Honey simulation with different viscosity formulations

Essay 3

Adam Alsegård*

TNCG13 - SFX: Tricks of the trade
Linköping University - Media Technology

December 5, 2017

Abstract

*This essay will look into the field of fluid simulation. More specifically how highly viscous fluids, such as honey, can be simulated in computer graphics. Two technical papers will be presented and compared. First **An Implicit Viscosity Formulation for SPH Fluids** by Peer et al. [1] who proposes a particle based Lagrangian method. The second paper is **Variational Stokes: A Unified Pressure-Viscosity Solver for Accurate Viscous Liquids** by Larionov et al. [2], who instead proposes a grid based Eulerian approach. The essay ends with a discussion of which approach is the best way forward.*

I. Introduction

Honey, mud, dough, toothpaste, tar, molasses. The examples of highly viscous liquids are numerous. While one could easily think that these materials can be simulated by a regular fluid simulation, with some tweaks to the viscosity term, that is not entirely true. To achieve the characteristic effects of these viscous liquids, such as buckling, meandering and rope-like coiling, new formulations has to be made from the Navier-Stokes equations. Two such novel formulations are presented in this essay. First an implicit formulation based on Smoothed Particle Hydrodynamics (SPH) proposed by Peer et al. [1] and then an implicit variational formulation of the Stokes problem in a grid based system proposed by Larianov et al. [2].

II. A particle based formulation

All modern particle based fluid simulations are built upon the SPH approach, in which viscosity is an important stability aspect. The particle based formulation can be either explicit (i.e. forces are applied explicitly) or implicit (i.e. a linear system for unknown velocities at the next timestep needs to be solved). When simulating highly viscous fluids an explicit approach require a very small timestep, which then is prone to overshooting, while an implicit formulation needs to solve a complex linear system. An implicit approach works with a significantly larger timestep however and therefore overshooting cannot occur and it is thus often preferred over an explicit formulations.

The implicit formulation proposed by Peer et al. [1] is more efficient and can handle a wider range of viscosities than previous methods. In highly viscous fluids the incompressibility and viscosity constraints can interfere. Previous methods solve this by performing two pressure projection steps. The main contribution of the proposed method is to reduce this to only one pressure projection step by addressing the interference of pressure and viscosity computation. While previous implicit methods assume that the input velocity field is divergence-free the proposed formulation instead reconstructs the velocity field from a target velocity gradient. This gradient encodes a desired shear rate damping and preserves the velocity divergence that is introduced by the SPH pressure solver to counteract density errors. The target gradient ensures that pressure and viscosity computations do not interfere and therefore only one

*adaal265@student.liu.se

pressure step is required. This is shown to be especially relevant for high viscosities and large timesteps.

Peer et al. does not employ a classic force-based formulation with standard physical parameters to build the linear system, instead they formulate a set of constraints that should be fulfilled at the next timestep. This is known as Position Based Dynamics (PBD). As in most PBD formulations Peer et al. make use of a non-standard parameter, here called the *shear viscosity parameter*, and use it in the range between zero and one.

The method solely relies on the plain Navier-Stokes equation for fluids. It does not account any rigid motion or elasticity.

i. Velocity gradient

Viscosity forces are usually computed from the divergence of a viscous stress tensor τ, which itself is computed from the strain-rate tensor D which in turn is the symmetric part of the velocity gradient $\nabla \mathrm{v}$. The viscosity forces always aim to minimize the entries of the strain-rate tensor. The different aspects in computing these forces makes this explicit formulation rather involved when used on highly viscous fluids.

Instead a novel implicit approach is used where a momentum-preserving velocity is computed based on a desired velocity gradient at the next timestep. This does not only predict a velocity gradient that obtains highly viscous materials but it also preserves the corrections of density deviations introduced by the pressure solver. Velocity changes due to pressure forces are computed first, then the viscosity solver computes velocities that account for viscosity, but do not influence the rate of volume change. Because the pressure and viscosity solvers do not interfere we do not need the final pressure solve (which is otherwise used to perturb the result of the viscosity solver.)

In the formulation proposed by Peer et al. the velocity gradient itself is decomposed into three components: the spin tensor R, the shear-rate tensor S and the expansion-rate tensor V.

$$\nabla \mathrm{v} = \underbrace{\frac{1}{2}(\nabla \mathrm{v} - (\nabla \mathrm{v})^T)}_{\mathrm{R}} + \underbrace{\frac{1}{3}(\nabla \cdot \mathrm{v})I}_{\mathrm{V}} + \underbrace{(\frac{1}{2}(\nabla \mathrm{v} + (\nabla \mathrm{v})^T) - \frac{1}{3}(\nabla \cdot \mathrm{v})I)}_{\mathrm{S}} \quad (1)$$

R describes the vorticity (i.e. the rate of rotation of a particle), V describes the density change at a particle (e.g. can be used to realize bulk viscosity) and S describes the rate of shear strain at a particle. The divergence of the velocity field is then computed as the trace of the velocity gradient.

When predicting the desired velocity gradient Peer et al. distinguish two different cases; when the particle density is above rest density and when it is below. When the particle is above rest density the key aspects is to preserve the expansion-rate tensor V and to control the shear rate with a non-physical parameter for shear viscosity ($0 \leq \xi \leq 1$).

$$\nabla^{\tau} \mathrm{v} = R + V + \xi S \quad (2)$$

When the particle density is below the rest density we have two cases. If the divergence of the velocity field is negative we have the same gradient as before, but if the divergence is positive then a maximum bulk viscosity is introduced by eliminating V from the predicted gradient (Eq. 2). This addresses a negative pressure issue that is normally not handled in SPH approaches. This formulation aims at stopping adjacent particles from moving away from each other and thus reduces the artificial volume gain that otherwise would occur.

Finally a first-order Taylor approximation is used to reconstruct the final velocities from the target gradient $\nabla^{\tau} \mathrm{v}_i$ resulting in a linear system. The average of the gradient of two adjacent particles are used to guarantee momentum-preserving velocity changes at the particles. This also accounts for the diffusion of the rotation rate that is particularly present in highly viscous fluids.

i.1 Boundary handling

The boundary handling is generally handled during the pressure solve but Peer et al. pro-

poses two further conditions to be handled during the viscosity solve. The first is *sticky boundaries*, motivated by how honey sticks to solid surfaces, which dampen the shear rates between fluid and solid particles. The second is *separating boundaries* which is the opposite, the fluid can flow freely, but not into the boundary.

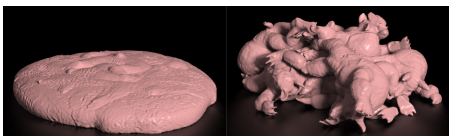

Figure 1: *Different viscosities. The left is simulated with $\xi = 0.8$ while the highly viscous setting to the right is simulated with $\xi = 0.2$. The computation time per frame was 11s for the left and 37s for the right image.*

ii. Discussion

The proposed formulation deviates from a physical model in two ways. It is parameterized with a non-physical constant ξ and the vorticity diffusion is not encoded in the formulation of the target gradient. The parameter ξ is a non-standard parameter and does not correspond to dynamic or kinematic viscosity. Nevertheless, ξ is physically motivated. It governs the damping of the shear rate in the fluid, thus it governs viscosity.

The formulation has multiple approximations, such as the first-order Taylor approximation, which result in numerical errors in spin, divergence, target gradient and momentum. However, the errors are small and the experiments provided by Peer et al. show that these errors are negligible.

Non-destructive velocity reconstruction (i.e. zero viscosity) works in the proposed solver but due to the smoothing of the velocity field in the reconstruction process this approach should not be used for low viscous fluids such as water. It also has a problem with rigid-like objects that require many solver iterations. For such materials this approach is significantly less efficient than viscoelastic solvers.

The solution is fully parallelized and implemented in Houdini. It can easily be added to a SPH solver as a post-processing step. The computation time of the viscous solver scales with the viscosity parameter ξ. With multiple examples Peer et al. demonstrate that they indeed can simulate a wide range of viscosities (see Figure 1), even water and rigid-like materials (even if this method is not the best candidate for those two). The proposed viscosity solver is particular appropriate for multiple phases (i.e. having materials with a range of values of ξ).

One of the biggest drawbacks of the paper is that no performance comparison has been performed whatsoever and that no mapping between the viscosity constant ξ to dynamic viscosity μ have been found or even tried..

A given example of future work is that the method heavily damps shear rates which means that the particle movement is heavily restricted. This leads to a visible pattern of the initial sampling. To improve this a Poisson disc sampling could investigated and maybe used instead.

III. A grid based formulation

In contrast to the previous Lagrangian approach Larianov et al. [2] has developed a novel Eulerian method where an unsteady Stokes solver can be used to simulate coupled viscous and pressure forces for grid-based liquid. Whereas Peer et al. and most other modern fluid simulators treat viscosity and pressure in separate solver stages, which reduces accuracy and yields incorrect free surface behavior, the implicit variational formulation of the Stokes problem proposed by Larianov et al. leads to a symmetric positive definite linear system. This in turn gives properly coupled forces, provides unconditional stability and treats difficult boundary conditions naturally through simple volume weights. Surface tension and moving solid boundaries are also easily incorporated.

The key contribution of the paper is that the standard Eulerian viscosity and pressure steps has been replaced by a novel unsteady Stokes step that reunites these forces and accurately recovers the missing effects of coupling.

A dominant source of error introduced by splitting is an improper treatment of the liquid's free surface. The correct bound-

ary requires the interface to be traction-free, $t = \sigma n = 0$ where t is the surface traction, σ is the total fluid stress tensor and n is the surface normal. However, this condition inherently requires tight pressure-viscosity coupling because σ is a sum of pressure and viscous shear stresses.

Unlike a steady-state Stokes solver, which yields indefinite linear systems, an unsteady (i.e. time dependent) solver can be used to yield a symmetric positive definite system. By discretizing a carefully chosen variational form they can also allow the free surface conditions to be enforced easily and implicitly as natural boundary conditions.

i. Fluid equations

The foundation for the formulation by Larionov et al. is the incompressible Navier-Stokes equations. From them they apply operator splitting and discretize in time, yielding a two step scheme. An intermediate velocity is arrived at first by solving advection and external forces.

They then simultaneously integrate pressure and viscous forces using backward Euler by solving

$$\frac{u_{n+1} - u^*}{\Delta t} = \frac{1}{\rho}(-\nabla p + \nabla \cdot \tau), \tag{3}$$

$$\nabla \cdot u_{n+1} = 0, \tag{4}$$

$$\tau = \mu(\nabla u_{n+1} + (\nabla u_{n+1})^T), \tag{5}$$

where ρ is density, u is velocity, p is pressure, τ is the symmetric deviatoric stress tensor, μ is the dynamic viscosity coefficient and t is time. These are the equations for an implicit unsteady Stokes flow allowing for spatially varying viscosity. The main contribution by Larionov et al. is a solver for this PDE.

ii. Boundary conditions

For an Eulerian approach it is difficult to derive appropriate discrete boundary conditions for non-grid aligned boundaries, especially for free surfaces. Larionov et al. solve this by using a variational framework that handles irregular geometry through natural boundary conditions.

There are two separate cases to be handled, one for free surfaces and one for solid surfaces. For the free surface case an integration step is performed over the liquid (i.e. non-air region). With the new formulation this is now handled automatically by the PDE, because they now have the coupling between pressures and viscous stresses that were absent in previous methods. For the solid case the integration is instead performed over the fluid (i.e. non-solid region).

Unfortunately both cases cannot be enforces simultaneously in the continuous case. However, by discretizing the integrals, following a variational finite difference framework, they can be combined. A FEM could be used instead but Larionov et al. choose the first alternative because it is simple, convergent and does not require stabilization.

By combining the two formulations at the discrete level Larionov et al. exploit the natural conditions for both boundary types simultaneously. The resulting sparse, symmetric indefinite system is then transformed into a symmetric positive-definite (SPD) system by finding the Schur complement and eliminating velocity, instead of stress which is the common way for steady Stokes.

A stress reduction and a null space elimination are also performed before arriving at the final boundary conditions that are used in the simulations.

iii. Discussion

The proposed method can be used as a replacement for the viscosity and pressure steps in standard fluid simulators. The method was integrated into Houdini and Larionov et al. have released a plug-in for Houdini 16.

Through experiments they show that the characteristic rope-like coiling is recovered (see Figure 2) and that the method preserves a higher degree of detail in viscous liquid surfaces than previous grid based methods, while still maintaining support for features such as moving boundaries and surface tension. They also prove that the method is convergent through grid refinement studies on analytical problems in two dimensions.

In their results Larionov et al. have multiple other examples of different simulation

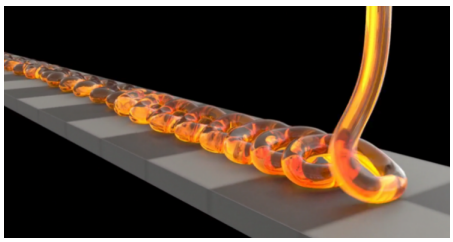

Figure 2: *A simulation of rope-like coiling.*

environments that their solution can handle, and they have also verified that the solver computes exact solutions, for time-dependent Stokes problems with linear solutions on irregular (i.e. non-grid-aligned) domains at least.

The solution is parallelized inside Houdini and a comparison of the computational cost has also been performed. In general they found that their solution was 1.6x to 3.4x times more expensive than Houdini's decoupled approach, which in turn is more expensive than per-component Laplacian viscosity. The increase is not surprising giving the much larger linear system.

However, even though the proposed method has a higher visual and physical fidelity, the increased cost indicates that this solver only should be used when realistic coiling effects, sharper surface details or greater physical fidelity is needed. Otherwise a decoupled variational viscosity solution or a Laplacian form should be used instead.

IV. Summary

This essay has become somewhat of an *Lagrangian vs. Eulerian* comparison. Larionov et al. [2] actually refers to Peer et al. [1] in their background chapter. According to Larionov et al. non of the existing Lagrangian (particle based) approaches "present quantitative evidence of accuracy" and because these methods so far leads to indefinite linear systems, in contrast to the positive definite system achieved by Larionov et al., the performance is much higher for Eulerian (grid based) approaches and Larionov et al. make the conclusion that this is why grid based methods are dominant in the industry.

Peer et al. on the other hand say that Eulerian methods require special treatment for free surface to allow for rotational fluid movements, and that they only for divergence-free velocity fields. Which was true when the paper was published but it is exactly these areas that Larionov et al. has improved upon.

Overall I would say that the grid based approaches has the upper hand, both in accuracy and performance, but it will be interesting to see what the next step will be for both of these approaches regarding highly viscous fluids.

References

[1] Andreas Peer, Markus Ihmsen, Jens Cornelis, and Matthias Teschner. An implicit viscosity formulation for sph fluids. *ACM Trans. Graph.*, 34(4):114:1–114:10, July 2015.

[2] Egor Larionov, Christopher Batty, and Robert Bridson. Variational stokes: A unified pressure-viscosity solver for accurate viscous liquids. *ACM Trans. Graph.*, 36(4):101:1–101:11, July 2017.