TOOMUCHVOLTAGE SOFTWARE INC.

Voxel-Based Hybrid Path-Tracing with Spatial Denoising

1) Introduction and Background

Real-time global illumination has been an elusive goal since the 1980's. With recent advancements in graphics hardware and computer graphics itself, it appears to be within reach.

We present a technique here that aims to produce global illumination for medium to large dynamic scenes with real-time to interactive performance on commodity hardware.

2) Methodology

The aim of our technique is not to exactly match the output from traditional path tracers.

Instead, we aim to produce a crude but very cheap version of the same output, fast enough to run in real-time on commodity hardware. All of this was built with dynamic scenery in mind.

To that end we've employed techniques such as rasterization based voxelization and frame by frame spatial and temporal denoising.

2.1) Preparation: Scene Voxelization, Partial Distance Transform and Gather Pass

We start off by voxelizing the scene. Static geometry is voxelized first and then combined every frame with re-voxelized dynamic geometry using a compute shader. This allows us to avoid atomicCompSwap() prescribed by [Crassin and Green 2012]. Three RGBA8 3D images are used as storage. Normals are stored in spherical coordinates (in 16 bits) and tangents and bitangents take an additional 8 bits. To store tangents and bitangents, an arbitrary ortho basis is formed using the normal. First 6 bits describe shadow casted by tangent on the arbitrarily generated X axis. Two additional bits describe whether tangents/bitangents are facing towards or away from the arbitrarily generated Y axis. As long as the arbitrary ortho basis is re-generated the same way at decode time, texture space tangents and bitangents should be recoverable (discounting lost precision due to bit packing.) Other attributes describe material properties.

At this point a partial distance transform is generated using the albedo 3D image. This will speed up path tracing which appears later in the pipeline.

Coarse voxels do not provide the best representation of complex scenery. Thus we perform a gather pass -- that also captures the attributes/materials describe above -- to use as our first bounce for path tracing later on.

Buffers Generated during Pre-process

2.2) Path Trace

Path tracing (with NEE and MIS) is done as a post-process. Emissive surface hints (along with surface area measurements) are uploaded via CPU-side UBO uploads. Ray traversal is accelerated via sphere-marching. BRDF of choice is Anisotropic Ward and is importance-sampled (see [Walter 2005]). The exact sampling scheme used is a less problematic approximation of [Walter 2005]'s as it avoids the tan() term.

NEE with sun light shining from above

2.3) Temporal Accumulation

Last 10 frames are kept for temporal back projection and irradiance accumulation. Disocclusion discards historical frames and fragment anisotropy weights historical contribution. A perfectly lambertian surface places equal weight on historical data while a perfect mirror only uses one frame. For any material in-between a Hermite curve (smoothstep) is used.

Hermite curve for relevance of historical data for a glossy surface

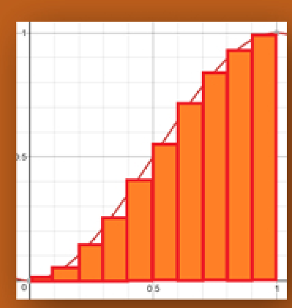

2.4) Spatially Denoise

Four passes make up the denoiser: the first two are radial and target tangent and bitangent directions in screen space (scaled by anisotropic roughness values.) The remaining two simply smoothen the results from the previous radial passes using a box blur. A radial cut-off alongside normal weighting prevents light bleeding. However, the cut-off distance increases further away from the viewer as light information becomes scarce.This results in some light bleeding.

The kernel formed by the first two passes forms a trapezoid.

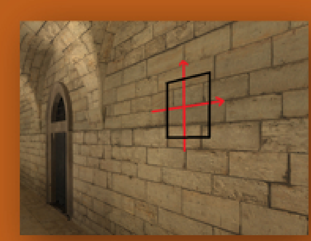

2.5) Upsample and Modulate

Since path-tracing is done at a lower resolution (640x480), the resulting image needs to be upsampled and modulated with the albedo image from the gather pass. This is done via an inside-out march to find the best candidate fragment to use for modulation with the center fragment. Position and normal similarities both factor into the candidate search.

Final Image

3) Results and Conclusion

A crude but cheap technique was presented to capture global illumination for moderate to large dynamic scenes with real-time to interactive performance. It is worth noting that the technique's trapezoidally shaped filter allows capturing of directional highlights such as the ones shown below. A performance comparison was done with a glossy version of Sponza at 1SPP against Baikal. Both techniques used NEE. Our technique consistently performed at 15 FPS while Baikal consistently performed at 6.7 FPS given the same resolution. However, it should be noted that Baikal is mainly optimized for use on Radeon hardware.

Highlights from Directional Roughness

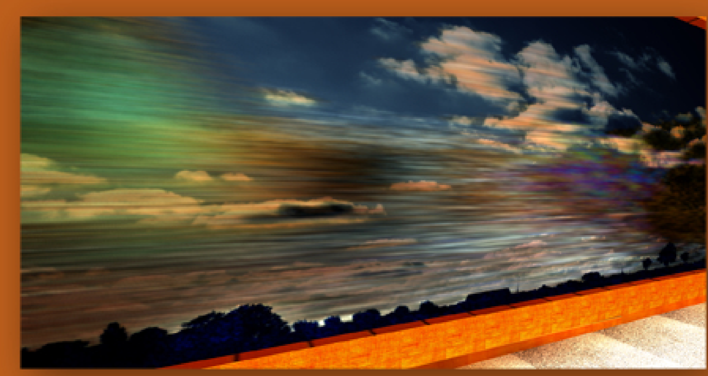

4) References

• Cyril Crassin and Simon Green. 2012. . CRC Press, Patrick Cozzi and Christophe Riccio. http://www.seas.upenn.edu/~pcozzi/OpenGLInsights/ OpenGLInsights-SparseVoxelization.pdf

• Walter, Bruce. "Notes on the Ward BRDF." (2005).