Concept for Rendering Optimizations for Full Human Field of View HMDs

Daniel Pohl*
Intel Corporation

Nural Choudhury†
Akanoodles Design

Markus Achtelik‡
Intel Corporation

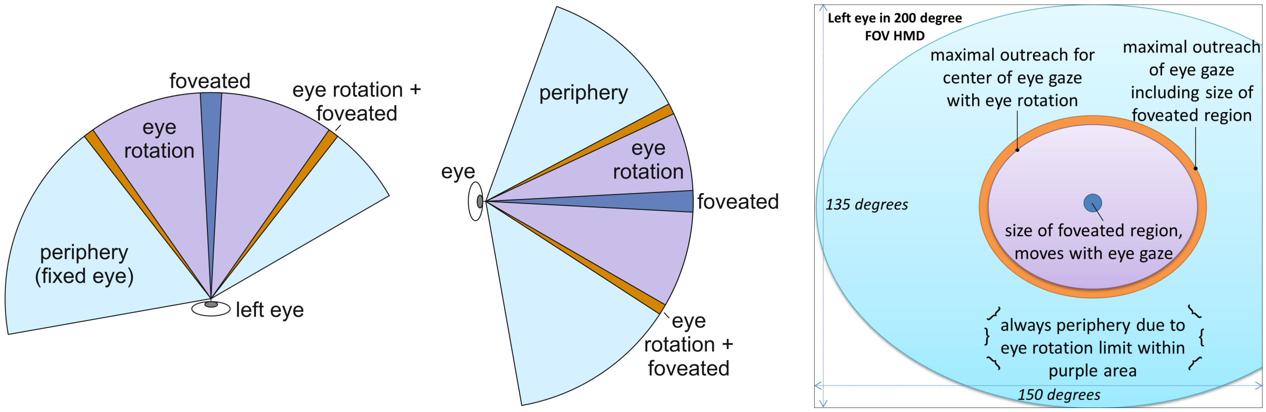

Figure 1: The human field of view including eye movement is displayed in the first two images. Using these properties for rendering in a head-mounted display with full field of view leads to the regions shown in the right image.

Abstract

To enable high immersion for virtual reality head-mounted displays (HMDs), a wide field of view of the display is required. Today's consumer solutions are mostly around 90 to 110 degrees field of view. The full human field of view for both eyes together has been measured to be between 200 and 220 degrees. Prototypes of HMDs with such properties have been shown. As the rendering workload increases with more pixels to fill the field of view, we propose a novel rendering method optimized for HMDs that cover the full human field of view. We target lower end HMDs where the cost of eye tracking would increase the price too much. Our method works without eye tracking, making use of certain human vision properties that appear once the full human field of view is covered. We achieve almost twice the rendering performance using our method.

Index Terms: Computing methodologies—Rendering

1 Introduction

With the introduction of affordable wide-angle HMDs like the Oculus Rift DK1 with 90 degrees field of view, it was possible for consumers to have immersive experiences at home. Even at the limited resolution of 1280x800 pixels, less than ideal photons-to-motion latency and the lack of head position tracking, the human brain got easily tricked into believing to be in another world. As vision is the human's strongest sense, a wide field of view is critical to deliver immersion. In 2017, the most popular consumer head-mounted displays share almost the same field of view between 90 and 110 degrees. However, there are several indications for a larger field of view for future head-mounted displays. For example, in a talk from Michael Abrash at Oculus Connect 3 in 2016, it was revealed that Oculus Research is setting their expectations of having head-mounted displays with 140 degrees field of view available in the year 2021[1]. Panasonic showed a prototype of a HMD with 220 degrees[2] and StarVR[3] works on a HMD with 210 degrees. A Kickstarter campaign from Pimax[4] promises to deliver an HMD with 200 degrees field of view in the price range of other consumer HMDs. Therefore, consumer HMDs enabling a full field of view are becoming accessible soon. Looking ahead in the market of such devices, it can be expected, that there will be higher-priced HMDs with many features like eye tracking, that can ease the rendering load. On the other side, it is likely that there would be a market for lower-priced HMDs where adding eye tracking would break the financial target. In this paper, we target these HMDs and offer new rendering optimizations based on the human vision system without the need for eye tracking.

2 Related Work

The concept to change rendering based on where a human looks has been used in experiments by Tong and Fisher [7] in a flight simulator which was projected onto a dome. Besides using a low quality image for the full dome, another high quality image at the eye gaze was used and blended in. Levoy and Whitaker [5] change the number of rays that are used in a volume ray tracer, depending on the eye gaze. Guenter et al. [3] showed a foveated rendering approach using an eye tracker in a setup with multiple monitors and adjusted the quality depending on the tracked eye gaze. Eye tracking in HMDs was demonstrated by Duchowski et al. [2]. Compared to the previous methods, we target lower end HMDs and will therefore work without eye tracking. Pohl et al. [6] introduced lens-based rendering optimizations for head-mounted displays in form of a sampling map which defines during ray tracing for every pixel which quality level (e.g. anti-aliasing) it should receive. Graphics hardware vendors integrated that concept for rasterizers in their SDKSs (NVidia VRWorks

*e-mail: daniel.pohl@intel.com
†e-mail: mail@akanoodles.com
‡e-mail: markus.achtelik@intel.com

[1]uploadvr.com/abrash-2021
[2]roadtovr.com/panasonics-220-degree-vr-headset-uses-4-screens-crazy-fused-lenses
[3]starvr.com
[4]kickstarter.com/projects/pimax8kvr/pimax-the-worlds-first-8k-vr-headset

"Multi-Res Shading"[5], AMD LiquidVR "MultiRes Rendering"[6]). Compared to our method, the prior optimizations are not based on human vision properties which come into effect, once the full field of view is displayed inside the HMD. Their optimizations are used because the sweet spot of the lens of the current generations HMDs is so small, that mostly the center area through the lens of the HMD is perceived sharp. The surrounding areas are blurred and cannot represent high detail which is the base of the prior methods. With higher quality lenses in future HMDs, it can be expected that the blur from lenses will be minimized.

3 Human Vision

As every human is different, numbers for the full field of view are often defined in the range between 200 and 220 degrees [4] [1]. The maximum amount of eye rotation to the left or right is 35 degrees, looking up at 25 degrees and down at 30 degrees[7]. The foveated region is around 5-6 degrees field of view [3]. Everything outside that area is peripheral vision which is perceived with significantly less sharpness. The peripheral area covers about 110 degrees field of view [8]. This results in the fact, that the full human field of view has always regions where it cannot see certain areas sharp, even when trying to rotate the eyes into that direction. When a human wants to see something in that region better, the head is rotated towards that direction. The properties remain at the new orientation: again, the most outer regions are not perceived sharp. Therefore, we can apply this concept to a HMD, where a fixed screen is around the user's eyes, independent of the head orientation of the user.

4 HMD Setup and Sampling Map

During this work, we did not have a HMD that covers the full human field of view. We assume having a target HMD which covers 200 degrees horizontal and 130 degrees vertical field of view. We assume an overlap of 50 degrees horizontally for each eye. That way, a single view for the left or right eye displays 150 degrees of content. Knowing the human peripheral vision properties and the maximum of eye rotations, we designed a sampling map for rendering. We laid out the values for the left eye as shown in Figure 1 (right). Based on that, we converted it into a sampling map: 8 bit values in the image indicate which level of detail is used by the renderer. The value 255 represents highest quality, 16 is an example of very low quality and 0 represents that rendering can be skipped in that area completely as it is never visible. The sampling map and its application is shown in Figure 2.

5 Rendering Setup and Benchmarks

We use our own implementation of a CPU-based ray tracing rendering system, partly accelerated by Intel Embree [9]. A benefit of ray tracing that the quality level can be easily changed on a per pixel basis. The renderer interprets a sampling map to determine the amount of supersampling per pixel. We use a workstation with a Dual-CPU (Intel Xeon E5-2699 v3). The rendering resolution is 1920×864 pixels for the stereoscopic image. We realize that in order to fill such a large field of view more resolution would be desired. We believe that with hardware-accelerated GPU rendering this could be achieved at higher frame rates. Nevertheless, in order to evaluate the performance impact of our method, the chosen resolution is sufficient. The content is the "island" level from *Enemy Territory: Quake Wars*[8]. The results are in Table 1.

6 Conclusion and Future Work

For higher priced HMDs with full human field of view, as known, foveated rendering with eye tracking can be used to render faster.

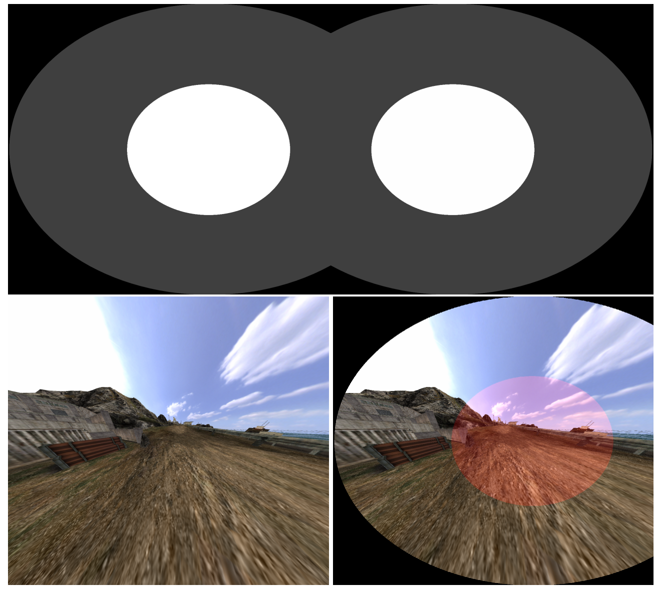

Figure 2: Top: sampling map for a HMD with 200 degrees FOV. The bottom left image shows the brute force 4x supersampled version of the left eye view. The bottom right image shows the optimized version using a sampling map, which defines that only the area within the red circle receives 4x supersampling. Within this area, the user can rotate the eyes and perceive sharp content within his foveated view.

Table 1: Benchmark in frames per second. Higher is better.

Supersampling	**Brute force**	**Our method**
4x	21	40
8x	16	32

However, for cheaper HMDs with full human field of view where adding eye tracking would be too expensive, our optimizations using human vision properties can be applied. A significant performance speed-up of almost 100 percent is achieved without the need for eye tracking. In future work, we would like to verify our results on a full field of view HMD and confirm in a user study that our optimizations can be used without differences in perceived quality.

[5] developer.nvidia.com/vrworks

[6] amd.com/en/technologies/vr

[7] www.hazardcontrol.com/print.php?fs=humanfactors&p=visual-acuity-and-line-of-sight, Copyright Nelson and Associates, Bryan, Texas

[8] id Software and Splash Damage

References

[1] G. Dagnelie. *Visual prosthetics: physiology, bioengineering, rehabilitation*. Springer Science & Business Media, 2011.

[2] A. T. Duchowski, E. Medlin, A. Gramopadhye, B. Melloy, and S. Nair. Binocular eye tracking in VR for visual inspection training. In *Virtual Reality Software and Technology*, 2001. doi: 10.1145/505008.505010

[3] B. Guenter, M. Finch, S. Drucker, D. Tan, and J. Snyder. Foveated 3D graphics. *ACM Transactions on Graphics (TOG)*, 31(6):164, 2012.

[4] D. O. Harrington. *The visual fields: A textbook and atlas of clinical perimetry*, vol. 299. CV Mosby, 1976.

[5] M. Levoy and R. Whitaker. Gaze-directed volume rendering. In *I3D*, pp. 217–223, 1990.

[6] D. Pohl, T. Bolkart, S. Nickels, and O. Grau. Using astigmatism in wide angle HMDs to improve rendering. In *2015 IEEE Virtual Reality (VR)*, pp. 263–264, March 2015. doi: 10.1109/VR.2015.7223396

[7] H. Tong and R. Fisher. *Progress Report on an Eye-Slaved Area-of-Interest Visual Display*. Defense Technical Information Center, 1984.

[8] H. M. Traquair and G. I. Scott. *Clinical perimetry*. Kimpton, 1957.

[9] I. Wald, S. Woop, C. Benthin, G. S. Johnson, and M. Ernst. Embree: A kernel framework for efficient CPU ray tracing. *TOG (Proc. SIGGRAPH)*, 33(4):143:1–143:8, 2014.