

Efficient Hybrid Volume and Texture Based Clouds

Laura Murphy
Pixar Animation Studios
lmurphy@pixar.com

Martin Sebastian Senn
Pixar Animation Studios
senn@pixar.com

Matthew Webb
Pixar Animation Studios
webb@pixar.com



Figure 1: Clouds in *Cars 3* and in *Incredibles 2*. ©Disney/Pixar

ABSTRACT

We developed workflows for *Cars 3* and *Incredibles 2* that maintain most benefits of volumetric clouds while leveraging the render speed and artistic control of traditional texture based approaches. For *Cars 3*, we rendered volumetric clouds using Houdini and RIS to create layered, reference plates that matte painters used to quickly paint sky texture maps. The Houdini based workflow let us intuitively scout sequences and set up paintings. We detail artist-friendly tools in Photoshop to allow for intuitive composition adjustments. On *Incredibles 2*, we replaced the Photoshop portion of the pipeline with Nuke and rendered three passes: key, fill, and world position. Standardizing the workflow around world aligned, high resolution, hemispherical textures greatly streamlined the process and provided downstream lighting department with greater control integrating the sky with the foreground.

CCS CONCEPTS

• Computing methodologies → Rendering; Volumetric models; Image processing;

KEYWORDS

clouds, sky, matte paint, houdini, compositing, volumes

ACM Reference Format:

Laura Murphy, Martin Sebastian Senn, and Matthew Webb. 2018. Efficient Hybrid Volume and Texture Based Clouds. In *Proceedings of SIGGRAPH '18 Talks*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3214745.3214797>

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).

SIGGRAPH '18 Talks, August 12-16, 2018, Vancouver, BC, Canada

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5820-0/18/08.

<https://doi.org/10.1145/3214745.3214797>

1 INTRODUCTION

For *The Good Dinosaur*, we created volumetric clouds for all the skies [Wrenninge 2015] [Bakshi et al. 2015] [Webb et al. 2016]. This workflow had several advantages over matte painted skies: aesthetic match with foreground elems, lightable by downstream lighting department, and proper perspective and parallax. The primary downside was significantly increased render times. We wanted to keep the look of rendered clouds, but with the render speed of the traditional matte painted skies. We also had upcoming show with look requirements (car shaped clouds) that would be difficult to achieve with strictly volume based skies. These motivated us to develop a pipeline where we could reintroduce the traditional texture based stages to the workflow.

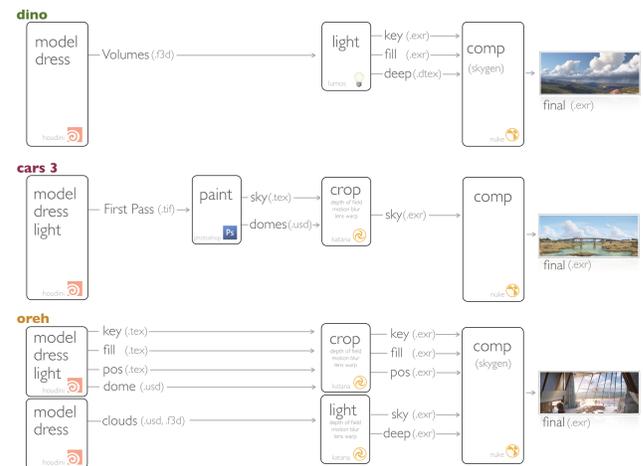


Figure 2: Workflows from *The Good Dinosaur*, *Cars 3*, and *Incredibles 2*. ©Disney/Pixar

2 CARS 3

For *Cars 3*, our workflow placed volumetric cloud renders into the hands of digital matte painters using Photoshop. Integrating our tools for setting up Photoshop matte paintings into Houdini allowed us to explore setups in 3D space with access to cameras and geometry across multiple shots as well as combining a 2D pipeline with our volume based toolset. To that end we dressed and lit clouds volumes in Houdini and rendered them into a layered series of panoramic reference plates using RIS. Starting with these plates matte painters were able to efficiently execute aesthetic adjustments down to the pixel level and add any details needed to complete the sky as a whole composition. Specifically painters could modify cloud shapes in ways difficult to achieve with a strictly volumetric representation, such as sky parts that had to look like cars. The individual Photoshop layers were then exported into separate textures and projected onto domes or cards, rendered from the shot camera in order to get correct RIS motion blur, depth of field as well as lens distortion. Layered domes and cards allowed for parallax between clouds but at a fraction of the cost of a full per frame volumetric cloud render.

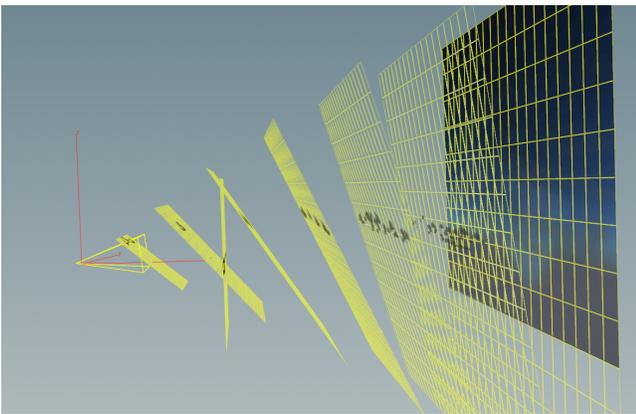


Figure 3: Sky texture layers projected onto cards, previewed in Houdini. ©Disney/Pixar

3 INCREDIBLES 2

Incredibles 2 had different priorities for its skies than *Cars 3*. We needed to make the skies HDR (high dynamic range) friendly and place more control over the sky into the hands of the lighters. Additionally, everything needed to land earlier since lighting was starting at the same time as animation.

We decided on a standardized dome texture workflow. We have one default sky dome for all sky shots, and for each sky setup we render three hemispherical lighting AOV textures (key lights, fill lights, and world position) from Houdini with RIS. To speed up turn around time of the typically 8k or 16k dome render and to mitigate losses caused by occasional render errors, we split the spherical render into crop tiles that could each converge in several hours. The resulting tiles are stitched into the full dome textures.

The three textures are rendered on the skydome as three render layers from shot cameras in Katana, which also applies the

motion blur, depth of field, and lens distortion. Skydome height is scaled down to increase perspective distortion of the texture which approximates a ceiling of clouds.

The rendered sky elements feed into Nuke, where lighters treat the compositing network as a live matte painting to more easily hit lighting/color notes. In Nuke our lighters fine-tune the colors, rebalance the key/fill passes, and simulate atmosphere with Pixar's SkyGen. If the volumes need any composition or lighting updates to the dome textures, we make our changes in Houdini and regenerate the dome textures. This workflow also allows us to maintain overrange values in clouds to ensure our skies are ready for HDR color grading.

For shots where the single dome sky approximation would not suffice (when clouds needed true parallax or interacted with set geometry), we promoted the clouds to full volumes via USD sublayers. Similar to *The Good Dinosaur*, these clouds were exported from Houdini and rendered directly as volumes for each frame [Webb et al. 2016].

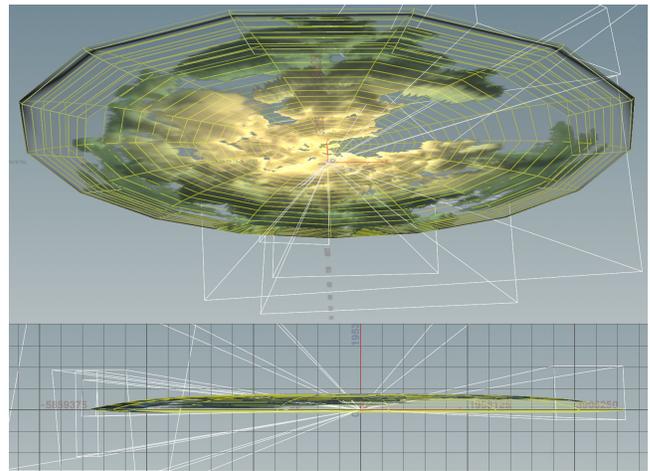


Figure 4: Default squashed skydome with a fill layer texture, previewed in Houdini. ©Disney/Pixar

REFERENCES

- Sanjay Bakshi, David Munier, Matt Webb, and Ana Lacaze. 2015. Disney/Pixar's The Good Dinosaur. ACM SIGGRAPH ASIA 2015 Computer Animation Festival - Production Talks & Panels. (November 2015).
- Matthew Webb, Magnus Wrenninge, Jordan Rempel, and Cody Harrington. 2016. Making a Dinosaur Seem Small: Cloudscapes in The Good Dinosaur. In *ACM SIGGRAPH 2016 Talks (SIGGRAPH '16)*. ACM, New York, NY, USA, Article 64, 1 pages. <https://doi.org/10.1145/2897839.2927405>
- Magnus Wrenninge. 2015. Art-directable Multiple Volumetric Scattering. In *ACM SIGGRAPH 2015 Talks (SIGGRAPH '15)*. ACM, New York, NY, USA, Article 24, 1 pages. <https://doi.org/10.1145/2775280.2792512>