Packet Traced Disc Rendering for Baking and LoD

Harrison McKenzie Chapter*
PDI|DreamWorks

Introduction

Level of detail is an important aspect of a modern production pipeline for handling complex scenes efficiently. At DreamWorks Animation, we were experimenting with new ways of reducing the rendering costs associated with our high density meshes. We introduce the pipeline we developed around OpenVDB and disc rendering, born out of our problems with geometric decimation.

Difficulties with Decimation

During our experiments with decimation we encountered one chief difficulty: while the decimators produced meshes with lower poly counts, the resulting faces often interact poorly with the subdivision and dicing schemes imposed by our proprietary REYES renderer. The resulting meshes often rasterized faster, but degraded shading performance by producing too many micropolygons.

To compound matters, we found that among artists there was little intuitive understanding of *why* poor shading results were being produced, given the reduced input complexity. Users expected a more direct, "less is faster" relationship between the number of control faces and overall frame time.

Volumes, Points and Discs

Using DreamWorks Animation's rich OpenVDB toolset, we developed a simple pipeline which converted shaded, high resolution point data into volumetric datasets. Those datasets were then sampled back into discs at lower resolutions to form possible LoD choices. Eliminating the mesh topology from the simplification process improved its predictability significantly.

Originally planned as a transitional step to rendering the simplified OpenVDB datasets directly, the discrete disc files yielded attractive properties themselves. It was easy to verify the rough look in a simple viewer, it was easy to modify, and it could be baked into from a variety of methods.

Rendering

By the disc cloud's nature, forming a reasonably closed surface requires significant overlap and intersection. The overlaps cause coverage work to be wasted on occluded surfaces, and the intersections make it hard to determine a good ordering which would skip the not-yet-occluded discs. To overcome the costs associated with disc-by-disc resolution of those conflicts, we instead packet traced the coverage for primary visibility.

On the front-end, a simple procedural groups individual discs into batches. When those batches enter visibility testing, a BVH is built over them. During the BVH build, we track all scanlines a disc's bounding box intersects, noting the minimum and maximum x coordinate.

During traversal, we compose packets of 64 rays representing each pixel's sub-pixel samples. Packets for any "live" pixels, as determined by the min/max x parameters, are traced through the BVH. The intersection records for the rays are mapped back into sub-pixel coverage masks, and proxy quads are inserted into the deep framebuffer cache for deferred shading.

For primary visibility, the discs directly take the place of micropolygons and undergo no tessellation beforehand. If the point density is chosen without care, undersampled shading can result. The cost vs. quality control afforded to the lighting lead, however, is substantial. Furthermore, the costs are intuitive and direct, following a simple "less is faster" rule in both visibility and shading.

Though most shading is baked into the discs, the deferred shading engine in the renderer allowed artists to combine the baked color with other shade time effects – such as specular lighting and dynamic shadows.

Conclusion

The level-of-detail solution we've developed and discussed here forgoes many of the automatic niceties one might expect. However, in that is a strength: its directness makes it easy to reason about, and the simplicity of its underlying representations make it easy to setup and customize.

The crew of *Mr. Peabody and Sherman* was able to adapt it into their pipeline mid-production, helping them finish the geometrically challenging elements at the end of the film in reasonable time and memory footprints.

Acknowledgments

Greg Denton, Eduardo Bustillo, Mark Newport, Amaury Aubel, Lita Cho, Shadi Almassizadeh, Dylan Lacewell

References

MUSETH, K., LAIT, J., JOHANSON, J., BUDSBERG, J., HENDERSON, R., ALDEN, M., CUCKA, P., HILL, D., AND PEARCE, A. 2013. Openvdb: An open-source data structure and toolkit for high-resolution volumes. In *ACM SIGGRAPH 2013 Courses*, ACM, New York, NY, USA, SIGGRAPH '13, 19:1–19:1.

OVERBECK, R., RAMAMOORTHI, R., AND MARK, W. R. 2008. Large ray packets for real-time whitted ray tracing. In *Interactive Ray Tracing, 2008. RT 2008. IEEE Symposium on*, IEEE, 41–48.

WALD, I., BOULOS, S., AND SHIRLEY, P. 2007. Ray tracing deformable scenes using dynamic bounding volume hierarchies. *ACM Trans. Graph. 26*, 1 (Jan.).