Headstrong, Hairy, and Heavily Clothed: Animating Crowds of Scotsmen

Paul Kanyuk Leon J.W. Park Emily Weihrich
Pixar Animation Studios

Figure 1: Scottish crowds with boistorous animation ©2012 Disney/Pixar. All rights reserved.

In an era of skinny jeans, superhero spandex, and business attire, we often forget that there was a time and place when kilts, tunics, and unkempt beards ruled the day, and where secondary motion on crowd characters was the norm. Such is the case for Brave, an animated film set in medieval Scotland, whose crowd shots comprise more than 25% of the show's total content and whose crowd characters are featured in that dreaded close to mid-ground range. With too many crowd characters to animate and simulate by hand, yet too few to allow for simplified geometry and rigging, the question becomes: "to cache or not to cache?" In this talk, we present a new answer to this age-old problem by combining retargeting, procedural variation, and "TidScene" geometry caching with a novel, clip-blending technique for cloth simulation. The result is an approach that efficiently facilitates both interactive crowd workflows and high fidelity secondary motion. Our Scottish crowds march, cheer, drink, brawl, and even cry with detail that is often indistinguishable from that of the main characters.

1 FSM Sequencing of Geometry Caches

Pixar, like many large animation studios, has explored a gamut of crowd simulation techniques, from finite state machines, to particle systems, to agent based simulators like Massive. When considering crowds in the close to mid-ground range, we've learned some hard lessons. Most importantly, animation applied to such characters with a crowds system must be easy to edit. Unfortunately, this consideration decreases the usefulness of several advanced crowds techniques, such as motion controllers, ragdoll physics, and full body procedural animation tools, whose dense data splines are ill-suited for animator editing. Spline simplification can help, but too often results in non-intuitively placed knots that are just as confusing. The old motion capture animator lament applies here as well: even if the simulation data is 90% correct, fixing the remaining 10% requires almost as much work as starting from scratch. This problem is especially acute for shows like Brave, in which the large number and close proximity of crowd characters necessitates touching-up an average of 10 characters per crowd shot. Our solution was to take a step back and revisit the old, time-honored method of finite state machine logic, while leveraging new techniques to gain more accurate and efficient results.

A second factor to consider when simulating close to mid-ground range crowds is the critical need for interactivity. To address this issue, we cached crowd character geometry using Pixar's proprietary TidScene format, a scene graph built on top of geometry data contained in a time-indexed Berkeley DB dictionary. Read performance for this format compares very favorably to the current industry leader, Alembic, especially at larger scales, and is thus suitable for interactive crowd previewing. TidScene caches were generated for each character animation clip, either a cycle or transition, and then sequenced by a technical director using our FSM-based animation clip applier, Wilma. Such sequencing of caches allowed for interactive previewing of the applied animation clips.

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

2 Animation Friendly Retargeting

Given that Brave has 80 distinct crowd character types, creating animation clips for each would have been intractable and we therefore had to pursue a form of retargeting. However, in order to produce clips that animators could still edit, it was necessary to adhere to several constraints. We found our retargeting method had to avoid introducing any new knots, preserve the temporal relationship between existing knots, and ensure that changes in knot value were the same for all knots in a given animation spline. To solve this problem, we used a combination of IK spline offsets to maintain foot-ground contact, and simple, animator-vetted, per-spline multipliers and offsets. The latter was determined by observing the steps involved in "cleaning up" the results of blindly applying animations created for our most average character to each of the 8 core body types. Offsets commonly helped keep limbs free of body intersections, while multipliers largely assisted in facial animation transfer. In practice, the majority of our animation clips were created for only 1 male character and then retargeted to the other 49. Likewise, a good portion of our clips were created by retargeting animation for 1 female character to the other 29. In a few select cases, we also chose to use 1 source animation clip for each of the 5 male body types.

3 Cloth Sequencing and Blending

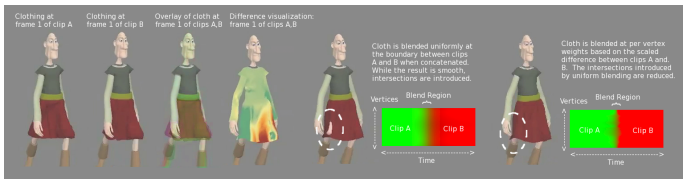

Figure 2: Blended Cloth Clips ©2012 Disney/Pixar. All rights reserved

Cloth simulation for crowds is often simply a matter of cloth rigging; pants, shirts, and most close-fitting garments need only to move roughly with the limbs for believability at a distance. Brave's story challenges this widely-held premise by featuring crowd characters with loose-fitting, multi-layered garments: a kilt, tunic, shirt, and sporran for the men, and a dress, shift, and sash for the women. For obscured, far-away crowds, our attempt to rig the cloth using subdivision surface projection provided sufficient results. However, for more noticeable crowd characters, we chose a different approach that would address the shearing and lack of secondary motion produced by the former. For each retargeted animation clip, we ran a fully-detailed cloth simulation and sequenced the result, along with its corresponding clip and cache, using Wilma. To alleviate discontinuities in the cloth motion at animation clip boundaries, we increased damping and added a "pre-roll" time to allow the cloth to settle before colliding it with moving limbs. To fix any further pops, we used a new technique to blend adjacent simulations, choosing large blend amounts for loose garments and small blend amounts for tight garments in order to minimize the introduction of intersections. To more amply blend without intersections, we utilized per-vertex blend amounts driven by the rate of increment between start and end vertex position [see Figure 2]. In cases where it was necessary to accommodate body pose changes made by animators following Wilma sequencing, we optionally ran another full cloth simulation for each touched-up character.

Although we developed a pipeline for importing trajectories from Massive for use in shots requiring complex navigation, the vast majority of the crowd shots in Brave were handled effectively using a combination of our Wilma system and animator touch-up. Our crowds pipeline as a whole, as demonstrated by the high quality and quantity of crowd shots, supports our dearly-learned conclusion: for close to mid-ground range crowds, using classic FSM-logic, spruced up with geometry caching, retargeting, and simulation blending, yields detailed yet editable animation, and therefore tremendous productivity.