxelize	Scan to VSV	Final Gather
Naive, at 512^3 [Wyman 2011]	250 ms	1400 ms	175 ms
Naive, at 128^3	128 ms	124 ms	99 ms
Imperfect VSVs, at 128^3	7 ms	5 ms	15 ms

References

HECKBERT, P., AND HERF, M. 1997. Simulating soft shadows with graphics hardware. Tech. Rep. CMU-CS-97-104, CMU.

RITSCHEL, T., GROSCH, T., KIM, M., SEIDEL, H.-P., DACHSBACHER, C., AND KAUTZ, J. 2008. Imperfect shadow maps for efficient computation of indirect illumination. *ACM Transactions on Graphics 27*, 5, 1–8.

WYMAN, C. 2011. Voxelized shadow volumes. In *High Performance Graphics*, 33–40.

*e-mail:chris.wyman@acm.org
†e-mail:zeng-dai@uiowa.edu

SIGGRAPH 2013, July 21 – 25, 2013, Anaheim, California.

ACM 978-1-4503-2261-4/13/07