Voxel Based Hybrid Path Tracing with Spatial Denoising

Baktash Abdollah-shamshir-saz
TooMuchVoltage Software Inc.
baktash@toomuchvoltage.com

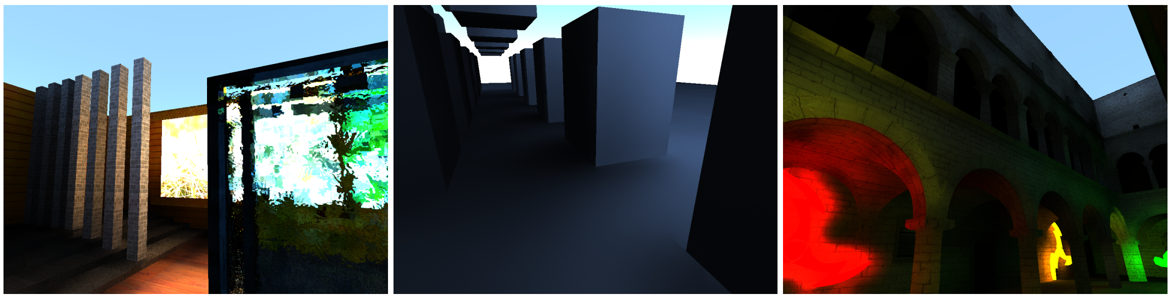

Figure 1: (left) a sample scene with emissive surfaces (center) a sample scene with most irradiance from the sky (right) The Sponza atrium with some emissive surfaces at the bottom

ABSTRACT

We present a novel rendering technique which allows real-time to interactive rendering of global illumination for small to large dynamic scenes on commodity hardware. This technique operates entirely within graphics processors, simulates many optical phenomena, can be adjusted via grid coarseness to trade off frame rate for reflection quality and is able to include GPU accelerated special effects such as tessellation. We achieve real-time and interactive performance on small and large scenes on a GeForce 970M GTX video card.

CCS CONCEPTS

• **Computing methodologies** → **Rendering**; **Ray tracing**;

KEYWORDS

voxel, path tracing, global illumination, real-time, denoising

ACM Reference format:
Baktash Abdollah-shamshir-saz. 2018. Voxel Based Hybrid Path Tracing with Spatial Denoising. In *Proceedings of I3D Symposium, Montreal, Quebec, Canada, 15-18 May 2018,* 2 pages.
https://doi.org/10.1145/8888888.7777777

I3D Symposium, 15-18 May 2018, Montreal, Quebec, Canada

ACM ISBN 978-1-4503-1234-5/17/07.
https://doi.org/10.1145/8888888.7777777

1 INTRODUCTION

Real-time global illumination for dynamic scenes has been an elusive goal ever since the publication of [Kajiya 1986]. With advancements in computer graphics and graphics hardware, many techniques have been implemented to tackle this challenge in real-time. [Kaplanyan and Dachsbacher 2010] was the first technique to demonstrate feasibility of real-time global illumination on limited hardware (gaming consoles) with a small memory footprint. [Crassin et al. 2011] was proposed as an alternative much closer to path-tracing. However, it also appeared to have a huge memory price tag (1024MB for the Sponza atrium). The technique proposed in [McLaren and Yang 2015] reduced this memory consumption by using viewer-centered cascades of size 32x32x32 (cached at distant cascades) and was explained in detail in [McLaren 2015]. Most real-time voxel-based approaches for global illumination uniformly cone trace against pre-illuminated (and often mipped) voxel data. Our approach instead combines ideas from [Crassin et al. 2011] and [McLaren 2015] to perform low sample-count Monte Carlo path-tracing followed by denoising as a post-process.

2 OVERVIEW

We start off by voxelizing the entire scene every frame using [Crassin and Green 2012] without the use of atomics. Stability is achieved by using a compute shader to only re-voxelize dynamic segments of a scene. The voxelized image encodes such information as refractive index, bi-directional roughness, albedo, emissivity, world space normals, tangents and bitangents. A novel encoding scheme is used to encode a tangent and bi-tangent into only eight bits (given the packed normal in 24 bits). We avoid memory concerns arising from this volume of information by avoiding directional voxel information which would require six times the memory usage. Thus, geometry has to be thick enough to avoid light leaks and other issues.

Once the scene is voxelized, we perform a gather pass to render the same information above into multiple screen-sized 2D floating point render targets. Subsequently, uni-directional path tracing begins as a post-process at one voxel above the surface. This elevation is performed using world space face normals to avoid self occlusion. A 512x512 random texture covers the space skewed along the z-axis to provide randomness everywhere in the scene for reflections and refractions. Fragments with higher roughness will trace more samples per pixel – with a maximum of six – for most overall roughness. Many laws of optics are observed: reflections happening inside the specular lobe respect bi-directional roughness in the tangent or bitangent directions. Schlick's approximation [Schlick 1994] decides the probabilities of rays reflecting or refracting, providing reflections at grazing angles. Normal maps are used liberally. Primitive normals, tangents and bitangents are adjusted during the tessellation process using the height map.

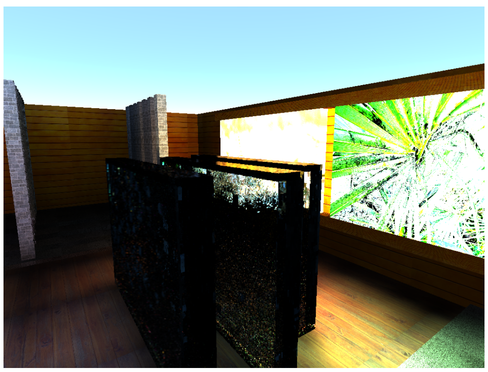

Figure 2: Reflections on top of glass at grazing angles.

The path traced output (which only holds light information) is then denoised through a two-pass, five-phase blur process. Each phase blurs a certain overall roughness bracket and above, and each pass either targets that bracket vertically or horizontally (to reduce sampling pressure). A fragment's roughness bracket is determined by the higher of its tan or bi-tan roughness. Fragments falling beyond a certain radial cut-off are excluded from contribution. Contributing fragments are weighted by the clamped dot product of their normals against the center fragment's normal. The kernel and the radial cut-off are increased for fragments further away from the viewer (every four feet in the virtual world) as scene illumination detail becomes scarce in the distance. This might result in some over-blurring in the distance.

The denoised information is then fed to the upsampler which does an outward search from the center fragment to find the best candidate for lighting modulation. The best candidate fragment is closest to the center fragment and shares a highly similar normal. Once the best candidate is found its light amount is modulated with the center fragment's albedo. This process allows us to path trace and de-noise the wooden scene above at 640x480 and upsample to full 1080p at 30+ frames per second on a GeForce 970M GTX.

3 RESULTS

A performance comparison was done between AMD Radeon's Baikal LBVH path tracer and our method's base implementation on the Sponza atrium. Our technique averaged 17.5 frames per second while Baikal averaged around 7.4 on a GeForce 970M GTX at 1024x768. However, it must be noted that Baikal was also tracing secondary rays towards the sun. This can be implemented using shadow maps and direct lighting for our implementation with minimal impact on performance. Our speed gain can be attributed to lack of intersection computations.

4 LIMITATIONS

Perfectly smooth specular surfaces will show blocky reflections since all intersections emanating from the gather buffer collide with voxels. Some blockiness is also exhibited around object silhouettes (most visible on glass) due to limited encoded surface information. This can be reduced or eliminated using a contouring scheme similar to what is described in [Laine and Karras 2011]. Further research is required to combine this with rasterization-based voxelization in an efficient manner.

5 CONCLUSIONS AND FUTURE WORK

With renewed emphasis on increasing memory bandwidth for recent video cards, sampling-heavy techniques such as this should become more popular and performant. Future improvements include bi-directional path tracing to reduce the number of samples required for good convergence, using signed distance fields for faster ray traversal and employing temporally stable denoisers such as [Schied et al. 2017] or [Mara et al. 2017].

REFERENCES

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

Cyril Crassin, Fabrice Neyret, Miguel Sainz, Simon Green, and Elmar Eisemann. 2011. Interactive Indirect Illumination Using Voxel Cone Tracing. *Computer Graphics Forum* 30, 7 (2011), 1921–1930. https://doi.org/10.1111/j.1467-8659.2011.02063.x

James T. Kajiya. 1986. The Rendering Equation. *SIGGRAPH Comput. Graph.* 20, 4 (Aug. 1986), 143–150. https://doi.org/10.1145/15886.15902

Anton Kaplanyan and Carsten Dachsbacher. 2010. Cascaded Light Propagation Volumes for Real-time Indirect Illumination. In *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D '10)*. ACM, New York, NY, USA, 99–107. https://doi.org/10.1145/1730804.1730821

S. Laine and T. Karras. 2011. Efficient Sparse Voxel Octrees. *IEEE Transactions on Visualization and Computer Graphics* 17, 8 (Aug 2011), 1048–1059. https://doi.org/10.1109/TVCG.2010.240

Michael Mara, Morgan McGuire, Benedikt Bitterli, and Wojciech Jarosz. 2017. An Efficient Denoising Algorithm for Global Illumination. In *Proceedings of High Performance Graphics*. ACM, New York, NY, USA. https://doi.org/10.1145/3105762.3105774

James McLaren. 2015. The Technology of The Tomorrow Children. (2015). http://fumufumu.q-games.com/archives/TheTechnologyOfTomorrowsChildrenFinal.pdf

James McLaren and Tao Yang. 2015. The Tomorrow Children: Lighting and Mining with Voxels. In *ACM SIGGRAPH 2015 Talks (SIGGRAPH '15)*. ACM, New York, NY, USA, Article 67, 1 pages. https://doi.org/10.1145/2775280.2792546

Christoph Schied, Anton Kaplanyan, Chris Wyman, Anjul Patney, Chakravarty R. Alla Chaitanya, John Burgess, Shiqiu Liu, Carsten Dachsbacher, Aaron Lefohn, and Marco Salvi. 2017. Spatiotemporal Variance-guided Filtering: Real-time Reconstruction for Path-traced Global Illumination. In *Proceedings of High Performance Graphics (HPG '17)*. ACM, New York, NY, USA, Article 2, 12 pages. https://doi.org/10.1145/3105762.3105770

Christophe Schlick. 1994. An Inexpensive BRDF Model for Physically-based Rendering. *Computer Graphics Forum* 13, 3 (1994), 233–246. https://doi.org/10.1111/1467-8659.1330233