

Exploiting the limitations of spatio-temporal vision for more efficient VR rendering

Gyorgy Denes
University of Cambridge
gyorgy.denes@cl.cam.ac.uk

Kuba Maruszczyk
University of Cambridge
kuba.maruszczyk@cl.cam.ac.uk

Rafał K. Mantiuk
University of Cambridge
rafal.mantiuk@cl.cam.ac.uk




Figure 1: Our method renders even-numbered frames at a lower resolution, saving on rendering time and transmission bandwidth. Low resolution frames are upsampled before being displayed; high resolution frames are compensated for the lost information. When such a sequence is viewed at a high frame rate, the frames are perceived as if they were rendered at full resolution.

CCS CONCEPTS

• **Computing methodologies** → **Rendering; Virtual reality; Perception;**

KEYWORDS

Temporal multiplexing, motion quality

ACM Reference Format:

Gyorgy Denes, Kuba Maruszczyk, and Rafał K. Mantiuk. 2018. Exploiting the limitations of spatio-temporal vision for more efficient VR rendering. In *Proceedings of SIGGRAPH '18 Posters*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3230744.3230760>

1 INTRODUCTION

Increasingly higher virtual reality (VR) display resolutions and good-quality anti-aliasing make rendering in VR prohibitively expensive. The generation of these complex frames 90 times per second in a binocular setup demands substantial computational power. Wireless transmission of the frames from the GPU to the VR headset poses another challenge, requiring high-bandwidth dedicated links.

Current approaches of rendering in VR with limited computational resources often utilise reprojection techniques [Beeler et al. 2016; Didyk et al. 2010; Vlachos 2016]. The assumption of these techniques is that the rendering cost can be reduced by drawing

only every k -th frame, and generating in-between frames by transforming the previous frame. However, as reprojection often produces unexpected distortions and artifacts when encountering rapid brightness change, occlusions and repeated patterns, these techniques should only be used as a last resort. Vlachos [2016] suggests that displaying a lower resolution frame is often preferred.

We propose a conceptually simple and robust technique for reducing both bandwidth and rendering cost for high-frame-rate displays by 25–49% with only marginal computational overhead and small impact on image quality. Our technique, Temporal Resolution Multiplexing (TRM), can also be applied to future high-refresh-rate desktop displays and television sets to improve motion quality.

2 OUR APPROACH

Temporal Resolution Multiplexing (TRM) takes advantage of the limitations of the human visual system: the finite integration time that results in fusion of rapid temporal changes, along with the inability to perceive both high spatial and temporal frequency signals at the same time. An illusion of smooth, high-frame-rate motion is generated by rendering a low-resolution version of the content for every even-numbered frame and compensating for the loss of information by modifying every odd-numbered frame. When such these frames are viewed at a high frame rate (> 90 Hz), the visual system fuses them and perceives the original, full resolution video. TRM also reduces motion blur artifacts as illustrated in Figure 2.

The proposed technique, although conceptually simple, requires much attention to details, such as overcoming dynamic range limitations, ensuring that potential flicker is invisible, display calibration and designing the solution that will save both rendering time and bandwidth.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGGRAPH '18 Posters, August 12–16, 2018, Vancouver, BC, Canada

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5817-0/18/08.

<https://doi.org/10.1145/3230744.3230760>




Figure 2: Simulation of perceived motion blur for full frame rate (120 Hz), reduced frame rate (60 Hz) and TRM with different resolution reduction factors.

3 IMPLEMENTATION AND FUTURE WORK

We first tested the technique in a custom OpenGL/OpenVR-based software. According to our measurements as shown in Figure 3, an unoptimized implementation of TRM offered a 19-25% performance improvement for a fill-rate bound football scene (Figure 4 left).

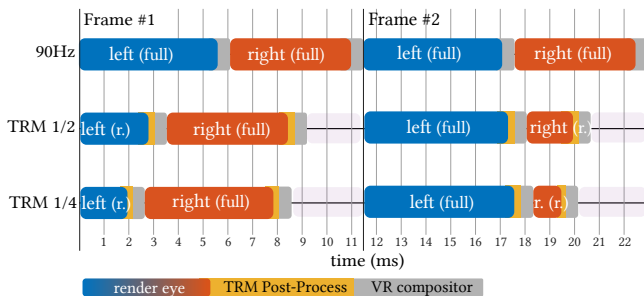


Figure 3: Performance measurement of standard 90 Hz rendering for two consecutive frames averaged over 1500 samples (top); compared with TRM with different resolution reduction factors (center and bottom). For each frame (full) indicates full-resolution, (r.) indicates reduced-resolution rendering.

As TRM is perceptually motivated, we validated its performance in a subjective experiment. We used pairwise comparison with sequential presentation for the virtual scenes in Figure 4. Participants compared TRM with baseline rendering and two alternative techniques: NCSFI (blur reduction) [Chen et al. 2005] and ASW (reprojection) [Beeler et al. 2016]. The experiment was performed on an HTC Vive and an Oculus Rift CV1. Nine paid participants with normal or corrected-to-normal vision took part (Figure 4). Results indicate that TRM was almost indistinguishable from full-resolution rendering and significantly better than NCSFI or ASW (Figure 5). In post experiment interviews we verified that observers perceived no flicker while using TRM on 90 Hz.

In conclusion, TRM is a robust, attractive alternative to dropped frames and naïve resolution reduction, and produces none of the obvious artifacts of reprojection techniques. It is easy to integrate into




Figure 4: Stimuli used for validation.

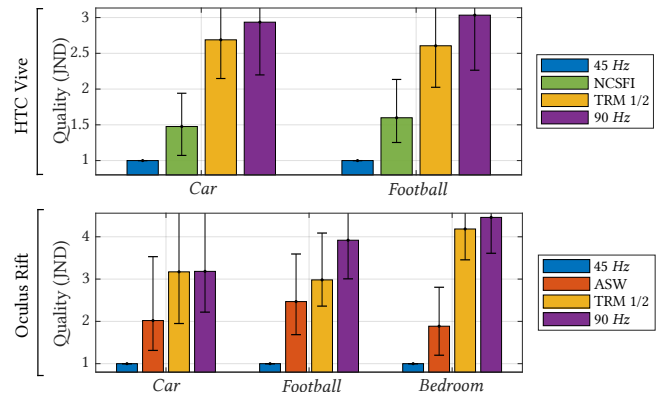


Figure 5: Results of pairwise comparison experiments. TRM with a resolution reduction factor of 1/2 is almost indistinguishable from full 90 Hz rendering. Similar state-of-the-art techniques NCSFI and ASW were deemed to have significantly lower quality by the participants.

existing rendering pipelines, fast to compute, and can be combined with other visual coding methods, such as chroma-subsampling and video codecs such as h.265 and JPEG XS to further reduce bandwidth.

Currently we are exploring the parameter space in order to determine the relationship between the display refresh rate and the maximum possible resolution reduction that results in no noticeable quality degradation.

REFERENCES

Dean Beeler, Ed Hutchins, and Paul Pedriana. 2016. Asynchronous Spacewarp. <https://developer.oculus.com/blog/asynchronous-spacewarp/>. (2016). Accessed: 2018-05-02.

Hanfeng Chen, Sung-soo Kim, Sung-hee Lee, Oh-jae Kwon, and Jun-ho Sung. 2005. Nonlinearity compensated smooth frame insertion for motion-blur reduction in LCD. In *2005 IEEE 7th Workshop on Multimedia Signal Processing*. IEEE, 1-4. <https://doi.org/10.1109/MMSP.2005.248646>

Piotr Didyk, Elmar Eisemann, Tobias Ritschel, Karol Myszkowski, and Hans Peter Seidel. 2010. Perceptually-motivated real-time temporal upsampling of 3D content for high-refresh-rate displays. *Computer Graphics Forum* 29, 2 (2010), 713-722. <https://doi.org/10.1111/j.1467-8659.2009.01641.x>

Alex Vlachos. 2016. Advanced VR Rendering Performance. In *Game Developers Conference (GDC)*.