The Stereoscopic Conversion Pipeline for John Carter

Scott Willman Gregory Keech John Groteluesched Michele Sciolette
Cinesite

1. Introduction

We present the stereoscopic conversion technology developed at Cinesite to convert over 1500 shots for Disney's John Carter.

Traditional stereoscopic conversion techniques can be broadly divided into two main categories: disparity grading techniques, where disparity maps are manually generated by skilled artists [Tucker-Fico at al, 2011] and geometry projection techniques where images are projected onto 3d geometry and rendered from a different point of view.

We developed a hybrid approach that allowed us to have the flexibility of the disparity grading techniques with the consistency and spatial coherency of a 3d geometry-based approach: we use accurate 3d geometry and camera tracking data to generate spatially correct depth maps that provide the basis for the image based part of our pipeline.

2. Implementation

The first stage of our pipeline is to generate accurate depth maps for each frame.

For live action elements this required two separate processes for each shot: a comprehensive camera and object tracking (rotomation) and a detailed traditional 2D rotoscoping of all elements. Efficiently managing these two independent processes was a key focus of our initial development effort and we created a large set of tools to streamline the process and ensure consistent outputs. Work on each shot started with our artists using an interactive tool to define which parts of the frame needed to be isolated and export these shot requirements into a standard format. Each artist in the rotomation and rotoscoping teams received a data file to help them visually and automatically check that their work matched the shot requirements. Being able to ensure consistency across these tasks allowed us to automate the setup for the following stages and played a key part in making the process scale to the required large volume of shots.

Using a 3D rendering system, we projected the 2D rotoscoping mattes onto each corresponding 3D object and rendered a depth map from the point of view of the original camera. The resulting object depth map often contained small errors or holes that we would fix using traditional image processing techniques.

For CG generated elements, depth maps were directly available from the rendering system, though we had to ensure that all scaling of depth maps was accurate so that all components of the final depth map were consistent with each other.

The rest of the pipeline is entirely based on two-dimensional image processing of the depth maps and image content.

One key part of the process consists of accurately and efficiently warping the images based on the depth maps. For a desired stereoscopic interaxial distance, we compute a disparity map so that for each pixel in the original view we know its corresponding location in the new view. This data cannot be used directly into a standard backward-warping tool, which causes invalid pixels to overwrite valid ones and trailing or leading edges to disappear, depending on which eye is being recreated.

To address this problem, we developed a proprietary depth aware forward-warping algorithm that generates a new 'warped' disparity map that can be used by a backward-warping tool available in a standard compositing package. As each pixel in the new view can be mapped to several source pixels in the original view, we use depth sorting to determine the valid source pixel choice. Given a requirement that our stereo images be free of vertical disparity, we were able to heavily optimize our code relying on the fact that we only needed to shift data horizontally.

One of the most difficult and time-consuming parts of the stereoscopic conversion process is the painting of occluded regions. We developed an occlusion mapping tool to identify all regions requiring paint and an automated depth aware inpainting technique that was extremely effective in dealing with small occlusions. In cases where the automated tool was unsuccessful, the occlusion maps allowed a paint artist to determine the exact size and location of all occluded areas, thus greatly reducing the amount of manual paint work required.

3. Conclusions

Our hybrid solution to the conversion problem proved to deal effectively with a variety of shot types and production workflows. Artists were able to quickly address all requests from the stereographers and generate a final stereo output of the highest possible quality.

Future developments of our stereo conversion process would be to use deep compositing to simplify the interaction with CG assets and allow volumetric renders to be placed within the live action environment more accurately.

References

Katie Tucker-Fico et al . 2011. Design and realization of stereoscopic 3D for Disney classics. In ACM SIGGRAPH 2011 Talks (SIGGRAPH '11). ACM, New York, NY, USA, , Article 12 , 1 pages.

SIGGRAPH 2012, Los Angeles, California, August 5 – 9, 2012.
ISBN 978-1-4503-1435-0/12/0008