

Large Scale Simulation and Surfacing of Water and Ice Effects in Dragons 2

Baptiste Van Opstal
DreamWorks Animation

Lucas Janin
DreamWorks Animation

Ken Museth
DreamWorks Animation

Mihai Aldén
DreamWorks Animation



Figure 1: Left: Water interacting with a huge new creature in “How to Train Your Dragon 2”. Right: Closeup of large-scale ice effect.

Introduction

“How to Train Your Dragon 2” introduces new creatures of truly massive scale, e.g. the Bewilderbeest shown above which measures approximately 600 ft from head to tail. This imposed unique challenges for the FX department when these creatures interact with environments like water or when they use their special ability to breath ice. While we could leverage somewhat on existing tools developed in previous productions, e.g. [Budsberg et al. 2013; Museth 2013a], it was soon clear that additional steps had to be taken to address the unprecedented scale of both the fluid simulations and the subsequent surfacing of the resulting animated particles.

Surfacing techniques

Both the water and ice effects rely heavily on particle-based fluid simulations and procedures. To this end we used Houdini’s FLIP solver, but with several important modifications in order to accommodate the massive scale. The specifics of these simulation pipelines will be discussed separately below, but for now it suffices to say that they generate on the order of 300 million particles per frame. The challenge is then to turn these massive particle clouds into representations of either water or ice surfaces. Our starting point was a proven work flow consisting of three fundamental steps: Particle rasterization, surface processing and mesh extraction. However, given the massive scales we had to develop additional techniques to reduce both the memory footprints and turn-around times. Most of the resulting tools are part of OpenVDB [Museth 2013b], which should allow for a significant degree of reproducibility.

A majority of our optimization techniques are based around novel ways to both produce and apply various types of volumetric masks during all three steps mentioned above. This includes masks derived from particle and level set properties, volumetric shadow masks from occluding geometry, frustum masks, and artistically guided level-of-detail masks. The applications of these masks are

numerous, but to mention just a few: culling of particles, localizing surface processing while preserving small details, segmentation of surfaces based on occlusion, and control of our new more aggressive adaptivity scheme for polygon extraction. To efficiently synthesize volumetric lighting effects, we used a mesh based rendering approach. Since this technique requires high quality manifold surfaces we developed a new adaptive meshing scheme based on [Nielson 2004].

Workflow for massive water simulations

To facilitate the large scale FLIP simulations we developed a cascading scales approach in which low resolution simulations were used to drive higher resolution simulations. More specifically, we adopted a work flow where each water simulation were performed at a total of four different voxel resolutions, two for the main base simulation and two for the white water simulation. The base simulation were optimized using various camera based level of detail techniques. The subsequent white water simulations were divided into regions and performed concurrently on multiple workstations to minimize memory footprint and optimize turn-around time.

Workflow for massive ice effects

An important hero-effect is the formation of large ice sculptures when sea water is regurgitate and frozen by the huge creatures. Besides the challenge of simulating and surfacing massive particle clouds this effect also called for a high degree of artistic direction. Consequently we adopted a hybrid approach where procedural velocity fields and shape proxies, which could easily be groomed, were used as sources and velocity inputs to FLIP simulations. Finally small scale details like misty layers, frozen clumps of snow and high-frequency spikes were achieved with secondary FLIP simulations and procedural displacements.

References

- BUDSBERG, J., LOSURE, M., MUSETH, K., AND MATT, B. 2013. Liquids in the crowds. In *ACM DigiPro 2013 Short Paper*, ACM, DigiPro.
- MUSETH, K. 2013. A flexible image processing approach to the surfacing of particle-based fluid animation. In *Mathematical Progress in Expressive Image Synthesis*, Springer, Symposium MEIS2013.
- MUSETH, K. 2013. VDB: High-resolution sparse volumes with dynamic topology. *ACM Trans. Graph.* 32, 3 (July), 27:1–27:22, <http://www.openvdb.org>.
- NIELSON, G. M. 2004. Dual marching cubes. In *Proceedings of the Conference on Visualization '04*, IEEE, Washington, DC, USA, VIS '04, 489–496.