Shaping Particle Simulations with Interaction Forces

Can Yuksel
DreamWorks Animation*

Kyle Maxwell
DreamWorks Animation*

Scott Peterson
DreamWorks Animation*

Figure 1. *A scene from How to Train Your Dragon 2 with pouring water simulated using our force model.*

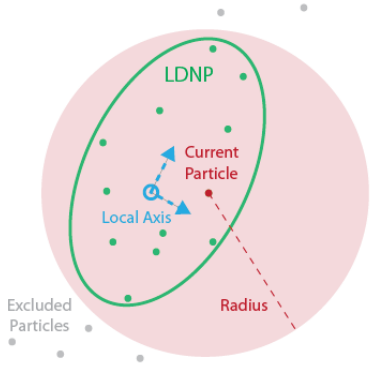

Figure 2. *A two dimensional example where the coordinate system is derived from the LDNP.*

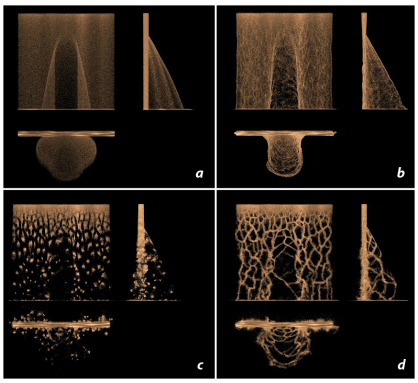

Figure 3. *(a) plain simulation and (b,c,d) with different sets of forces.*

Introduction

In this talk, we propose a force model to simulate visually challenging behaviors in particle simulations, such as pouring sand, or water splashes. Traditionally, forces derived from noise or fluid dynamics are used to emulate these effects. Noise-based forces are fast, but lack the ability to capture the intrinsic qualities of the these phenomena. While fluid dynamics accurately depict fluid properties, they are computationally expensive and difficult to art direct.

We present a new type of force which is calculated from only the intrinsic distribution of the particle system. We efficiently approximate these interactions by using a particle-to-distribution approach. Due to the force-based formulation, this approach can be used alone or coupled with traditional methods for adding detail to particle simulations. Our optimized algorithm enables us to simulate millions of interacting particles with wide areas of influence.

Force Generation

The force per particle is calculated using a local distribution of neighboring particles (LDNP). The eigenvectors of the LDNP's covariance matrix provide a basis for a local coordinate space which is used in describing the force. The user sets a feature size by choosing a search radius, effectively defining the borders of the LDNP per particle.

In addition to the search radius, there are user parameters that control the amount of force in up to ten different directions. The first three force directions push (or pull) the particles along the axes of the local coordinate space. Each of the next three directions attracts the particle to one of the local axes. Three more directions cause the particle to rotate around the local axes. The last direction attracts the particle to the LDNP center.

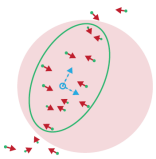

Figure 4. *Two dimensional example of one of the force directions (red arrows). The LDNP is shown for one particle, but is calculated per particle.*

Some of these forces create very distinct intuitive shapes. The force toward the axis of the largest eigenvector aligns all the particles into a long chain by reducing the variance in all but the major direction. In practice, we most often use this chain force for simulations of splashes.

Optimization

Traditionally, particle-to-particle interaction operations are computationally expensive. We start by using a kd-tree to identify particle neighbors. We then take the whole set of particles and merge those that are closer together than a fraction of the search radius. The new merged particles represent the distribution that approximates the original particles. To improve accuracy, we use k-means clustering utilizing the merged points as initial seed.

For a large search radius, this method drastically reduces the computation time because more particles are removed. This method ensures that kd-tree complexity is a function of the search radius but is nearly independent of the number of particles. Given a fixed search radius, the computation time scales almost linearly with the number of simulated particles.

Figure 5. *Energy effect using a mix of interaction forces.*

Results

Because the local distribution based forces accentuate already existing patterns, they are relied on by artists at DreamWorks Animation to enhance a large portion of the particle simulations done at the studio.

References

Blow, Jonathan. "Statistics." My Friend, the Covariance Body (from The Inner Product). 17 Jan. 2004. Number None.
http://number-none.com/product/My%20Friend,%20the%20Covariance%20Body