High Resolution, Low Latency Image Synthesis for Head Mounted Displays

A novel patented method using perceptual rasterization resulting in high-resolution field-of-view foveated images with low latency.

Category
Software

Learn more

Image credit: I Stock (bymuratdeniz)

Background

Head-mounted displays (HMDs) have requirements beyond typical desktop display-based systems: HMDs must maintain low and predictable latency and must cover a significant proportion of the user's field of view at high resolution.

Today's HMD technology has challenges which result in breaks-in-presence, simulator sickness, and reduced performance:

Current HMD graphics pipelines struggle to produce the images required. They are limited by the number of pixels that can be pushed from the GPU to the display, which strains the raw bandwidth of even the latest GPUs.

Traditional graphics pipelines compute images at a single snapshot time and ignore how the display is driven image will be perceived. Many displays support low persistence by scanning the image illumination, but this results in some parts of the screen appearing at lower latency than others.

The single pass perceptual rasterization method presented offers an improved approach for a graphics pipeline to deliver efficient, high-resolution foveated images with low latency. Solutions such as this may help further market adoption and realize the impactful future of virtual and augmented reality by achieving more comfortable viewing.

This approach overcomes limitations of warping with respect to disocclusions, object motion and view-dependent shading, as well as geometric aliasing artifacts in other foveated rendering techniques.

Technology Overview

UCL researchers suggest a rasterization pipeline tailored towards the needs of HMDs, where latency and field-of-view requirements pose new challenges beyond those of traditional desktop displays. Instead of image warping for low latency, or using multiple passes for foveation, they show how both can be produced directly in a single perceptual rasterization pass. They do this with per-fragment ray-casting. This is enabled by derivations of tight space-time-fovea pixel bounds, introducing just enough flexibility for the requisite geometric tests, but retaining most of the simplicity and efficiency of the traditional rasterization pipeline (Figure 1).

The unique combination of techniques included in this patented method include:

Rolling Rasterization: a method for GPUs to synthesize rolling images where the time at each pixel depends on its display location.

Foveation: a method to synthesize images with spatially varying pixel density. Important areas are larger and cover more pixels while less relevant areas are smaller, with fewer pixels.

Watch here: https://youtu.be/Xc46CaMYnCc

Further Details:

Interactive webpage

Benefits

Reduced latency with rolling images

High resolution foveated images

Solution can be implemented in graphics pipeline or directly in hardware

Can be used/configured with any type of HMD

Applications

The perceptual rasterization technique can be used for displays in VR/AR, aerospace, automotive, military applications, production studios, and others. The improved resolution at high field-of view with lower latency enhances user experience, supporting an increased adoption of such displays across multiple applications.

Opportunity

Exclusive or non-exclusive licensing opportunities.

Figures

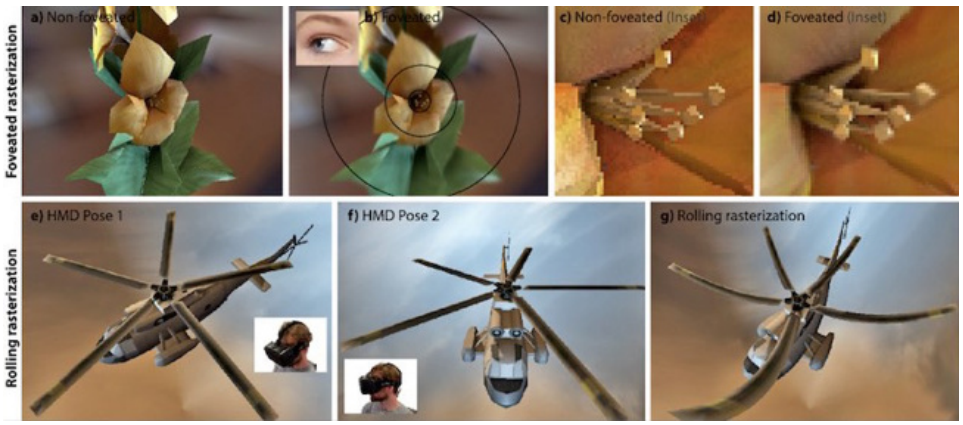

Patents

International patent application number PCT/GB2019/051520

IP Status

Patent application submitted

Seeking

Licensing