Ambient Occlusion Volumes

Morgan McGuire
Williams College and NVIDIA

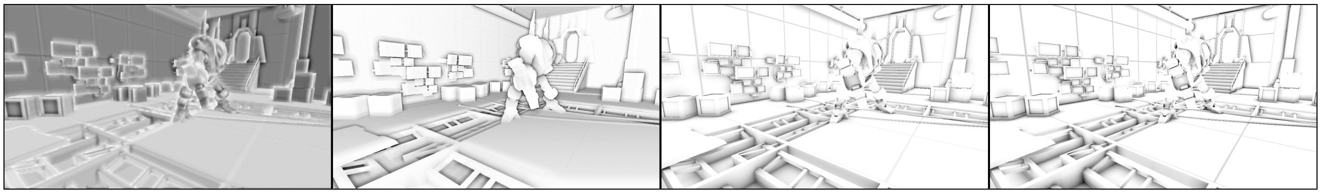

[Kajalin 2009]	[Szirmay-Kalos et al. 2009]	Ray Traced Reference	**Ambient Occlusion Volumes**
256 samples/pix	64 samples/pix	5000 samples/pix	(analytic)
16 ms	61 ms	556228 ms	31 ms

Figure 1: 1.4 M-triangle scene at 720p resolution on GeForce 280 GT using geometry from *Marvel Ultimate Alliance 2* by Vicarious Visions/Activision. The new Ambient Occlusion Volume algorithm produces quality comparable to ray tracing, in real time and with no noise. In contrast, many alternative screen-space algorithms offer extremely fast AO approximations but cannot converge towards the correct result, even when extended to high sampling rates as shown here.

Introduction Ambient illumination is an approximation to the light reflected from the sky and other objects in a scene. **Ambient occlusion** (AO) is the darkening that occurs when this illumination is locally blocked by another object or a nearby fold in a surface. Both ambient illumination and occlusion are qualitatively important to perception of shape and depth in the real world. Artists have long recognized AO, and specifically seek to reproduce the real phenomena such as corner darkening and contact shadows. Because AO is so important to shape perception, practitioners in many disciplines rely software that simulates AO to visualize structures. Ambient-occlusion renderings of building facades, medical data, and engineering parts reveal the spatial relationship between complex details. CG artists also a prefer untextured, AO-only rendering as a method for visualizing geometry.

While the problem of computing an explicit AO factor has been studied for over a decade, most previous AO algorithms seek extreme quality or extreme performance. Ray tracing methods are the gold standard for quality, but take minutes to evaluate. Recent screen space methods (e.g., [Mittring 2007; Kajalin 2009; Bavoil and Sainz 2009; Szirmay-Kalos et al. 2009]) are well suited to console games because they execute in only a few milliseconds at HD resolution. However, they sacrifice significant quality for that performance and do not improve under the higher sample rates possible on today's GPUs. We build on the screen space notion of AO as a deferred shading pass for performance, but develop and analytic solution inspired by radiosity and Crow's shadow volumes and thus yields higher quality. By enabling high-quality AO in real-time, it is well suited to modeling applications and CAD. With screen-space subsampling it can match the performance of screen space methods with high quality for games on contemporary GPUs.

Algorithm For each polygon $P = \{\vec{p}_0, ..., \vec{p}_{k-1}\}$ in the scene, the a geometry shader forms a corresponding **ambient occlusion volume** V. This volume is the analogue of a shadow volume for ambient light: it is a polygonal prism containing P dilated by the maximum perceivable occlusion distance δ and then extruded along its face normal by δ. The algorithm rasterizes each V over a previously-rendered z-buffer with depth write disabled. For each pixel covered, a pixel shader reconstructs the corresponding world-space position $\vec{x}$ (and normal $\vec{n}$ from a geometry buffer), clips P to the local tangent plane defined by $\vec{x}$ and $\vec{n}$, and then evaluates:

$$\mathrm{AO}_P = \frac{1}{\pi} \sum_{i=0}^{k-1} \mathscr{A}\left(\mathrm{S}(\vec{p}_i - \vec{x}), \mathrm{S}(\vec{p}_{(i+1) \bmod k} - \vec{x})\right), \quad (1)$$

where $\mathscr{A}(\hat{a}, \hat{b})$ is the signed area of the projection of the disk segment $(\vec{0}, \hat{a}, \hat{b})$ into an elliptic segment on the tangent plane at $\vec{x}$. This is the cosine-weighted solid angle subtended by P, which is also known as the form factor.

These AO_P values accumulate into a draw buffer. When all P have been evaluated, that buffer contains the net ambient occlusion at each pixel. For scenes where many thin objects are stacked, the volumes can interpenetrate and produce darker results than ray tracing. Where objects have thickness at least $\delta/2$, the analytic result is exact.

An early version of this work appeared in Williams College CS technical note 2009-01 and a more recent version has been submitted to ACM ToG [McGuire 2010].

References

BAVOIL, L., AND SAINZ, M. 2009. Multi-layer dual-resolution screen-space ambient occlusion. In *ShaderX*^{7}, W. Engel, Ed. Charles River Media, March.

KAJALIN, V. 2009. Screen space ambient occlusion. In *in ShaderX*^{7}, W. Engel, Ed. Charles River Media, March.

MCGUIRE, M., 2010. Ambient occlusion volumes. submitted to *ACM Transactions on Graphics*.

MITTRING, M. 2007. Finding next gen: Cryengine 2. In *SIGGRAPH '07: ACM SIGGRAPH 2007 courses*, ACM, New York, NY, USA, 97–121.

SZIRMAY-KALOS, L., UMENHOFFER, T., TTH, B., SZCSI, L., AND CASASAYAS, M. 2009. Volumetric ambient occlusion. *IEEE Computer Graphics and Applications*. Preprint—to appear.

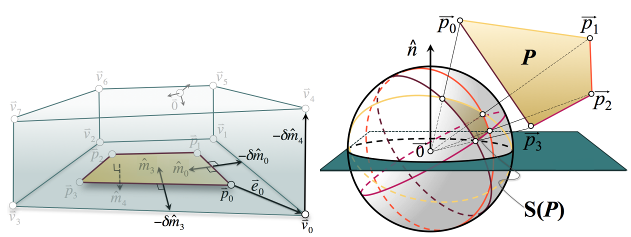

Figure 2: *Left*: AOV for a single quadrilateral P and *right* its projection to $\mathrm{S}(P)$ on the unit sphere about $\vec{x}$. This preserves its ambient occlusion, which is proportional to the projected area of $\mathrm{S}(P)$ in the tangent plane defined by $\vec{x}$ and $\vec{n}$.