Synthesising Panoramas for Non-Planar Displays: A Camera Array Workflow

Esan Mandal
DreamWorks
esan.mandal@dreamworks.com

Amy Kwa
DreamWorks
amy.kwa@dreamworks.com

Figure 1: Stitched and UV mapped panoramic image

ABSTRACT

In this talk we present a production workflow to generate panoramic high-resolution images for location-based entertainment and other semi-immersive visualization environments. Typically the display screens at these installations are an integral part of the surrounding architecture and have arbitrary non-planar surfaces. Our workflow is designed to minimize the distortions caused by the screen shape and optimize rendering of the high-resolution images while leveraging our existing feature film pipeline which, uses a standard perspective linear-projection camera model.

CCS CONCEPTS

• **Computing methodologies → Camera**; **Panorama**;

KEYWORDS

panorama, non-planar display, theme parks, camera, high-resolution

ACM Reference Format:
Esan Mandal and Amy Kwa. 2018.*Synthesising Panoramas for Non-Planar Displays: A Camera Array Workflow* . In *Proceedings of SIGGRAPH*. ACM, New York, NY, USA, 2 pages.

1 MOTIVATION

The workflow addresses the following challenges and constraints when generating such large format panoramic images:

Panoramic coverage : In a semi-immersive display the FOV typically goes beyond 90 deg, and may go all the way to 360 deg. A single super-wide perpsective projection camera for panoramic coverage results in substantial lens distortion.

Arbitrary screen shape : The non-planar arbitrarily shaped display inherently results in screen space distortion.

Irrelevant pixels : Standard perspective cameras produce a rectangular image, which means that we have to over-render to fit the arbitrary screen shape within the rectangular bounds of the camera format. This over-rendering results in numerous irrelevant pixels being rendered that do not get displayed.

High-resolution : The total number of pixels rendered for this short running high resolution content becomes comparable to that of a feature film, necessitating optimization.

Feature Pipeline : We want to use our existing feature pipeline so as to leverage all the existing assets and take advantage of artist's proficiency and expertise of our tools.

2 BACKGROUND

The visualization environments typically use a multi-projector system to achieve the coverage of large area displays. The system also abstracts the configuration of the projectors so that they are inpdependant from how the images are created via an intermediary projection mapping. The rendered images are provided as a UV map of the screen surface, which translates between the coordinate space of the rendered images and the coordinate space of the projectors perspective. The abstraction allows complete creative freedom of the rendering camera's layout while the physical projectors can be positioned as best suited to the physical optics and architecture.

The projection system software generates a 3D scene that accurately replicates the screen architecture and the physical projectors as 3D cameras. In this 3D scene the content images are UV mapped onto the digital screen as a textured image sequence. The software renders the textured screen from each camera, and feeds them into the corresponding projectors to be displayed on the actual physical display screen. This process aids in mitigating seam and overlap issues inherent to multi-projector systems [Raskar et al. 2018].

3 METHODOLOGY

Our approach is built on a specialized camera rig and a pipeline for post-processing the rendered images into screen space UV maps.

3.1 Camera Rig

The camera rig is designed to provide the desired coverage and allow complete freedom for creative direction. After analyzing single and camera array rigs we found that, depending on the screen shape, the single camera renders can have 5 times more irrelevant pixels. Moreover, to achieve the same fidelity the single camera image required 6 times the resolution. (Fig. 2). To illustrate our workflow we will use the specifications of a successfully completed project.

The screen An accurate screen specification and a digital 3D model is obtained since the camera rig is custom built based on the shape and size of the screen.

© DreamWorks Animation, L.L.C. 2018. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive version was published in SIGGRAPH 2018 Talks,
https://doi.org/10.1145/3214745.3214804.

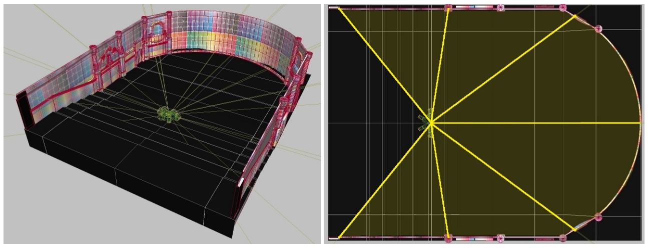

Figure 2: The screen architecture and camera layout.

Camera position The camera is positioned at a predetermined "sweet spot", which represents an idealized viewing point where all parallax and perspective illusions sync up and result in the best viewing experience.

The camera array Multiple nodal cameras with overlapping frustums and approximately equal FOV [Kopf et al. 2007], are positioned at the "sweet spot" to cover the whole screen as tightly as possible. We used six cameras for this project.

Camera FOV and orientation The focal length and orientation of the cameras are adjusted to best achieve the full coverage of the screen surface. The pixels rendered for each degree of FOV needs to be consistent so as to ensure that the pixel resolution is evenly distributed across the panorama.

Camera format The camera format is based on the desired resolution of the relevant pixels that gets displayed on the screen. To obtain a vertical image resolution of 2048px the cameras vertical resolution was 2190px so that it fully accommodates the curved screen shape. The horizontal resolution at 3900px was the full horizontal panoramic coverage divided by the number of cameras.

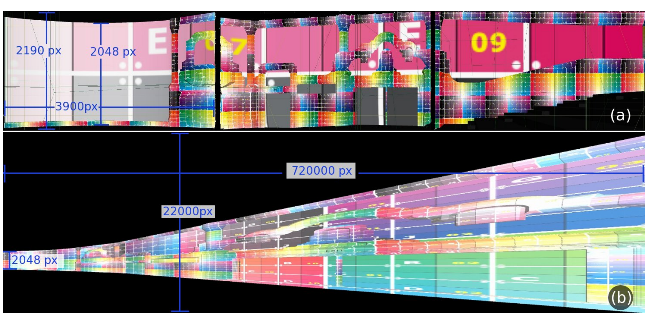

Figure 3: a) Camera format renders of the rig b) For comparision, camera format render from a single camera.

Film Back Offsets Since the cameras are nodal and the position is locked, the "film back offset" attribute is used to adjust each individual camera's framing optimally. We developed custom translators to convert the film-back offset between packages and integrate it into our pipeline.

The Reticle The extra coverage of cameras mean that, just by looking through the camera view, an artist cannot tell if a character action is going to fall within the screen space or not. We created a "Camera Reticle" that provided the artist with a geometry mask parented to the camera that occludes the irrelevant pixels from the POV of the camera array.

3.2 Image Post-Processing

The post-processing workflow assembles the individual images rendered using the camera array rig [Szeliski 2004] and generates the UV map in screen space. The screen geometry and the camera rig are accurately recreated in Nuke. Each of the rendered images are projected through their corresponding cameras onto the screen geometry and the projected image on the screen geometry is scanline rendered into a UV map. The alignment of the render camera with the idealized audience viewing position and the projection mapping of the images onto the screen geometry, minimizes the distortion introduced due to the shape of the screen surface. This setup was implemented as a Nuke gizmo and the post-processing was added to our rendering workflow as an automated process.

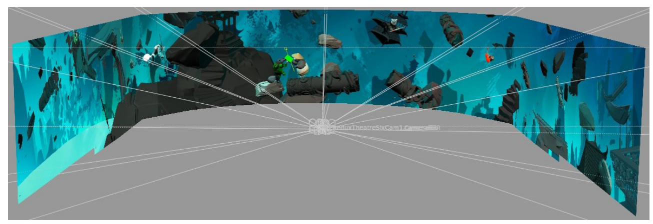

Figure 4: The 3D scene setup in Nuke with the rendered images project onto the screen.

4 RESULTS AND FUTURE WORK

The camera rig and the post-processing workflow was implemented with minor additions to our pipeline which enabled artists to use legacy characters and assets. We have delivered high-frame rate content as large as 22800X2048px for various entertainment installations and theme park rides across the world. Our content has been displayed on screens of various shapes and sizes at Motiongate Dubai, Shrek's Adventure(London), Universal Hollywood, ACMI Exhibition, DreamPlace and other locations.

In the future, as we move to a Ray-tracer based pipeline we can take advantage of the felxibility of the different camera projection models (spherical, cylindrical etc.) which, will facilitate the use of a single Ray-traced camera that won't have the drawbacks of a single perspective camera. A single camera rig will simplify the whole production significantly. We are also exploring Ray-traced cameras that take as input any arbitrary geometry as a film back. Mapping the shape of the screen to the camera's film-back geometry will ensure that only relevant pixels are rendered.

REFERENCES

Johannes Kopf, Matt Uyttendaele, Oliver Deussen, and Michael F. Cohen. 2007. Capturing and Viewing Gigapixel Images. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2007)* 26, 3 (2007).

Ramesh Raskar, Jeroen Baar, Srinivas Rao, and Thomas Willwacher. 2018. Multi-Projector Imagery on Curved Surfaces. (01 2018).

Rick Szeliski. 2004. *Image Alignment and Stitching: A Tutorial.* Technical Report. 89 pages. https://www.microsoft.com/en-us/research/publication/image-alignment-and-stitching-a-tutorial/