Creating Diversity and Variety in the People of Kumandra for Disney's Raya And The Last Dragon

Avneet Kaur
Walt Disney Animation Studios
Burbank, USA

Erik Eulen
Walt Disney Animation Studios
Burbank, USA

Johann Francois Coetzee
Walt Disney Animation Studios
Burbank, USA

Figure 1: Lineup showing a subset of the diverse crowd combination variants prior to the addition of the final look treatment for clans.

ABSTRACT

In Walt Disney Animation Studios' "Raya and the Last Dragon", the fantasy world of Kumandra is composed of five lands, representing five parts of a dragon. All aspects of the character designs were inspired by the many cultures of Southeast Asia. Each land is inhabited by a unique clan and the crowds assets need to reflect this diversity and variety both within and between the clans. This was achieved by introducing a novel approach that is modular in both design and construction of the assets. Key aspects include strategic reuse and refit, and new look techniques for creating additional variation between clans. We also employed a tracking and management system for visually validating the assets which played an important role in the efficient use of the data downstream. In addition, an extremely collaborative workflow between all departments involved was critical, including the Visual Design, Character Asset, and Crowds Simulation departments. The overall enhancements to the workflow made it possible to creatively generate the thousands of crowd assets with the desired art direction for the film.

CCS CONCEPTS

• **Computing methodologies** → **Physical simulation**.

KEYWORDS

crowd assets, model, rig, look, cloth, hair

ACM Reference Format:
Avneet Kaur, Erik Eulen, and Johann Francois Coetzee. 2021. Creating Diversity and Variety in the People of Kumandra 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/10.1145/3450623.3464660

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

ACM ISBN 978-1-4503-8373-8/21/08.
https://doi.org/10.1145/3450623.3464660

1 INTRODUCTION

In "Raya and the Last Dragon", Kumandra comprises five lands that together form the shape of a dragon - Fang, Heart, Talon, Spine, and Tail. Character designs in this fantastical world were inspired by the many cultures of Southeast Asia, and the crowds assets on the show needed to reflect this unique diversity. It is also important to the story that the clans populating each land be visually distinguishable. It was thus necessary to create rich variations within a clan while still maintaining the integrity of each clan's unique look. A novel challenge in this film was the unique wrapped clothing designs. Kumandra also needed to be populated with many different species in addition to the human characters, including mythical dragons.

Our traditional crowd assets workflow typically involves designing each individual character asset as a unit, including its look, hair, clothing, and accessories. This base workflow needed to be extended and optimized to meet the challenges posed by the large scale and distinguishing nature of the crowds in the five lands. To this end, we developed a novel modular approach all the way from design through asset creation and assembly, coupled with strategic reuse and refit, asset management and tracking, and optimization of workflows.

2 APPROACH

2.1 Modular Approach to Character Design

In addition to the usual factors for achieving variation like distinct body types and face types, we also employed a new modular approach in the design of the clothing, hairstyles, and accessories. This was a departure from previous shows where the concept of modularity was not explicit in the design phase, and only came later in execution. The first phase of this new end-to-end modular approach started early in pre-production, during which the Assets team collaborated very closely with the Character Designers. Together we formulated a strategic plan for integrating modularity

into the design of the crowd assets from the ground up. One goal was to enable the costume teams to create modular garment components such that every top could be matched with every bottom to provide maximum variety in the outfit creation. All designs were created with the belt at a fixed location to which the top and bottom garment components were anchored. This led to the possibility of hundreds of outfit combinations per crowd element.

2.2 Modular Build for Foundational Assets

Given the modular design specifications for the clothing, garment components were constructed for adult, teen, and child crowd archetypes. Each individual clothing garment was tailored separately, making a digital wardrobe of possible outfit combinations for each archetype. To explore all possible combinations in the digital wardrobe, the Assets team created a grid visualizing the outfits that could be created by mixing and matching top and bottom garments. This allowed the Character Designers to carefully choose a subset of outfits that are visually appealing while maximizing variety.

Along with the base garment mesh data, simulation cloth rig data (e.g. maps, constraints) is also packaged with each modular garment component. To enable re-use of data in a non-destructive way across design or topology iterations, all of this data is represented in uv space. Similarly, our in-house retopology tool leverages the same uv space to create the final quad topology of the render meshes, allowing topology changes of simulation meshes without affecting the work of downstream departments. This heavy use of uv space, coupled with the new modular approach to building our costumes, minimized over building, reduced the data footprint, and made it possible to quickly and efficiently absorb upstream data changes. Similar to cloth, foundational data for look, hairstyles, and accessories (e.g. headgear, jewelry, footwear) was also authored on the archetypes.

2.3 Assembly of Individual Crowd Assets

Once all the asset data (rigs, outfits, wigs, accessories, and look) is created for the foundational archetype, a recipe specifies how to assemble the variant combinations from the modular outfit, wig, and accessories components. All of the asset data is then refit from the archetype element to the next level in a crowd hierarchy, where each "silhouette" represents a different body type. The asset data is further refit to each "leaf" level character and its variants. Each leaf has a different height and face type, and leaf assets are the ones ultimately deployed in the film. The archetype, silhouette, and leaf level characters in the crowd tree share the same topology (see Figure 2). Refitting is performed via inheritance using our in-house crowd assembly system which consists of several department-based refitters (modeling, rigging, cloth, hair and look). Upstream and downstream data changes are tracked by the refitters, making it possible to seamlessly co-mingle user and procedural data across hundreds of elements and their variants non-destructively. The challenges posed by the complexity of the assets, in particular the wrapped garments, necessitated refitting from a very solid foundation and avoiding changes after refits at the silhouette and leaf hierarchy levels. Hence all data authoring for tailoring and performance refinement, including look pattern work, is done at the archetype level. Similarly, all other asset data and the recipe

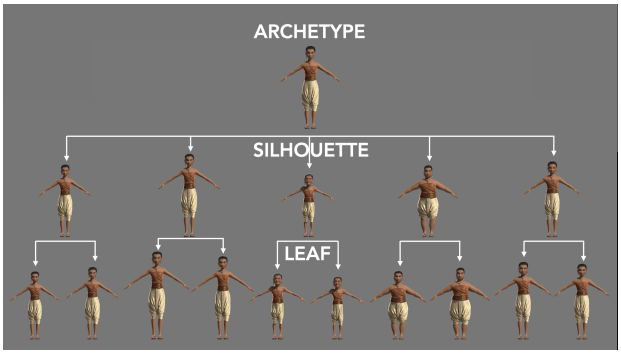

Figure 2: Crowd tree hierarchy used for refitting asset data from archetype to silhouette to leaf levels via inheritance.

for creating combination variants also reside at the archetype level, and are inherited by the leaf level variants. This enables us to make design and performance updates easily and efficiently at the archetype level, and quickly mass propagate these changes down the archetype tree. Rendered lineups of the 450+ leaf level combination variants were created for visualization and validation of each unique crowd asset used in the film. These lineups (Figure 1) were assembled into a crowd catalog that showcased the hundreds of unique combination variants available for a crowd asset. This also allowed the Art Directors to strategically design new combination variants from existing outfit and wig variants.

2.4 Additional Look Treatment for Clans

In order to make the clans visually distinguishable, each was designed to have a unique color and texture palette. The look pipeline has historically used primvars to drive color variation via swatches or baked texture maps. Given the large number of crowds assets, we needed an efficient way to achieve the clan-specific looks. The existing primvar workflow was leveraged, however instead of per asset textures, a single material was designed such that the same texture could be utilized on every crowd character on the show. This was achieved by utilizing an RGB texture map where each channel is used as a value to select from a predefined color palette. A specific color ramp curve authored for each clan is used to modulate the input value. This made it possible to stay within the space of the palette, and achieve modulation via procedural color shifts. The addition of the look treatment on top of the base asset combinations allows for 243,000 possible crowd character variations.

3 CONCLUSION

In total, 18,987 human crowd character instances were actually used in the film and were created from 450+ unique crowd asset combination variants. The overall crowds workflow also extended successfully to the dragon crowds, allowing the Art Director to mix and match dragon grooms, fins, tails, and color for a total of 1,480 visual variations. Collective ownership and trust empowered the various asset teams to work closely and effectively while experimenting with design ideas and iterating on workflows. This deep collaboration resulted in a new holistic modular approach for creating crowds assets, making it possible to achieve the desired diversity and variety in the people of Kumandra.