

Rendering *Fast & Furious: Supercharged* in Stereo for a 400ft Curved Screen

Rob Pieké*
MPC



Figure 1: Example frame from the final delivery, rendered as a single 27544 x 2160 image - ©2015 Universal. All rights reserved.

Abstract

For *Fast & Furious: Supercharged*, we were required to produce a 90-second shot which would be projected in stereo on an enormous U-shaped screen surrounding an audience in a series of long tram cars at Universal Studios Hollywood.

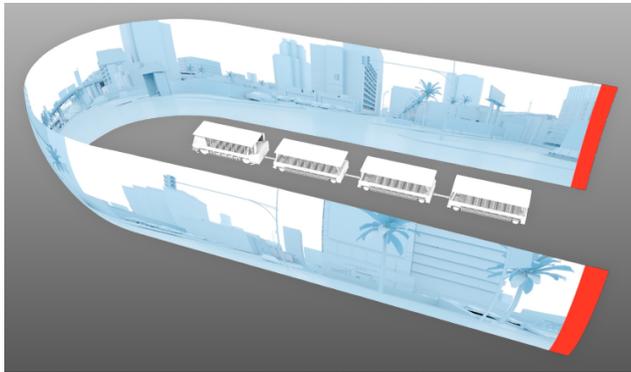


Figure 2: Mock-up of the physical build at the theme park - ©2015 Universal. All rights reserved.

This raised three main technical challenges: anticipating the projection warping and inverting it at render-time, artistically recomposing the perceived depth of the scene to provide a better audience experience, and ensuring a consistent stereoscopic effect.

To overcome these challenges, we developed a custom camera representation which allowed the camera to change its location and orientation over the domain of the rendered image, independent of any temporally-based motion for blurring. This further necessitated modifications to both our animation and rendering workflows, as well as a new validation process to emulate the viewing experience.

Keywords: camera projection, theme park rides, pipeline

Concepts: •Computing methodologies → Computer graphics; Rendering;

*e-mail: rob-p@moving-picture.com

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). © 2016 Copyright held by the owner/author(s). SIGGRAPH '16 Talks, July 24-28, 2016, Anaheim, CA, ISBN: 978-1-4503-4282-7/16/07 DOI: <http://dx.doi.org/10.1145/2897839.2927395>

1 Screen-Driven “warpCam”

Due to the non-planar nature of the screen, we needed a way to correlate the pixels in our rendered images with physical locations when projected.

We created a high-contrast calibration image and had it projected onto the screen during a laser scanning session. With the calibration features preserved in the scan result, we were able to measure these points and create a highly detailed mapping between the UV space of our rendered images and the 3D position of the actual screen. This process is shown in the figure below.

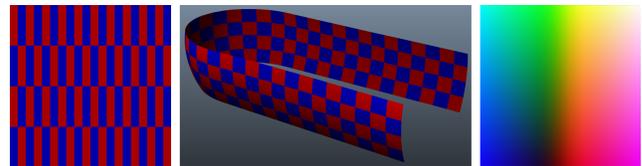


Figure 3: From left to right: calibration image; laser scan; UV/XYZ correlation map - ©2015 Universal. All rights reserved.

Using this map, we extended RenderMan and mental ray with a custom camera we referred to as “warpCam”. The standard perspective projection was bypassed and, instead, the direction for each pixel’s camera ray was driven from our measured data. This resulted in images which appeared heavily warped when viewed directly on a computer monitor, but which appeared straight and proper when projected on the final screen.

2 Reducing Perspective

For many parts of the screen, we knew some of the audience would be looking forward at it while some would need to turn to look backward, resulting in different expectations for whether they should see the back or front of an object at that location.

We reduced this effect by pulling the virtual camera away from the screen, to create a more orthographic projection. Effectively considering the screen as three sections: the left side, the right side, and the nose-cone, we modified the warpCam to pull away from the two side sections, and blend between these offset locations for the nose-cone – rendering every pixel from a subtly-different viewpoint.

Our animators required a compelling preview as to where their work would appear on-screen, so we reproduced the warpCam inside Maya. As the animators rotated the camera to look at different parts of the screen, the camera’s origin would automatically translate to match the location it would have during final rendering.

Pulling the camera backwards resulted in “unsafe areas” where the camera could end up inside objects which the audience was meant

to perceive as being behind the camera, or where objects could appear on multiple parts of the screen at the same time. To compensate for this, we offset the origins of the camera rays along their viewing direction, both automatically to minimize the unsafe area, and with artist control in case we started clipping visually-important features.

3 Omnidirectional Stereo

Using a standard perspective projection results in a loss of stereoscopic effect when looking directly to the sides which, on our curved screen, is where the bulk of our content was directed to be. This challenge is similarly described by [Simon and Beckhaus 2003], who suggest rendering a number of image “slices”, rotating the stereo camera pair around a nodal point to capture each slice.

We extended our warpCam model to include horizontal offsets. After first translating and orienting our camera’s nodal point for each pixel in the image being rendered (using the methods mentioned in the previous sections) we further applied locally-horizontal offsets for each eye, resulting in a globally-compelling stereo experience.

4 Validation

Assessing our rendered images was often not practical to do directly, as the content at this point was significantly warped (again, so as to later appear straight and true when projected). We built a virtual version of the \cap -shaped screen in Nuke and used the data driving our warpCam to map the rendered images to this geometry. We were then able to generate a second set of images for review, representing the arbitrary vantage points of the audience in the tram cars looking at different parts of the screen.

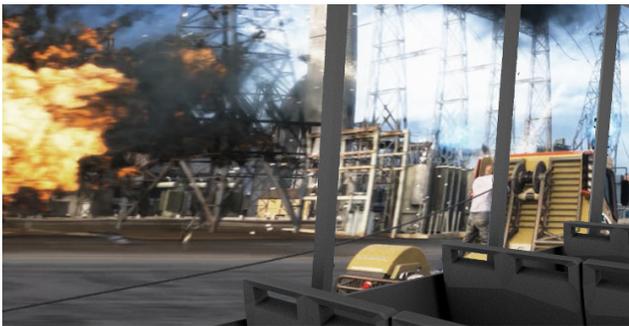


Figure 4: *Emulating the viewer experience - ©2015 Universal. All rights reserved.*

Through these “experience-camera” renders, we could assess both how much of the action would be occluded by other viewers and how much distortion riders would be subjected to depending on where they sat. This information was directly fed back into further creative discussion and composition.

The validation images also identified to us where naïve image translation (to adjust convergence and perceived depth) was introducing vertical disparity. We compensated for this effect with further image warping, using our knowledge about the relationship between our image and the physical screen.

Acknowledgments

A special thank you must be given to Thierry Coup and Chris Shaw for their creative vision, and to Diego Guerrero, Alessandro Sabioni and Martin Riedel for their relentless testing of the technology we built for this project.

References

SIMON, A., AND BECKHAUS, S. 2003. Omnidirectional stereo surround for panoramic virtual environments. In *ACM SIGGRAPH 2003 Sketches & Applications*, ACM, New York, NY, USA, SIGGRAPH ’03, 1–1.