Environment Rendering Optimization for Pixar's *The Good Dinosaur*

Daniel Garcia
Pixar Animation Studios

Inigo Quilez
Oculus Story Studio

Dave Dixon
Pixar Animation Studios

Ariela Fedorov
Pixar Animation Studios

Matt Kuruc
Pixar Animation Studios

Susan Fisher Fong
Pixar Animation Studios

1 Stochastic Offscreen Simplification

The vast environments for *The Good Dinosaur* cover miles of terrain with procedurally generated geometry and instanced vegetation assets. Early sequence profiles revealed large memory costs attributable to tracing offscreen trees. Given the variety of vegetation in a set, many assets were not repeated enough to be instanceable and some instances only appeared offscreen. Our rendering procedurals employed both stochastic and frustum pruning. However, ray traced objects were excluded from frustum pruning or a large buffer was applied to minimize shadow popping. Stochastic pruning was still applied but as if the object was onscreen.

As the original stochastic simplification algorithm used inputs similar to renderer dicing metrics [Cook et al. 2007], we looked to recent additions to RenderMan which optimize dicing offscreen geometry for ray tracing. One of these additions applies a linear detail reduction factor based on distance outside the frustum. We similarly compute the distance for our models and map this to a multiplier on the percentage of aggregate geometry sent to the renderer. This multiplier on the percentage kept is smooth under camera motion and compatible with stochastic simplification's area preservation. This feature reduced the memory cost in some shots by almost 5GBs, a significant portion of our 16GB-32GB targets.

2 Analytical Spherical Harmonics

While we can trace shadows for focused light sources like the sun, unfocused ambient sky dome shadowing is still a large expense for a blurry effect. Precomputed spherical harmonics visibility pipelines, like our neighboring asset approach [Kuruc et al. 2014], approximate this. However, to get a consistent lighting work flow for geometry that is generated entirely at render time, we needed a technique to compute spherical harmonics coefficients without pre-baking.

To achieve this, firstly, we observe while incoming irradiance normally has to be integrated numerically (convolved) with the basis functions, for simple types of light sources like directional lighting, an analytical solution exists. Secondly, spherical harmonics integration is a linear operator. Scaling and accumulation of directional light sources yield the same result as scaling and accumulating the individual analytical coefficients.

Procedural objects encode an estimation of incoming light as a linear combination of analytical directional light sources. Based on the local position of the element within the structure, virtual light/occlusion sources could be placed at positions of largest aperture. The system allowed for creating local or global shaping based on the size of the geometrical structure to which the position of the virtual light sources would follow. Occlusion is achieved by dim light sources for the center of canopies or clumps of branches. Transmission and back lighting is achieved with soft light sources on the backside of the less occluded direction. Computing the analytical coefficients at render time required only a few dozen multiplies for each element, a minimal cost during procedural generation.

Figure 1: *Large scale environments like those found in this test shot provided the crew with many rendering challenges ©Disney/Pixar*

3 Automatic Level of Detail Cameras

Instanced geometry requires multiple discrete levels of diceable detail in our hybrid REYES/Tracing pipeline. Always instancing the most expensive representation of a tree isn't an option as it would lead to overshading. We can choose an optimal instance via screen space area for a single frame, but transitioning levels between frames can cause popping. To solve this, we place a static camera akin to a dicing camera to compute detail. Our tool identifies shots with minimal camera motion and automatically locks detail calculations to a single representative frame. This historically accounts for about 70% of our films' shots. Shots without a representative frame are flagged for manual validation and fixes.

4 Conclusion

The growing complexity of our ray traced environments required improving our toolkit with extensions to traditional techniques. Stochastic simplification with a frustum distance multiplier, analytical spherical harmonics for procedural geometry, and automatic level of detail cameras reduce memory peaks, speed up shadow computations, and minimize artifacts in our rendering pipeline.

References

Cook, R. L., Halstead, J., Planck, M., and Ryu, D. 2007. Stochastic simplification of aggregate detail. In *ACM SIGGRAPH 2007 Papers*, ACM, New York, NY, USA, SIGGRAPH '07.

Kuruc, M., Tsang, N., and Nettleship, T. 2014. Art-directable canopies in pixar's vegetation pipeline. In *ACM SIGGRAPH 2014 Talks*, ACM, New York, NY, USA, SIGGRAPH '14, 69:1–69:1.

SIGGRAPH 2015 Talks, August 09 – 13, 2015, Los Angeles, CA.
ACM 978-1-4503-3636-9/15/08.
http://dx.doi.org/10.1145/2775280.2792545