Real-Time Foveated Raytracing for Virtual Reality Using Modern Acceleration Techniques

CAMERON MCWILLIAMS, University of Washington

1 ELEVATOR PITCH

Previously developed foveated raytracing rendering techniques for virtual reality will be demonstrated, employing acceleration features available in the latest generation of consumer graphics cards. A scene will be rendered for real-time viewing using a virtual reality headset. For each frame, a sampling distribution over the visual field will be constructed, informed by the user's gaze and a model of visual accuity over the retina. The scene will be sampled with the constructed distribution using raytracing techniques assisted by modern hardware acceleration. In the initial design, samples will be shaded with Phong lighting, raytracing being used only to sample geometry. The resulting samples will be resampled to form a final raster frame to be displayed. The real-time performance of this rendering method will be evaluated. As a first stretch goal, the raytracing will be extended to the shading. As a second stretch goal, the rendering will be further accelerated using more advanced sampling techniques, like reprojection and denoising.

2 EXTENDED OVERVIEW

With the improvement in quality of head-mounted displays, the immersive potential of virtual reality's visuals is increasingly being limited by the ability to generate photorealistic visuals. The offline generation of photorealistic images on consumer hardware was only considered tractable within the last decade, and real-time generation of images with more than very limited or emulated photorealistic effects is an invention of the past few years, still yet to see widespread use. With virtual reality demanding resolutions and framerates above that of desktop applications in order to provide a comfortable sensory experience for users, photorealism is sacrificed in exchange for performance. However, as the performance of graphics hardware improves new techniques become tractable for virtual reality applications.

Raytracing has long been the standard technique for generating photorealistic images in animation and film. However, full raytracing was formerly considered impossible even for the lesser demands of desktop applications, and thus implementations of raytracing in virtual reality have relied on acceleration techniques. A common acceleration technique in both raster and raytraced virtual reality rendering is reducing the render resolution at higher retinal eccentricities (i.e. "foveation"). Prior research has used sampling with raytracing performed on standard compute cores, with successful real-time performance only being achieved for simple geometry sampling (see e.g. Weier et al. [2016]), making raytracing not a substantially more useful technique as compared to raster rendering. With the recent development of dedicated raytracing acceleration in graphics hardware, the tractability of employing raytracing has increased. Fully ray-traced illumination has been demonstrated in desktop applications, indicating that modern graphics devices may be capable of producing images for virtual reality using primarily raytracing techniques. Extrapolating demonstrated results on previous generation graphics devices to advertised performance of current generation devices suggests that foveated raytracing is perfectly tractable using already developed techniques.

Raster foveation requires the use of resampling from a retinal polar coordinate system. Comparatively, raytracing provides a great advantage in this regard, as the sampling distribution used during

Author's address: Cameron McWilliams, sndwv@uw.edu, University of Washington.

raytracing acts as the retina space sampling as well. As a natural consequence of the foveated sampling, not all pixels of a frame will receive a sample. Previous methods of constructing a full raster image include artificial intelligence-based reconstruction, and accumulation and spatial reprojection. Siekawa et al. [Siekawa et al. 2019] proposed a simpler technique, seeing sufficient quality from simple interpolation between samples. The periphery of the human vision is particularly sensitive to changes in contrast and motion, so raytracing techniques must incorporate some temporal filtering technique. Applications in the industry have primarily used some variant of a technique called Spatio-Temporal Variance Guided Filtering [Schied et al. 2017], however as noted by Siekawa et al., even simple temporal averaging provides a passable quality result.

2.1 Technical Challenges

This project involves addressing the following key technical challenges.

- Constructing a scene able to fill the field of view that incorporates opportunities for visual quality to be improved with ray tracing.
- Developing a raytracing rendering engine using modern hardware acceleration for raytracing.
- Implementing a foveated resampling technique for final raster frame generation.
- Developing simple benchmarks of the rendering technique's performance.
- Optionally, implementing additional ray tracing techniques to improve visual quality of the scene.
- Optionally, improving performance with more advanced acceleration techniques, such as spatial reprojection or spatiotemporal filtering.

2.2 Key Risks and Mitigations

The following key risks and potential implementation alternatives have been identified.

- Hardware acceleration alone may not succeed in providing a real-time experience at current head-mounted display framerates. If this cannot be sufficiently mitigated by lowering framerate or sampling density, demos may emulate varying levels of (theoretical) graphics device performance via sampling of a raster render.

3 HARDWARE AND SOFTWARE

This project requires the following hardware.

- [Requested] HTC Vive Pro Eye. Open to alternate eye-tracking solutions. No eye-tracking solution as a last resort.
- [Personal] PC with RTX-capable graphics card: Provided by student. VR-capable laptop with RTX 2070 Max-Q.

The following hardware would be useful for the project.

- [Requested] Testbench with RTX-capable graphics card with performance exceeding 4 gigarays per second (i.e. Desktop RTX 2060 or better).

4 DEVELOPMENT PLAN

The aim is to complete this project over three weeks, with the following major milestones.

- **March 4**: Complete raytraced geometry sampling.
- **March 6**: Complete sample generation and resampling pass.
- **March 7**: Implement eye tracking.
- **March 9**: Complete demo scene.
- **March 11**: Design and run benchmarks.

- **March 14**: (Optional) Implement additional raytracing techniques.
- **March 16**: (Optional) Implement additional filtering/acceleration techniques.
- **March 18**: Submit report.
- **March 19**: Prepare materials for the demo session.
- **March 20**: Demo session.

REFERENCES

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.

Adam Siekawa, Rados law Mantiuk, and Rafa l Piorkowski. 2019. Foveated Ray Tracing for VR Headsets. (2019).

Martin Weier, Thorsten Roth, Ernst Kruijff, André Hinkenjann, Arsène Pérard-Gayot, Philipp Slusallek, and Yongmin Li. 2016. Foveated Real-Time Ray Tracing for Head-Mounted Displays.