Populating the World of Kumandra: Animation at Scale for Disney's "Raya and the Last Dragon"

Alberto Luceño Ros
Walt Disney Animation Studios
Burbank, USA

Kristin Chow
Walt Disney Animation Studios
Burbank, USA

Jack Geckler
Walt Disney Animation Studios
Burbank, USA

Norman Moses Joseph
Walt Disney Animation Studios
Burbank, USA

Nicolas Nghiem
Walt Disney Animation Studios
Burbank, USA

Figure 1: Example crowds: large scale, close up, dragons.

ABSTRACT

Walt Disney Animation Studios' 59th film "Raya and the Last Dragon" takes place in the fantasy world of Kumandra. We look at the challenges of casting and choreographing diverse people, creatures, and props to bring a varied spectrum of cultures and scenes to life: tense gatherings, intimate palace interiors, bustling floating markets with integrated boat traffic, panicked crowds, everyday life, and magical movements of dragons. The sheer diversity of characters presented a new set of obstacles that inspired us to extend our crowd system to efficiently address our increasingly diverse needs. Using production examples and results, we look at how our existing workflows and pipeline were leveraged and augmented to support these efforts. We discuss solutions for art-directed, simulated, and procedural approaches using our in-house Houdini-based system called Skeleton Library [El-Ali et al. 2016].

CCS CONCEPTS

• **Computing methodologies** → **Procedural animation**.

KEYWORDS

crowd, crowd pipeline

ACM Reference Format:
Alberto Luceño Ros, Kristin Chow, Jack Geckler, Norman Moses Joseph, and Nicolas Nghiem. 2021. Populating the World of Kumandra: Animation at Scale for Disney's "Raya and the Last Dragon". In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH '21 Talks), August 09-13, 2021.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3450623.3464648

SIGGRAPH '21 Talks, August 09-13, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-8373-8/21/08.
https://doi.org/10.1145/3450623.3464648

1 INTRODUCTION

In "Raya and the Last Dragon", crowd characters are essential to bringing the world of Kumandra to life. The film brought many challenges to the Crowds team, prompting us to revisit our pipeline and expand on existing workflows and procedures.

2 WORKFLOW

2.1 Data pipeline

As shown in Figure 2, Crowds consumes character and environment assets in the form of crowd bind assets from the Character team and animation clips or cycles from Animation. Crowds artists work primarily in Houdini, which is integrated with our in-house particle-based animation instancing and editing system called Skeleton Library; this system scales to large numbers of characters at interactive speeds. Using our custom instancer, Crowds artists publish completed crowds elements that include Look and can be rendered in Hyperion, Disney's in-house path tracer. From there, the Technical Animation department selects and promotes characters that need cloth and hair simulation, saving the characters into a new enhanced element. Characters can be "promoted" to the Animation department to enhance their performance. Animation also

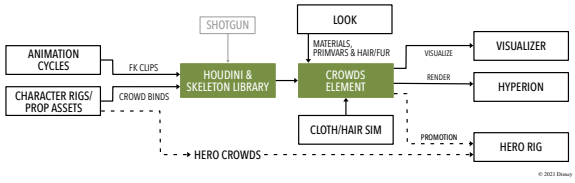

Figure 2: Data Pipeline

authors hero crowd characters which do not go through the Crowds department at all.

2.2 Asset management

Kumandra is divided into 5 lands, each named after a part of the dragon: Heart, Fang, Spine, Talon and Tail, and each with their own distinct cultures and idiosyncrasies. This presented new challenges for the Character and Crowds teams in designing how data is structured, managed, and consumed. Our goal was to visually differentiate characters from each land, while maximizing efficiency and asset utilization. Our Shotgun-based labelling and tagging system allows crowd placement tools to filter based on Shotgun tags, providing a powerful mechanism for casting assets in shots. Creating a new label to distinguish each land's permitted geometric variations enabled shared assets across lands while ensuring that design aspects were preserved. The most effective method for distinguishing lands came from Look variations on outfits. Our primvar system assigns values to each mesh of the instanced assets, driven by user-specified particle attributes. Primvar values are typically used to control variations in textures and color swatches. We introduced a new primvar that acts as a switch in the materials and drives the selection of textures and palettes for each land.

2.3 Animation Refitting

Marrying animation clips with character assets presents challenges when deploying large numbers of assets. The process of transferring animation data authored on one asset to others can be time consuming and unreliable. Our character asset hierarchy is designed to maximize efficiency of asset building and animation clip authoring and propagation. This is made possible by minimizing rig structure variation across geometric variants of assets. Our refitting process propagates data along the asset hierarchy, which consists of three levels of characters. Animation clips are authored and reviewed on "archetypes", then automatically transferred onto "silhouettes" which represent different body types. After the results are reviewed and edited, the animation is directly inherited by "leaf" level characters which are ultimately deployed in shots. These leaf assets share rig structure with their respective silhouettes, guaranteeing compatibility. This system allows a large number of leaf level geometric variations (in the thousands) to consume a manageable number of animation variations for each clip without any loss of quality.

2.4 Live Effects Data Pipeline

Kumandran dragons run in the air by magically stepping on raindrops. The Effects Department was tasked with creating the visual magic effect of foot impact on rain. This required transfer of data between the Crowds and Effects departments. Since there are several pipeline steps separating the two departments, a baked geometric representation would no longer be valid by the time it reaches Effects. Since both departments are Houdini-centric, we could leverage two HDAs for Effects to generate the data live. The first reads relevant data from the crowds element while the second saves the trail of each dragon's foot as a curve with attributes for detecting foot impact. Thanks to that method, Effects artists did not rely on baked data and could get updates without sending the task back to Crowds to recompute. Additionally, it allowed for continued iteration on the trail generation and foot impact detection algorithms.

3 SHOT DESIGN AND EXECUTION

For shot work, we leverage the power of Houdini which coexists with custom pipeline and Skeleton Library tools. The output data is always in the same format, providing compatibility and allowing artists to use any combination of approaches for a task.

3.1 Art-directed

Artists often rely on a heavily art-directed approach to meet the specific needs of each shot. Our system provides both procedural and hand-crafted approaches for placement and casting of characters. On top of that, there are tools to alter the character performance through procedural or manual processes such as Look-Ats, FK offsets, and IK constraints. This enables artists to shape crowds easily and maintain a high level of control and precision in creating vignettes, without having to promote characters to the Animation Department. This workflow is very efficient for small to midsize and close-up crowds where high quality is required and compositional considerations are critical.

3.2 Procedural

For more complex shots and large scale animation, procedural methods are often instrumental. One challenging sequence required us to choreograph dragons running on raindrops, prompting the development of a tool that given a curve, instances and animates a group of dragons along it. The tool provides high level controls to artists to manage aspects of the flock such as the overall shape, the amount of noise, and speed. For another sequence that takes place in a floating market, we developed a boat system that automatically populates scenes with lanes of boats along curves and fills them with props and passengers. The system handles primary and procedural secondary motion: it adds rocking motion to boats and passengers and synchronizes the rowing cycle of characters.

3.3 Simulation

Finally, we leverage simulation for some performances. For sequences including large groups of fish and beetles, we used methods including Boids, Position Based Dynamics, and Rigid Body Dynamics. Resulting simulated data is compatible with the rest of our system and can be further enhanced using any of the art-directed and procedural approaches discussed above. This ensures that artists can address artistic or story notes efficiently and gives them the controls necessary to produce the most compelling results.

4 CONCLUSION

Two thirds of the film's sequences included crowd characters in some form. A total of 259 shots feature characters animated using this enhanced workflow, bringing the final asset deployment count to 68814 after cleanup and culling. A total of 1035 animation clips were deployed, generated from 447 unique archetype level clips. 5138 geometric variations of assets were used. Overall, our system provided an efficient and flexible framework to handle the large scale of animation required for the film.

REFERENCES

Moe El-Ali, Le Tong, Josh Richards, Tuan Nguyen, Alberto Luceno Ros, and Norman Moses Joseph. 2016. Zootopia Crowd Pipeline. In *ACM SIGGRAPH 2016 Talks* (Anaheim, California) *(SIGGRAPH '16)*. Association for Computing Machinery, New York, NY, USA, Article 59, 2 pages. https://doi.org/10.1145/2897839.2927467