

ARM

Expanding Virtual Horizons – The Future of Mobile VR

Sam Martin – Senior Graphics Architect, ARM

If you've not watched someone experience Virtual Reality (VR) for the first time, I highly recommend it. Whether seven or seventy, the moment they get that headset on they all perform the same contortionist act, neck craning to peer around their virtual space like they're at the bottom of a well. Next comes the sharing dilemma, 'look at that!' 'You should see this!' It reminds me of the early days of the Walkman when we joyfully shouted in each other's faces, forgetting the music was only in our heads. This is a good snapshot of current VR: it's new, exciting and so totally absorbing you immediately want to share it. Indeed, Facebook's acquisition of Oculus suggests they see a significant potential for VR in social media. It's easy enough to imagine a time when it would be possible to sit around in a virtual cafe, chatting with lifelike avatars of friends and relatives, or browsing a feed of VR videos and 360 photos from your favorite social network, all while sitting on the sofa in your pajamas. Predictions of our VR future are everywhere. We shall have to wait to see what, among the myriad of crystal-ball gazing ideas on how we will use and consume VR, will eventually transpire. What is clear is that there are few other technologies that have the same potential reach and impact as VR, and its future cousin, Augmented Reality (AR). In this paper I will describe the market and challenges for VR as I see them today, along with the technology that enables it.

VR Gaming and Media

Gaming is seen as the first use case for VR, although it's one of a long list of potential uses. For games, VR offers the promise of another level of immersion, of new game mechanics, experiences and control mechanisms that simply do not exist in the traditional outlets. When mobile gaming first appeared there was much talk of porting console games directly to mobile. This largely failed to materialize. The differences in control mechanisms and form factor, the diversity and usage patterns of users and differing expectations of the mobile market demanded a re-think. The change from PC/console to mobile had a much larger impact on the games themselves than initially anticipated. Today's mobile games still share brands with PC and console titles, and on high-end devices we are able support the same graphical content, but the games industry has learnt how to adapt traditional gaming and has developed entirely different types of games for the mobile market. They have done so with some style as mobile gaming revenue is comparable to, and in some cases exceeding, PC/console revenue.

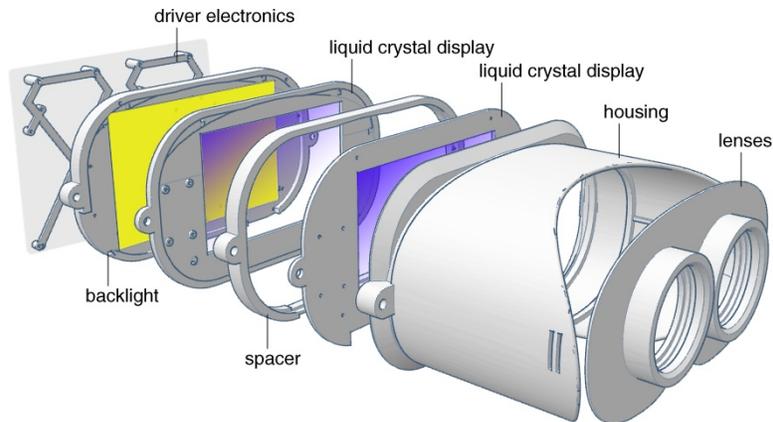
We are now on a new round of discovery with VR gaming, learning what does and doesn't work and exploring the unique opportunities it presents. It will be an exciting few years watching this content evolve. Early experiences suggest adjusting to the totally immersive nature of the technology while being shot at by aliens can be quite frightening and gentler forms of entertainment are more compelling than their descriptions might suggest at first glance. There will of course be many who will happily scare the living daylight out of themselves or seek out the world's biggest virtual rollercoaster, so I would not expect a one-size fits all result. The diversity in the Oculus store today suggests 'Experiences' and 'Concepts' will be new forms of entertainment in their own right. Add even the smallest element of interaction to one of these and you have a whole new style of 'game'.



nDreams® Perfect Beach VR Experience

It's hard to find a creative industry that isn't thinking about the potential impact of VR movies, photos and live broadcast. For broadcast media, it's an easy sell. Stand in the middle of the pitch at your favorite sporting event, attend major news events in person, sit beside a leading actor as he receives a long-awaited Oscar, or stand front and center at the inauguration of the next president of the US. Journalists, theatre companies, artists, musicians, educational centers and museums are already busy exploring the possibilities. Professional cameras for capturing 360° video and photos are already on the market, although standardized video and image containers and metadata have yet to mature and could do much to help. There is also considerable interest in consumer-oriented VR cameras, providing VR-enabled capture of the much-loved cat birthday videos, family holiday snaps and baby moments we routinely share on social media. "You'll do it, Beyoncé will do it", remarked Chris Cox, Facebook's Chief Product Officer, while speaking recently about sharing VR video. Both Facebook and You Tube already support 360° VR-ready video streams and projects such as Google Jump aim to push up the quality bar.

Not all VR video is captured in the traditional sense of live-action recording. Oculus Story Studio are exploring animated story telling through game engine technology; rendering rather than replaying the story. The idea of "VR Storytelling" has captured many people's imaginations. However, as those in the juncture between the film and games industry learnt rather painfully when first trying to incorporate their film making experiences of storytelling to the immersive experience of games, mixing any form of storytelling with interactivity can be an awkward combination. Even so, there are some sound technical reasons for considering rendering as an option. A back-of-an-envelope calculation will tell you that a stereoscopic high definition VR film could easily be an order of magnitude more demanding to stream and decode than regular HD video. If captured, compressed and streamed, much of the work would be wasted as roughly 5/6 of the 360° sphere would go largely unseen. By taking the users view direction into account you can do considerably better, which favors rendered content and would give any new form of view-adaptive video technology an edge. We may see a rise in GPU-centric video technologies, such as those under development by OTOY and 8i, or advances from a related and emerging field of light field cameras. It is a fast moving area and the industry is keen to keep up.



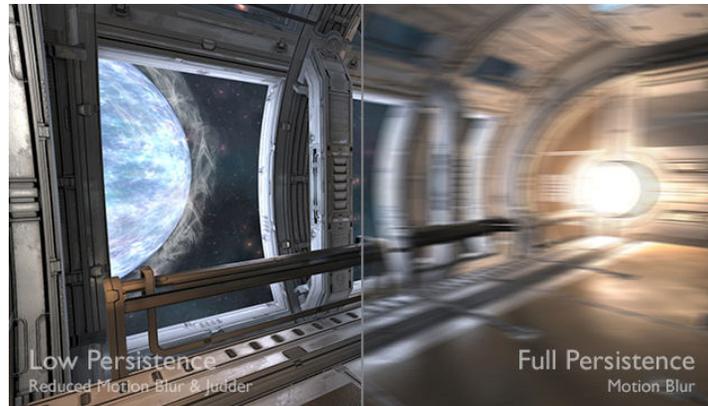
[Huang et al's Light Field Stereoscope](#)

Challenges for High Quality VR

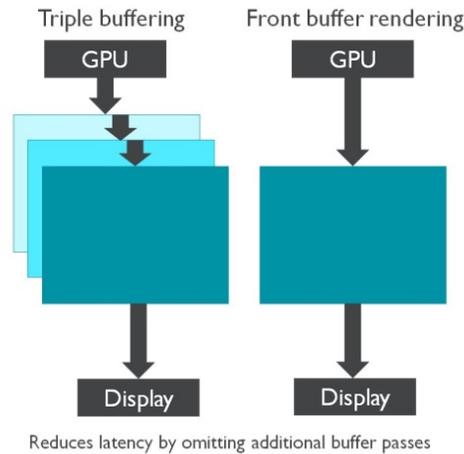
High quality mobile VR is demanding not only on GPUs but across the system. It requires close cooperation of many components, notably low-latency, zero-copy paths between video, GPU and display, as well as careful power management to ensure predictable timing within thermal limits, which is especially important with mobile VR. Linux is not a strict real-time operating system but the low-latency demands of VR require close to real-time behavior. Therefore, small changes in inter-thread communication and thread priorities, DVFS policy and other system integration choices can all affect the robustness of the critical path of “motion to photons”. Out of the box, an Android application may be triple buffered by the OS and have around 50 ms of latency just for the graphics submission path. This produces a very plausible result and is fine for short viewing experiences, but the lag between your head position and what you see is noticeable. After a while, this discrepancy can become uncomfortable and lead to nausea, even for rather sedate content. While still an impressive experience, the emphasis for the industry is to bring high quality VR to the mass market in order to unlock its broad potential.

The generally accepted maximum “motion to photons” latency is < 20 ms. Below this point the latency is unnoticeable. At the standard panel refresh rate of 60 frames per second (16 ms/frame) this is a tough target but is achieved by the Samsung GearVR devices through a combination of technologies, principally:

- The “asynchronous timewarp” process
- Rendering direct to the display frontbuffer
- Use of a low-persistence OLED display



Once per application-frame, the application will take the latest head tracking data and render the current view of the world into an 'eyebuffer' – a texture logically separated into halves for the left and right eye. Applications are encouraged to run at 60 FPS if possible, but this is hard to guarantee in practice and others may need to target 30 FPS. In order to provide a much tighter frame rate guarantee and reduce effective latency, the asynchronously running timewarp process is responsible for taking the latest eyebuffer available and writing it directly to the frontbuffer just head of the display reading it. The timewarp process uses a separate high priority OpenGL ES context to do this rendering, enabled by the 'context_priority' extension on ARM Mali GPUs. Work submitted through the high priority context will interrupt any existing application rendering on the GPU and run its own workload to completion before restoring the previous application rendering work. This allows the timewarp to submit work just in time for the display to read the data back out again. Since the timewarp runs closer to the display read time than the application rendering did, it also has more up-to-date head tracking data. It can use this to adjust the eyebuffer contents to better match the actual head orientation. It can only adjust for changes in head rotation, not position or changes in scene animations, but in practice this reduces the rotational latency below the noticeable threshold, providing a smooth and transparent experience. Drawing just-in-time to the frontbuffer in this manner best suits the low-persistence OLED displays. For LCD displays with a global backlight the display must illuminate the image at a single point in time. Low-persistence OLED displays are instead illuminated for a short period as the framebuffer is read out. When held horizontally and seen with a slow motion camera you see a wave of illumination moving left-to-right across the panel for each frame. This allows the head rotation to vary during the scan, further reducing latency.



The demands of VR do not end with latency. The magnification of the panel by the headset lenses highlights the need for high resolution panels. Even at a native resolution of 1440p, the illumination patterns on the panel are clearly visible. Driving displays larger than 1440p at 60+ FPS places considerable demand on the system and on the available external memory bandwidth. It becomes increasingly important that exchange of framebuffers between video, GPU and display is compressed. For this purpose, all ARM IP uses the lossless ARM Framebuffer Compression (AFBC) technology to minimize framebuffer memory traffic. It is also important that the application is able to efficiently render to multiple views in order to efficiently create the eyebuffer without unnecessarily duplicating work between left and right eyes. On platforms that expose it, Mali supports the 'multiview' OpenGL ES extensions that allow the application to submit the draw commands for a frame once to the driver and have the driver instantiate the necessary work for each eye. This greatly reduces the CPU time required in both the application and driver, and on Midgard Mali GPUs we further optimize the vertex processing work, running the parts of the vertex shader that do not depend upon the eye once and sharing the results between each eye. The appearance of Vulkan, the new cross platform API from Khronos, opens up further possibility for mobile VR with the potential to further reduce CPU overhead and put more power back in the hands of the developer.

VR is both ground breaking and fast moving, but thanks to the ubiquity and capability of modern mobile devices, it is not the exclusive preserve of expensive bespoke desktop or console systems. To truly realize its potential impact on our world, mobile VR must reach mass-market scales, and we are perfectly placed as thought leaders in the area, to enable this growth. Working closely with our partners, we will help deliver high quality VR experiences to maximize the potential of this technology and change your virtual world.