INTERNATIONAL ORGANISATION FOR STANDARDISATION
ORGANISATION INTERNATIONALE DE NORMALISATION
ISO/IEC JTC 1/SC 29/WG 11
CODING OF MOVING PICTURES AND AUDIO

ISO/IEC JTC 1/SC 29/WG 11 m41018
Torino, IT – July 2017

Source: **Jules Urbach (OTOY)**

Title: **An Introduction to ORBX**
Status: **Input contribution**

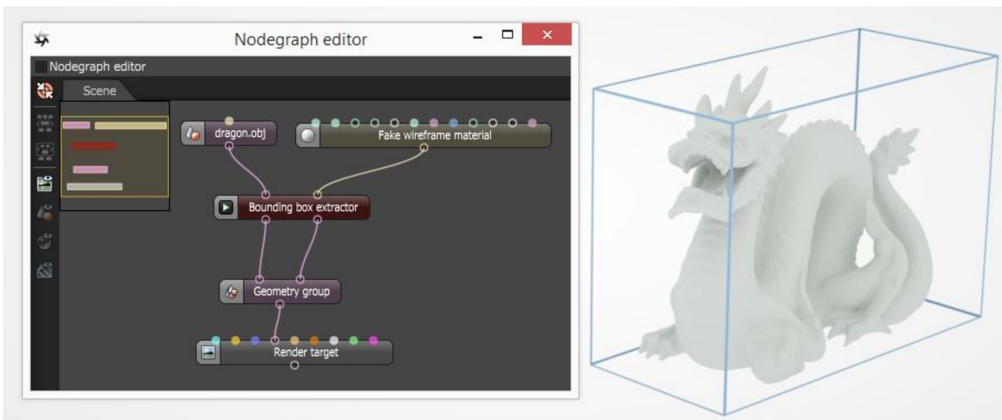

Introduction

ORBX was designed as a node based rendering schema, initially to cover OTOY's own requirements for storing and streaming holographic media (8D light field + depth) and interactive content.

In 2011 ORBX was refactored into an interoperable format for industry partners to leverage in their own software and services as a vendor and device agnostic framework for the above.

Notable 3rd party companies backing ORBX include: Autodesk, Mozilla (2013), Warner Bros. (2014), Oculus (2015), Disney, HBO, Discovery (2016), Facebook, Samsung and Unity (2017) [2-10].

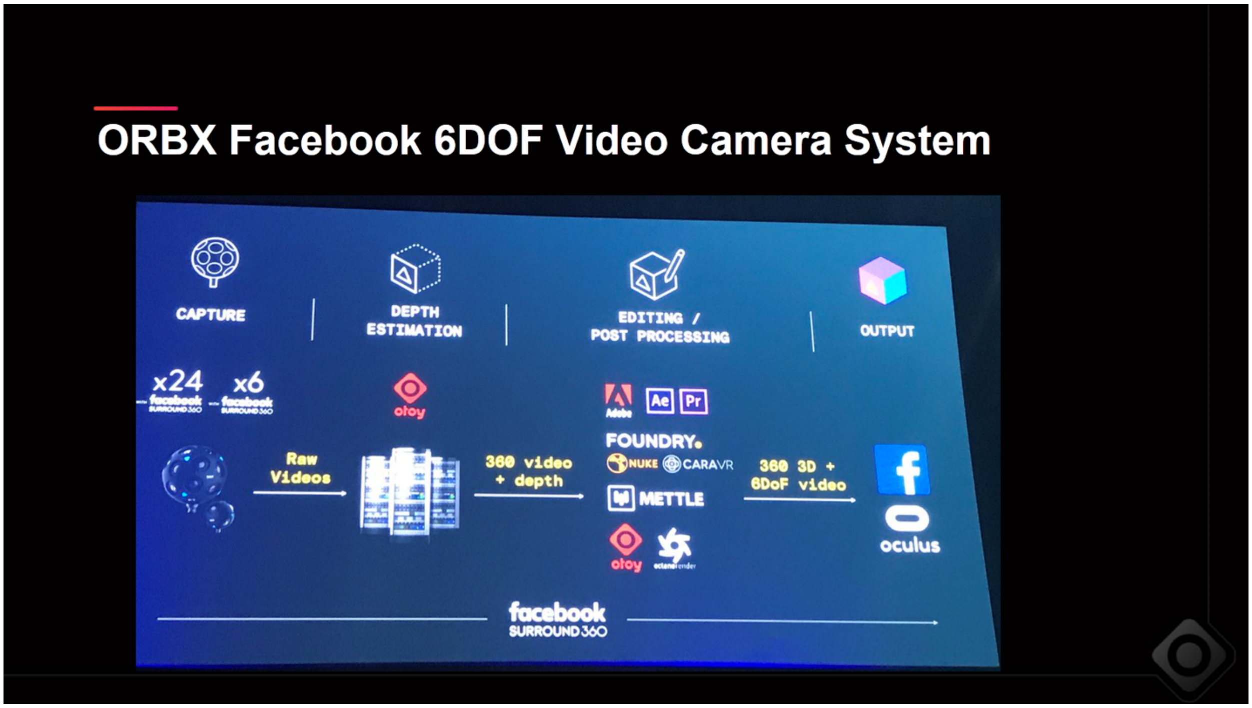

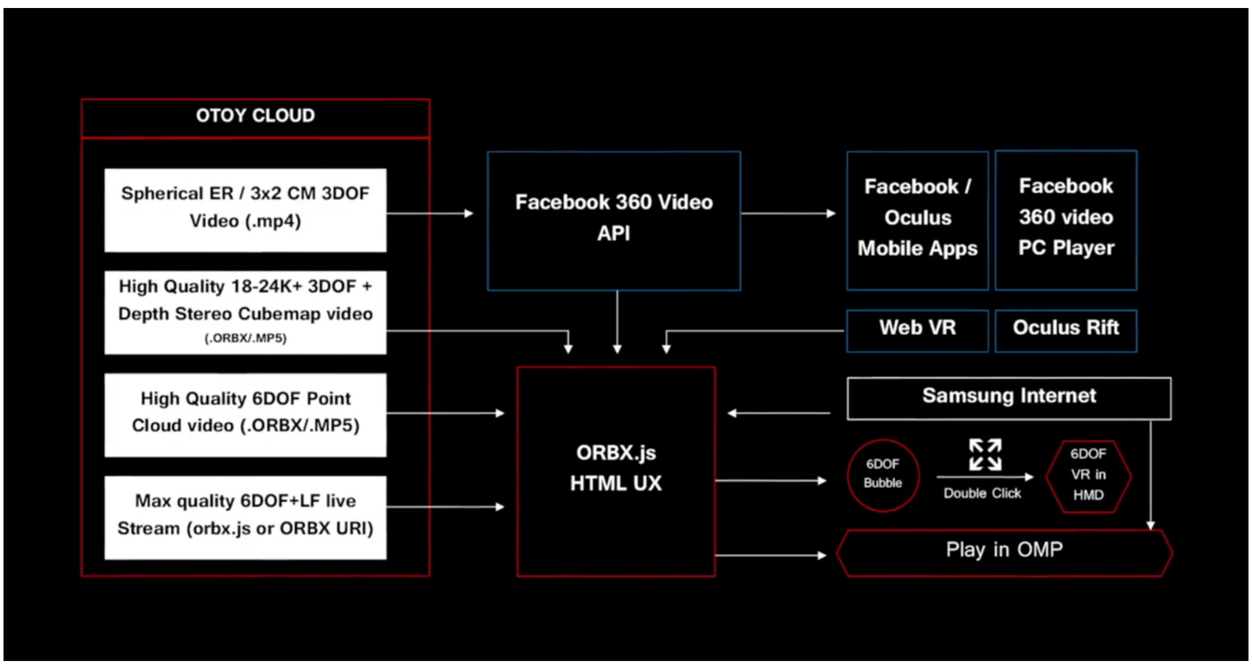

At the core of a serialized ORBX file is a simple virtual disk system and a single high level XML index file representing a root render graph to be evaluated, with optional inputs (camera view projection and position for example) and at least one output (usually a render target buffer with two or more dimensions).

The XML render graph schema in ORBX is a valid PBR node system, with all asset dependencies accounted for, in a form that can be evaluated by existing OSS in a complete, spectrally and physically correct rendering operation, with results in one or more outputs with the following formats and spatial-temporal dimensions:

- 2D Image/Video (thin lens, panoramic 3DOF mono/stereo, 6DOF 2D+depth deep layers)
- 3D Baking (PTC, OpenVDB, .FBX/glTF/.ABC)
- 4D-8D multidimensional ray bundles and buffers (light field/reflectance fields)

To this day, there is no single open format that encapsulates, in concept or in practice, the simple principle that ORBX is designed to achieve – a complete and full scheme for rendering a photorealistic scene on par with cinematic and architecture visualization requirements.

ORBX, by design, outside of the XML root graph and container, is an umbrella for numerous (and validated) open source sub-formats – EXR, Alembic, OSL, OpenVDB, glTF [20-24] that have proven themselves as industry standards covering important, albeit incomplete, portions of the render graph. OTOY is actively looking to new standards to add to the ORBX scheme as they emerge and provide more procedural depth to the render graph – including MaterialX (ILM/Foundry/Autodesk) [25], MDL (NVIDIA) [26] and USD (Pixar) [27].

We also wish to extend ORBX compute power further through portable domain specific native C99 compute modules (beyond Lua+FFI [28]) that can extend the principles of OSL surface, volume and displacement 'closures' with more procedurally computable scene graph extensions such as dynamics & physics evaluation, Inverse Kinematics, as well as domain and hull 'sharers' for arbitrary, but deterministic, procedurally generated scenes (e.g. to RenderMan DSO scene modules [29]). Possible virtualization systems for safely executing ORBX C99 modules in browsers and other JIT VM targets could leverage WebAssembly [30] and WebGPU [31].

A basic overview of the framework

As we look into the future of media beyond legacy video and image formats, and consider volumetric, immersive and free viewpoint mixed reality content, the concept of compressing pixels becomes less and less practical.

ORBX is intended to provide a scene graph framework. This is what we need beyond simple 6DOF (e.g. RGB+D), naïve light field and point cloud compression. The ultimate form of compression in a volumetric experience is simple – define a procedural way of rendering the scene – like a video game does – and trade data for compute. In the case of the famous Cornell box[11], a few lines of WebGL shader code and textual data on a phone browser can provide infinite and physically correct arbitrary viewpoints for such a scene[12][13]:

"The Cornell Program of Computer Graphics has become best known for its research on physically based rendering. We believe that computer graphics simulations will never become predictive of reality unless we correctly model the physics of light reflection and light energy propagation within physical environments."[14]

ORBX, which itself is based on PBRT (The Academy Award winning book on rendering) [15], extends this simple system to cover all possible scene types.

Given today's advances in GPU rendering on mobile devices (iPad 10.5 inch can rasterize 10 million triangles at 60 Hz while AR Kit is running), any scene that can be rendered locally and photo realistically with a free viewpoint should be represented procedurally where possible.

OTOY has also tested ray tracing ASIC from PowerVR which perform at 10x the ray tracing speed of current GPUs – which make procedural photorealistic rendering on a single mobile ASIC perfectly achievable [16]:

Further advances in TPU (i.e. Machine Learning ASICs) will make de-noising ray traced scenes trivial after less than 10 samples [17]:

A secondary goal of ORBX, besides providing a source description for rendering a scene procedurally, vs from geometry or voxelized buffers, is to serve as a target for computational photography.

Today, ORBX is the target format for all 8D reflectance field captures performed by OTOY's Light Stage service. Light Stage is more than a volumetric '4D' capture. It uses polarized light strobes to capture micro surface details at all light frequencies – including normals, specular lobes, scattering and more – all the things needed to provide a robust and near perfect CGI duplicate of the scanned asset in a spectrally correct PBR rendering.

In 2015 we performed a high density light field capture of 1728 views (16 Megapixels per view) which was submitted to both JPEG and MPEG in April 2017, as a reference scene for evaluating light field and volumetric codec efficiencies [18][19]. At MPEG 119, OTOY is presenting an updated representation of the scene - efficiently re-encoded as an ORBX scene (i.e. a fully CG asset with correct BRDFs for materials) that is compact enough to be fully navigated on an iPhone with AR Kit at 60 Hz, or an iPad at 120 Hz.

The need for a formal standard

ORBX continues to be adopted as an informal industry standard for 3D modeling applications, rendering platforms, and for creation and storage of VFX assets. We anticipate that current and future stakeholders who rely on ORBX would benefit from the formalization of ORBX as an MPEG format for the interchange of 3D photorealistic assets. Moreover, the formalization of ORBX as an MPEG format would help bootstrap the development of VR and AR services (e.g. by large scale commercial network providers) for these types of assets.

In support of creating such a new MPEG standard, OTOY is willing to contribute the ORBX specification under licensing terms consistent with a royalty-free strategy.

Supplemental material

We are providing an updated PDF (July 2017) describing the ORBX container schema. The latest and up to date description of the ORBX render graph schema can be viewed on http://docs.otoy.com/.

References

1. Jules Urbach, Arianne Hinds, "Sample data representation for real and synthetic content based on ORBX", M40667, April 2017, Hobart, AU.
2. https://www.roadtovr.com/beaming-the-matrix-into-your-eyes-otoy-ceo-on-the-future-of-real-time-lightfield-rendering/
3. https://venturebeat.com/2017/01/14/otoys-adds-lightfield-baking-and-other-tools-to-its-unity-integration/
4. https://uploadvr.com/facebook-otoy-volumetric-camera/
5. https://home.otoy.com/otoy-facebook-release-revolutionary-6dof-video-vr-camera-pipeline-tools/
6. https://home.otoy.com/making-vr-web-browsing-better/
7. https://home.otoy.com/hbo-and-discovery-invest-in-otoy-to-create-universal-publishing-platform/
8. https://home.otoy.com/otoy-joins-disney-accelerator/
9. http://variety.com/2014/digital/news/warner-bros-dc-entertainment-otoy-developing-batman-virtual-experience-1201324169/
10. https://blog.mozilla.org/blog/2013/11/05/mozilla-otoy-and-autodesk-work-to-deliver-high-performance-games-and-applications-on-the-web/
11. https://en.wikipedia.org/wiki/Cornell_box
12. http://www.graphics.cornell.edu/online/measurements/reflectance/index.html
13. http://madebyevan.com/webgl-path-tracing/
14. http://www.graphics.cornell.edu/online/box/
15. http://www.pbrt.org/
16. https://home.otoy.com/otoy-and-imagination-unveil-breakthrough-powervr-ray-tracing-platform/
17. https://www.facebook.com/groups/OctaneRender/permalink/806910102819403/
18. https://www.roadtovr.com/exclusive-otoy-shows-us-live-captured-light-fields-in-the-htc-vive/
19. https://uploadvr.com/otoy-light-field-demos/
20. http://www.openexr.com/
21. http://opensource.imageworks.com/?p=alembic

22. http://opensource.imageworks.com/?p=osl
23. http://www.openvdb.org/
24. https://en.wikipedia.org/wiki/GlTF
25. http://materialx.org/
26. http://www.nvidia.com/object/material-definition-language.html
27. https://graphics.pixar.com/usd/docs/index.html
28. https://luajit.org
29. https://renderman.pixar.com/resources/RenderMan_20/proceduralPrimitives.html
30. http://webassembly.org/
31. https://webkit.org/blog/7380/next-generation-3d-graphics-on-the-web/