Real-time Control of Newtonian Fluids

Károly Zsolnai and László Szirmay-Kalos

Budapest University of Technology and Economics, Hungary

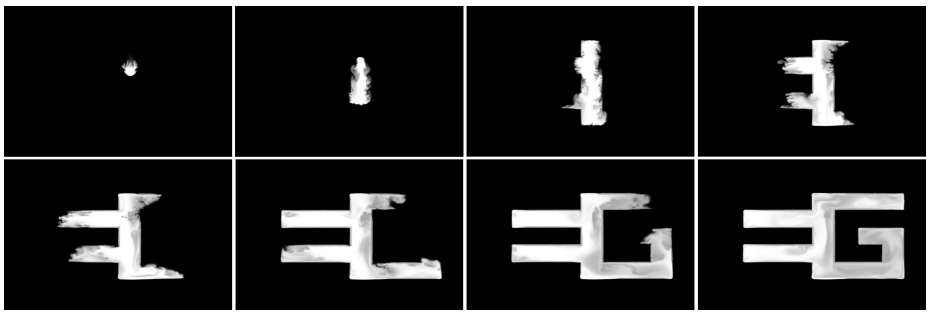

Figure 1: *Animating the fluid into the EG logo with the proposed method.*

Abstract

In this paper we address the fluid control problem, where, alongside simulating the motion of fluids, an arbitrary density distribution (a shape of any kind) is given, and forces are exerted on the system with the intention that the fluid would sooner or later take this form. Prior research work has shown that the problem is challenging because every region is tightly bound to its neighborhood, therefore a force that acts upon a point will also have effect on nearby regions, making the controlling process unpredictable. Utilizing optimization and control theory, current state of the art methods are able to give visually pleasing results at the cost of 5 to 7 minutes of computation time per frame. We present a novel solution for the fluid control problem with certain restrictions, making it possible to solve it in real-time.

1. Introduction

Fluid simulation means the mimicking of real fluids by solving the governing equations with parameters and boundary conditions reflecting a real scene. *Fluid control*, on the other hand, is the determination of parameters in a way that the resulting fluid motion follows a prescribed behavior.

Let us consider a fluid element of unit mass. Its motion is described by Newton's second law, stating that the acceleration $\vec{a}$, which is the derivation of velocity $\vec{u}$ is proportional to the total force:

$$\vec{a} = \frac{\mathrm{d}\vec{u}}{\mathrm{d}t} = \vec{F}_{\mathrm{int}} + \vec{F}_{\mathrm{ext}}$$

where the force is expressed as the sum of internal forces $\vec{F}_{\mathrm{int}}$ and external forces $\vec{F}_{\mathrm{ext}}$. Taking into account the specific phenomena of fluids, this equation becomes the first Navies-Stokes equation:

$$\frac{\partial \vec{u}}{\partial t} + (\vec{u}\cdot\nabla)\vec{u} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\vec{u} + \vec{F}_{\mathrm{ext}}, \qquad (1)$$

where $\vec{u}$ is velocity field, ρ stands for the density, p for pressure, ν denotes the kinematic viscosity of the fluid, and the internal force is decomposed to the force of pressure and the internal friction. Advection term $(\vec{u}\cdot\nabla)\vec{u}$ shows up because the fluid element is not followed in the Eulerian viewpoint, but the location in focus is fixed to the lattice points of a static grid. This Navier-Stokes equation expresses the conservation of momentum. In addition to this, we should also enforce the conservation of mass to have a realistic simulation. For a detailed introduction to fluid dynamics we refer the reader to [CM93] and [Bri08].

Considering fluid control, we have to find an external force field that makes the density field converge to a target density in time, while retaining only natural movement in the fluid domain. Applying an external force field intuitively

means that we have little spoons of infinitesimal size, which we use to stir the fluid in the chosen directions, doing it many places at once. Stirring is determined by the difference between the obtained and the target density field according to the *negative feedback* principle of control theory. Adding external forces accelerates the fluid and keeps it in motion. However, when the target density is obtained, it should be maintained in a natural way. The simplest way to do this is to increase the viscosity, i.e. the friction to dissipate the kinetic energy and to make the fluid stop. Note that when the density reaches the target density, control force $\vec{F}_{\text{ext}}$ becomes zero, having dissipated the kinetic energy by friction, velocity $\vec{u}$ also becomes zero, which makes pressure differences ∇p zero as well, finally freezing the fluid. On the one hand, "freezing" some parts of the fluid by increasing the viscosity prohibits the convergence of other parts since such actions increase the constraints to be satisfied. On the other hand, stopping the flow in itself is unrealistic since fluids are supposed to be in motion.

In this paper, we present a novel method to simplify the optimization approach while improving its convergence by controlling not only the friction but also the total internal force. We aim at a dynamic balance where the objective is satisfied by a constantly moving fluid. The new method is able to solve the fluid control problem in real time while retaining realistic looking fluid flows.

2. Related work

Solving the Navier-Stokes equations means the calculation of time differentials and applying a stable integration scheme. Jos Stam proposed a stable advection formula [Sta99]. Simulating fluids on a finite resolution grid has serious drawbacks. Simulation of turbulent flows would require the modeling of high frequency changes in the velocity field. Unfortunately, this is quite costly as increasing the grid resolution raises the computational cost of the simulation significantly. Fedkiw et al. proposed a way to reinject the lost energy to the system in the form of an external force field [FSJ01], called *vorticity confinement*. The second drawback of the finite grid approach is that we obtain the velocity information between the grid point by bi- or trilinear interpolation. Unfortunately, the Navier-Stokes equations are not linear, therefore higher order advection methods, such as the Back and Forth Error Compensation and Correction Advection [DL03] and MacCormack Advection [SFK*08] yield better quality solutions. It is also possible to obtain additional realism by using the Wavelet Transform [KTJG08].

Several solutions exists to solve the fluid control problem, such as Jos Stam's adjoint method [MTPS04], or the control of fluids on triangle meshes [RTWT12]. The algorithm of Shi and Yu adds a long-range force field to even out the distribution of the fluid on a macro level, and a short range field to carve out the fine details [SY04]. Our approach builds upon this method, so we briefly summarize it.

2.1. Short-range force fields

The short-range force field is to be constructed to carve out the fine details of the target distribution locally. In [SY04], Shi and Yu have shown that the minimization of the following functional of short-range force $\vec{F}^S$ would make a desirable field when the time is stepped forward:

$$c_1 \sum_{\vec{x}} \left(\Delta\rho(\vec{x}) - \lambda\Delta\rho^a(\vec{x})\right)^2 + c_2 \sum_{\vec{x}} \left(DIV(\vec{x})\right)^2 -$$
$$c_3 \sum_{\vec{x}} \left(\frac{\vec{F}^S(\vec{x})}{\left|\vec{F}^S(\vec{x})\right|} \cdot \frac{\nabla\rho(\vec{x})}{|\nabla\rho(\vec{x})|} \right)^2 -$$
$$c_4 \sum_{\vec{x}} \sum_{\vec{y}} \cos(\theta(\vec{y}) - \theta(\vec{x})) +$$
$$c_5 \sum_{\vec{x}} \sum_{\vec{y}} \left(\left|\vec{F}^S(\vec{y})\right| - \left|\vec{F}^S(\vec{x})\right| \right)^2. \quad (2)$$

The $c_1, \ldots, c_5$ coefficients are used to assign different weight values to the terms. We give an intuitive interpretation of the formula: the *first term* is responsible to ensure that the difference between the current and the target density is minimal by maximizing the amount of density change in excess density areas, and minimizing it at areas that match the target density well. The mass flow has to be divergence-free throughout the simulation and control process, which can be enforced by using the long-range force field and evaluating the projection step afterwards to keep the divergence-free property, i.e. to preserve the mass during fluid motion. The *third term* is the dot product between the normalized direction vectors of the force field and the gradient of the density field. As $\cos(\theta(\vec{x}))$ represents the orientation of the short-range force field in point $\vec{x}$, the *fourth and fifth expressions* are to minimize the amount of directional variance and the magnitude of the applied short-range force field.

2.2. Long-range force fields

If some part of the fluid domain has excess density, meaning that the density at point j, ρ_j, is higher than the target density ρ_j^t given by the input distribution, the region will transport density by exerting force towards the direction of its neighborhood for those who have lower density than the target. As in physics, the exerted force weakens with the square distance. This long-range force field is responsible for arranging the density field so it can start converging to the target on macro level:

$$\vec{F}_i^L = \sum_j \frac{\left[\rho_j - \rho_j^t\right]^+ \vec{r}_{ij}}{\left|\vec{r}_{ij}\right|^3}, \quad (3)$$

where $\vec{r}_{ij}$ is a vector that points from grid point i to j and superscript $^+$ denotes the operation which keeps the original value if it is positive and replaces negative values by zero. Constructing this force field has a quadratic computational cost of the number of discrete points, which is too demanding. Luckily, it isn't necessary at all, since as we increase the

distance will make the force decay, we can safely define a maximum distance that is to be computed from every point.

2.3. The new method

The most important attribute of the new approach would be not to use short-range force field due to its computational costs, therefore omit carving out some of the fine details, but design a different, cheaper force field that is able to mobilize large amounts of density towards the target distribution while still retaining a realistic looking flow.

In our version we have also relaxed the evaluation of the long-range forces to regions that have nonzero target density,

$$\vec{F}_i^L = \sum_j \frac{\left[\rho_j - \rho_j^t\right]^+ \vec{r}_{ij}}{\left|\vec{r}_{ij}\right|^3}, \quad \forall j: \rho_j^t > 0, \tag{4}$$

to reduce its cost to be proportional to only the volume where the target distribution is nonzero, as opposed to the original method, where it is evaluated in the full simulation domain.

It can occur that upon reaching convergence the excess density remains around the outer side of the boundary of the target distribution. To clean up these details around the boundary, an additional force field may be applied that pushes the density back into the domain of the target distribution. The use of this *cleanup force* field is entirely optional, and it would consist of vectors that are oriented from points that contain density but have zero target density and are near the boundary to every other point that has nonzero target density. Formally:

$$\vec{F}_i^C = \sum_j \frac{\rho_j \vec{r}_{ij}}{\left|\vec{r}_{ij}\right|^3}, \quad \forall i: \rho_i^t > 0 \text{ and } \forall j: \rho_j > 0 \wedge \rho_j^t = 0. \tag{5}$$

Depending on the target density distribution type, Neumann- or Dirichlet-type boundary conditions may be also applied.

Deciding what should happen after the target density is reached is the crucial task of fluid control. The most straightforward solution is to set ν to a very high value to "freeze" the fluid in the convergent subdomains. The results will remain correct, but not very lifelike and generally unconvincing. Here we address the shortcomings of this approach by introducing a *scaling factor s* not only for the friction but for the total internal force of the fluid. This idea may sound quite counterintuitive: why speed up the fluid at regions where it already looks correct? It would be reasonable to say that it is the exact opposite of what should happen. On a microscopic level, freezing the fluid domain would always be the best choice: if we have only a few points in space, freezing them by assigning a very high kinematic viscosity value upon reaching the target distribution will ensure that they will remain in the correct state at all times. However, let us consider a simple example on a macroscopic level, where we have a fluid domain of significant size where the target distribution can be reached only by going through a narrow choke point. At this region, the fluid will start freezing, therefore it will prevent any further fluid movement, making it impossible to get density through. This scenario will not be restricted only to narrow choke points: for almost every practical case on closed shapes, the closer we are to the state of convergence, the higher the probability for this to happen.

Intuitively, freezing the fluid can be associated with the "after you're done, just stop and rest" behavior, as opposed to increasing the total internal force, which would mean more like "after you're done, start helping others". This behavior will not only allow the fluid to flow through narrow choke points, but effectively transfer density to neighboring regions of poor convergence and preserve fluid movement after the target density is reached.

Putting it all together, the external or control force is the sum of long-range and the cleaning-up forces:

$$\vec{F}_{\text{ext}} = \vec{F}^L + \vec{F}^C,$$

The Navier-Stokes equation is modified and the internal force is scaled up by s where the actual density is close to the target density. The scaled up internal force will maintain some motion even close to the converged state when the control force drops to zero. The modified equation is:

$$\frac{\partial \vec{u}}{\partial t} + (\vec{u}\cdot\nabla)\vec{u} = s\left(-\frac{1}{\rho}\nabla p + \nu\nabla^2\vec{u}\right) + c\vec{F}_{\text{ext}},$$
$$s = \begin{cases} 1+\delta, & \text{where } \left|\rho^t - \rho\right| < \varepsilon, \\ 1, & \text{elsewhere.} \end{cases} \tag{6}$$

where c is a constant factor to control the magnitude of the control force field (and therefore the speed of convergence) and δ is small value to keep the fluid in motion after convergence when the actual density is equal to the target density within error threshold ε. This technique is capable of guiding the fluid towards the target distribution in real-time.

3. Results

The properties of the proposed technique are as follows:

- Both the fluid simulation and the control algorithm can be run in parallel as they take 18 and 14 milliseconds at most respectively on an 512^2 grid with 20 Jacobi iterations, therefore it is a real-time solution.
- It yields remarkably fast convergence speeds.
- It is to be used with target distributions of low-variance for a high measure of realism, or it is to be used on more complex, higher-variance distributions with the aid of biased diffusion at the cost of less realism and more computational demands.
- Relying only on local data, it can be extended to simulations of any dimension with favorable amount of computational overhead.

We compared the new method to an uncontrolled simulation, where Neumann-type boundary conditions are defined but no control force field is applied. Even on the simplest

scenarios, this uncontrolled flow does not deliver a desirable degree of convergence (upper image of Figure 2). This method is also unable to handle any amount of density which is not near or inside the target distribution domain. The results using the new method were rendered in real-time and are shown in the lower image of Figure 2, which, using roughly the same amount of density, was able to achieve fast convergence in the same amount of time.

4. Conclusions

In this paper, we have proposed a novel algorithm to the fluid control problem, where an initial state and a target distribution are given for the fluid, with the intention that it would sooner or later take the form of this distribution. A solution to this problem is an external force field that is changing in time, describing the forces that have to be exerted on the fluid to take a given shape. It is desirable that this force field makes the fluid converge in a short amount of time while retaining only natural movements in the fluid domain.

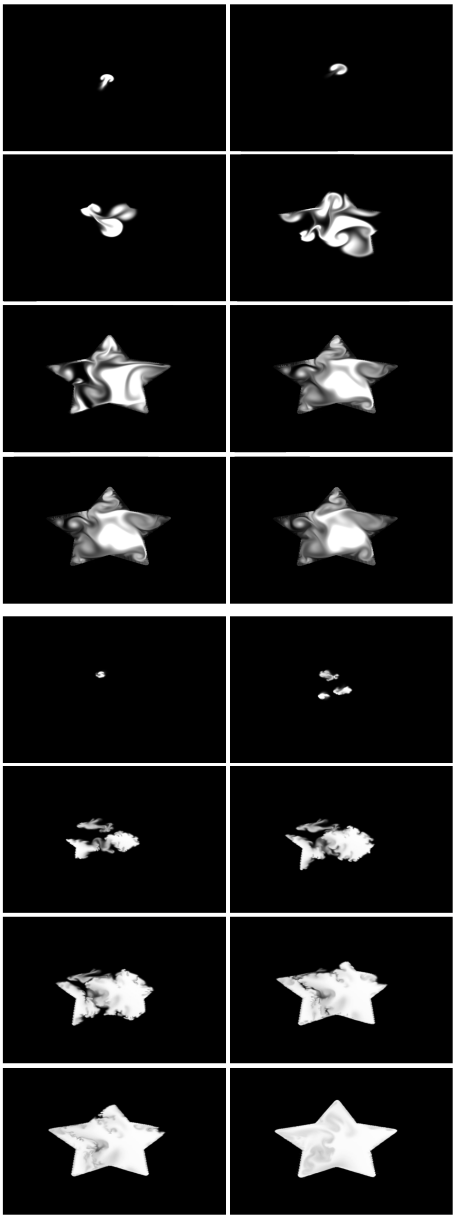

Figure 2: *An imitation of the solution to the fluid control problem where no control forces, only boundary conditions are used (upper). The example shows that even if the fluid is locked inside the domain of interest, it is highly unlikely that it would suddenly flow into the shape of a star. The proposed method provides good coverage of the target density, and is aware of the regions of poor convergence, which are constantly helped out by nearby regions (lower). Roughly the same amount of density is used in both cases.*

References

[Bri08] BRIDSON R.: *Fluid Simulation*. A. K. Peters, Ltd., Natick, MA, USA, 2008.

[CM93] CHORIN A., MARSDEN J.: *A Mathematical Introduction to Fluid Mechanics*. Texts in Applied Mathematics. Springer-Verlag, 1993.

[DL03] DUPONT T. F., LIU Y.: Back and forth error compensation and correction methods for removing errors induced by uneven gradients of the level set function. *J. Comput. Phys. 190*, 1 (Sept. 2003), 311–324.

[FSJ01] FEDKIW R., STAM J., JENSEN H. W.: Visual simulation of smoke. In *ACM SIGGRAPH 2001* (2001), pp. 15–22.

[KTJG08] KIM T., THÜREY N., JAMES D., GROSS M.: Wavelet turbulence for fluid simulation. In *ACM SIGGRAPH 2008 papers* (New York, NY, USA, 2008), SIGGRAPH '08, ACM, pp. 50:1–50:6.

[MTPS04] MCNAMARA A., TREUILLE A., POPOVIÆ Z., STAM J.: Fluid control using the adjoint method. *ACM TRANS. GRAPH. (SIGGRAPH PROC 23* (2004), 449–456.

[RTWT12] RAVEENDRAN K., THUEREY N., WOJTAN C., TURK G.: Controlling Fluids using Meshes. *SCA '12: Proceedings of the 2012 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (July 2012), 1–8.

[SFK*08] SELLE A., FEDKIW R., KIM B., LIU Y., ROSSIGNAC J.: An unconditionally stable maccormack method. *J. Sci. Comput. 35*, 2-3 (2008), 350–371.

[Sta99] STAM J.: Stable fluids. In *Proceedings of SIGGRAPH 99* (1999), Computer Graphics Proceedings, Annual Conference Series, pp. 121–128.

[SY04] SHI L., YU Y.: Inviscid and incompressible fluid simulation on triangle meshes. *JOURNAL OF COMPUTER ANIMATION AND VIRTUAL WORLDS 15* (2004), 3–4.