

# Capturing Thin Features in Smoke Simulations

Magnus Wrenninge\*  
Sony Pictures Imageworks

Henrik Fält  
Sony Pictures Imageworks

Chris Allen  
Sony Pictures Imageworks

Stephen Marshall  
Sony Pictures Imageworks



Figure 1: Absolem from Alice In Wonderland. ©2010 Disney Enterprises Inc.

## 1 Introduction

The perceived quality and realism of a fluid simulation is directly related to a few technical and measurable parameters: the resolution of the simulation, and the numerical diffusion of the algorithm used. A high resolution simulation can resolve fine details and vortices well, and low numerical diffusion ensures that the energy in the simulation and the sharpness of the features are not lost as the simulation progresses. The simplest fluid simulation is concerned only with the velocity of the fluid and considers fluids that fill the entire domain (gas simulation). When simulating smoke or some other quantity suspended in the fluid, one considers both the velocity of the fluid as well as the concentration of density at each point in the domain. The smoke field may influence the motion of the fluid through gravity, expansion, or other forces, in which case it is said to be coupled to the simulation.

It is common to use the same resolution for both the velocity and smoke field, as the voxels (in the case of finite differences) are aligned, making for a straightforward implementation of the solver. In this case, the size of the smallest representable fluid vortex is the same as the smallest representable smoke feature. However, when considering the perceived visual quality of a smoke simulation we have found that capturing fine details in the smoke field is of greater importance than capturing small-scale fluid motion. That is, the viewer will accept a simulation with sharp smoke features as realistic, even though the feature scale of the fluid vortices is large.

## 2 Decoupling smoke field resolution

We present a system in which the smoke field resolution (the visual quality) is higher than the velocity field's. We do this for two main reasons: the pressure solution of the velocity field is much more expensive to compute than the advection of the fluid and smoke fields, and it allows control over the perceived viscosity of the fluid independently of the sharpness of smoke details.

The technique for coupling the high resolution smoke field to the velocity field is simple: Before applying gravity or expansion forces, the high resolution smoke field is resampled using an averaging filter kernel into a field with resolution identical to the velocity field. This ensures that smoke features smaller than a simulation voxel have an appropriately reduced effect. Using the low resolution representation of smoke we can then run the velocity field simulation using standard techniques.

Our system uses the FLIP algorithm [Zhu2005] to calculate the fluid motion of the gas, which has excellent numerical diffusion properties. The velocity field uses the common MAC structure, and the smoke field is stored cell-centered in a sparse voxel structure (Field3D's Sparse-Field), ensuring that memory use is proportional to the number of voxels filled with smoke.

## 3 Sparse particle-based advection

Using a high resolution smoke field solves the problem of resolving fine features. However, any grid-based advection scheme will suffer from numerical diffusion due to repeated interpolation. Another problem stems from the time complexity of these algorithms, as they scale with the resolution of the simulation domain. Preferably, we would like the advection time to be proportional to the number of voxels occupied in the smoke field, so that empty space is 'free'.

We use a dual eulerian/lagrangian representation of the smoke field: sparsely placed particles are used to advect the smoke field, and a high resolution voxel grid is used to couple the smoke field to the fluid simulation. The primary representation is the particles, which is the state that is kept from frame to frame and the conversion of particles to the high resolution voxel representation is done by splatting. Diffusion and other effects are computed on the high resolution voxel grid and changes to the voxel grid representation are applied back to the particles using the FLIP technique (using only the derivative of smoke density). Our technique has very low numerical diffusion, as forward advection of particles is diffusion-free and the FLIP update has very little inherent diffusion (depending on the FLIP blending parameter).

A major drawback with using a particle representation of smoke is that the granular nature quickly becomes apparent as the smoke is advected and spreads. To solve this problem we employ an adaptive smoothing step: after each advection is performed the particle count for each voxel is recorded, and if it falls below a given threshold, one or more of the contained particles are 'split', adding an extra particle and dividing the smoke quantity between the two.

The sparse nature of the particle representation gives our technique a time complexity that is linearly proportional to the number of voxels occupied with smoke, regardless of the resolution of the smoke field. In production scenes the technique is often 20-30 times faster than semi-lagrangian advection at similar resolution, and gives a much higher quality visual impression.

## 4 Use in production

Our technique for smoke simulation was used extensively to create the Mushroom Forest environment in Alice in Wonderland. All of the smoke elements surrounding Absolem, the smoking caterpillar, were entirely simulated, and took advantage of the decoupled velocity field and smoke field resolutions. Because the implied scale in the scene is very small, fine vortices were often visually disturbing, and the resolution that gave the best result was usually in the range from  $50^3$  to  $150^3$ . Conversely, in order to maintain sharp visual features at full 2K film resolution the smoke fields were of very high resolution, with final elements typically in the range from  $1000^3$  to  $4000^3$ .

## References

ZHU, Y., AND BRIDSON, R. 2005. Animating sand as a fluid. ACM Transactions on Graphics 24, 3 (Aug.), 965–972.

\*e-mail: magnus@imageworks.com