

Taming Liquids for Rapidly Changing Targets

Lin Shi Yizhou Yu
University of Illinois at Urbana-Champaign

Abstract

Following rapidly changing target objects is a challenging problem in fluid control, especially when the natural fluid motion should be preserved. The fluid should be responsive to the changing configuration of the target and, at the same time, its motion should not be overconstrained. In this paper, we introduce an efficient and effective solution by applying two different external force fields. The first one is a feedback force field which compensates for discrepancies in both shape and velocity. Its shape component is designed to be divergence free so that it can survive the velocity projection step. The second one is the gradient field of a potential function defined by the shape and skeleton of the target object. Our experiments indicate a mixture of these two force fields can achieve desirable and pleasing effects.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation; I.6.8 [Simulation and Modeling]: Types of Simulation—Animation

Keywords: Fluid Simulation, Fluid Control, Feedback Forces, Geometric Potential

1 Introduction

Recent progress on physics-based fluid simulation in the graphics community has produced stunning dynamics comparable to their counterparts in the real world. Nevertheless, animators not only need realism, but also need to achieve certain design goals when producing animations. In the past, liquid actors and animals, such as those in "The Abyss" and "Terminator 2: Judgment Day," have been created with realistic rendering, but without the desirable liquid dynamics. Recent techniques on liquid control [McNamara et al. 2004; Rasmussen et al. 2004] have incorporated physics-based fluid dynamics. But they are not particularly convenient to use. Fluid control methods should be evaluated with the following criteria:

- **Control capability.** When given a set of user-specified constraints, such as key frames or target shapes, a control method should be able to force the movement of the fluid to approximately satisfy the constraints.
- **Ease to use.** As an authoring tool, the method should be able to produce desirable fluid animations without too much computational cost or user intervention.
- **Fluid-like motion.** The control method should not overconstrain the fluid. The natural movement of the fluid should be preserved as much as possible.
- **Stability.** The controlled movement of the fluid should be stable without obvious oscillations.

In this paper, we focus on controlling liquids to match rapidly changing target shapes which represent regular non-fluid objects. The physical properties of liquids and the spatially varying large



Figure 1: A controlled simulation with a water horse emerging, running and collapsing.

velocities and accelerations over the target shape pose challenges on the control method. First, the large magnitude of the velocities and accelerations of the target shape demands a stable and well-balanced fluid control method to track the fast-changing global shape while maintaining natural fluid motion. Second, when compared to gases, liquids have a clear boundary at their interface with the air and their volume is strictly incompressible. These properties imply that undesirable artifacts introduced by a control method can be easily noticed.

With the unique challenges posed by the aforementioned problem, we present a simple but effective control technique. The basic idea is to apply feedback forces (Section 3) in addition to a potential field that is induced by the geometry of the target object (Section 4) to reduce the amount of discrepancy between the fluid shape and the target shape at each frame of an animation. However, only considering differences between the two shapes becomes inadequate for rapidly changing targets and undesirable oscillations of the fluid around the target shape become obvious. As a result, our method applies additional feedback forces to compensate for differences between corresponding velocities as well to prevent excessive overshooting. Our fluid control method was originally inspired by the optimal control theory in [Bullo 2000] which is based on both potential shaping and feedback control.

Our control method meets all of the aforementioned criteria. It is capable of tracking rapidly changing targets while maintaining fluid-like motion. Given a target animation, computing the spatially varying feedback and potential forces are automatic and fast. There are only three tunable parameters representing the overall strengths of the force fields. Thus, our control technique does not incur much extra computation and user intervention beyond a regular fluid simulation. As a result, it typically only takes a few hours to author a fluid animation with controlled behavior.

2 Related Work

This work has been made possible by much of the previous work on fluid simulation [Stam 1999; Fedkiw et al. 2001; Foster and Fedkiw 2001; Enright et al. 2002; Goktekin et al. 2004; Carlson et al. 2004], and the level set method [Sethian 1999; Osher and Fedkiw 2001]. In particular, our liquid simulation follows the implementation in [Enright 2002]. In the following, we are going to focus on fluid control methods which are most relevant to this paper.

Foster and collaborators initiated fluid control in graphics. Foster and Metaxas [1997] proposed embedded controllers that allow animators to specify control parameters in a fluid animation. In [Foster and Fedkiw 2001; Lamorlette and Foster 2002], animator-designed space curves and surfaces can be applied to control the motion and structures of fluids. The tangents of the curves or the normals of the surfaces indicate the directions of motion. In light of this approach, a control method based on relatively dense particles (one particle per grid cell) [Rasmussen et al. 2004] has produced impressive results. The control particles can not only control the velocity of the fluid in their local neighborhoods, but also control viscosity, the level set of the liquid surface, and the velocity divergence. Although the results being production quality, the keyframes for all the particles need to be generated from an animation package, which involves a significant amount of work from an animator, including tuning spatially varying velocity, viscosity, and shape. In addition, at regions with dense control, natural fluid motion may be compromised. These techniques do not allow the user to directly specify higher-level objectives, such as matching the density or shape of a target object.

Recently, a systematic and inspirational optimization framework for controlling both smoke and liquid simulations through user-specified keyframes was introduced in [Treuille et al. 2003; McNamara et al. 2004]. By viewing fluid simulation as a composition of functions, simulation derivatives can be surprisingly obtained from the fluid solver itself. However, since the derivatives with respect to each control parameter need to be computed throughout an entire simulation, this approach is very expensive. The adjoint method was adopted in [McNamara et al. 2004] to significantly improve the efficiency of these derivative evaluations for each iteration during the nonlinear optimization. Nevertheless, a number of iterations are still necessary to obtain a good solution. As a result, this framework is not particularly animator friendly since the optimization takes a long time and can only be performed on relatively coarse grids. Furthermore, the objective function of the optimization forces the fluid to tightly match the target density, sacrificing its natural appearances. The method introduced in this paper is much more efficient and can handle grids with two orders of magnitude more voxels. In addition, our method can maintain much of the natural motion of fluids.

An efficient and novel technique to match smoke density against user-specified distributions was reported in [Fattal and Lischinski 2004], which involves a custom-designed driving force term and a smoke gathering term. It has demonstrated much shorter computing time than nonlinear optimization. However, the driving force term requires a nonuniform density distribution for its gradient estimation. It is not obvious how to generalize the force term for liquids since liquids have a uniform internal density. Meanwhile, another efficient method based on the level sets of the smoke was introduced in [Shi and Yu 2005]. It can effectively control the smoke shape without the forcing and gathering terms in [Fattal and Lischinski 2004]. However, it relies on a compressible fluid model to eliminate visual artifacts caused by the constraints on the smoke boundary. Such a compressible model becomes inappropriate for liquids which are strictly incompressible. The method in this paper applies a divergence-free feedback force field for shape differences to avoid this problem. It also applies additional velocity feedback forces to




Figure 2: The initialization of the shape feedback force on the liquid interface. The initial force \mathbf{f}_1 at \mathbf{x}_1 outside the target object is defined according to (3); \mathbf{f}_2 at \mathbf{x}_2 inside the target is defined according to (4).

achieve stability.

Yet another simple and clever control scheme for smoke simulation was presented in [Hong and Kim 2004], which exploits a potential function based on the shape of the target object. The negative gradient of the potential function serves as an extra force field which tries to reduce the overall potential of the smoke region. Upon its convergence, the shape of the smoke region coincides with the target shape, which represents the lowest energy configuration of the potential function. The original paper [Hong and Kim 2004] focuses on smoke and only has results for static target shapes. When the target shape has complex and rapidly changing motion, it is unclear how to keep the liquid in pace with the target. In addition, when the configuration of the liquid is inconsistent with the target shape, undesirable oscillations around the target shape become obvious because of the inertia of the liquid. The method in this paper overcomes these problems.

3 Feedback Control Forces

During a controlled liquid animation, our algorithm applies a feedback control force field to reduce the shape discrepancy between the liquid and the target. Since the target surface may have large accelerations, control forces only based on shape differences become insufficient to guarantee both convergence and natural motion. Inspired by control theories, such as [Haugen 2004; Bullo 2000], we decided to apply feedback forces to compensate for velocity differences as well. As a result, the complete feedback force has two terms,

$$\mathbf{f}_{feedback} = \mathbf{f}_{shape} + \mathbf{f}_{velocity}, \quad (1)$$

where \mathbf{f}_{shape} compensates for shape differences and $\mathbf{f}_{velocity}$ compensates for velocity differences. In control terminology, \mathbf{f}_{shape} corresponds to proportional control and $\mathbf{f}_{velocity}$ corresponds to derivative control. There is no obvious definition of a force field corresponding to integral control in the context of shape matching. Therefore, we replace the role of integral control with a geometric potential field which will be discussed in Section 4.

At any moment t , the liquid boundary and target shape are represented by two signed distance functions, $d_L(\mathbf{x}, t)$ and $d_T(\mathbf{x}, t)$, respectively. $d_T(\mathbf{x}, t)$ is positive when \mathbf{x} is outside the target shape, and $d_L(\mathbf{x}, t)$ is positive when \mathbf{x} is outside the liquid volume.

3.1 Velocity Feedback

To effectively apply this feedback force, we design a force field throughout the liquid volume. At any point \mathbf{x} either inside the liquid volume or on the liquid boundary, the velocity feedback force is simply

$$\mathbf{f}_{velocity}(\mathbf{x}) = -\beta (\mathbf{v}_L - \mathbf{v}_T), \quad (2)$$

where β is the gain for derivative control, \mathbf{v}_L represents the velocity of the liquid at \mathbf{x} and \mathbf{v}_T represents some target velocity at the same point. We use $\beta = 25$ in our experiments.

Ideally, the target velocity, \mathbf{v}_T , should be the velocity of the target shape. For any point at the interior of the target shape, this velocity can be uniquely determined. However, when a point on the liquid boundary is outside the target shape, the velocity of the target shape at that point is not well defined. Nevertheless, the signed distance function of the target shape is computed everywhere, and it moves with the target shape. We use the signed distance function to propagate the velocities of the target shape to all the points in the work space. Thus, a point outside the target shape is assigned the velocity of the closest point on its surface. In our experiments, such a propagation produced desirable target velocities.

3.2 Shape Feedback

This feedback force only concerns the differences between the target boundary and the liquid boundary. However, the effect of such a force field can be easily diminished by the projection step [Stam 1999; Foster and Fedkiw 2001] which forces the entire velocity field to be divergence-free. That means the force field needs to be divergence-free to survive the projection step. To obtain such a force field, we take the following steps. We first initialize the forces on the boundary. Further optimization is performed to guarantee that the flux of the forces on the boundary is zero, which is a necessary condition for the liquid to be volume-preserving. At the end, a divergence-free force field at the interior of the liquid can be solved by considering the forces on the boundary as the boundary condition.

Force Initialization on the Liquid Boundary. Let us focus on points on the liquid boundary. As shown in Fig. 2, when such a point \mathbf{x} is outside the target shape, the initial force at this point is formulated as

$$\tilde{\mathbf{f}}(\mathbf{x}) = -\alpha d_T \frac{\nabla d_T}{\|\nabla d_T\|}; \quad (3)$$

otherwise, the initial force is formulated as

$$\tilde{\mathbf{f}}(\mathbf{x}) = \alpha d_L \frac{\nabla d_L}{\|\nabla d_L\|}, \quad (4)$$

where α is the gain for proportional control. Note that the force follows the opposite direction of ∇d_T when the point is outside the target so that the liquid can be pulled toward the target faster. For the portion of the liquid surface that is inside the target shape, faster expansion is achieved by using the normal direction of the liquid surface. In our experiments, we usually set $\alpha = 625$.

Force Optimization on the Liquid Boundary. On a discrete grid, let $\{\mathbf{n}_i\}_1^m$ be the normals of the m surface faces on the liquid boundary, and $\{\mathbf{f}_i\}_1^m$ be the finalized feedback forces on the boundary. To ensure that the total flux of these forces is zero, we impose the following constraint,

$$\Phi_{\mathbf{f}} = \sum_i \mathbf{f}_i \cdot \mathbf{n}_i = \sum_i \sum_{j=1}^3 f_i^j n_i^j = 0, \quad (5)$$

where $\mathbf{f}_i = [f_i^1 \ f_i^2 \ f_i^3]^T$ and $\mathbf{n}_i = [n_i^1 \ n_i^2 \ n_i^3]^T$. The area of the faces has been eliminated from (5) since it is the same for all the grid cells.

Let $\{\tilde{\mathbf{f}}_i\}_1^m$ represent the initial forces on the liquid boundary. We would like to adjust such an initialization by minimizing the following objective function while satisfying (5),

$$\sum_i \|\mathbf{f}_i - \tilde{\mathbf{f}}_i\|^2, \quad (6)$$




Figure 3: A potential field is defined with respect to the target object's shape and skeleton. The skeleton is assigned a distance value d_0 while the surface of the target is assigned a distance value d_1 with $d_0 < d_1$. The potential at any point \mathbf{x} within the liquid volume is defined according to (10) and (9).

where the squared differences measure the deviation from the initialization. Such a constrained minimization problem can be solved by introducing a Lagrange multiplier for (5). Suppose $\Phi_{\tilde{\mathbf{f}}} = \sum_i \tilde{\mathbf{f}}_i \cdot \mathbf{n}_i$. The final solution of the boundary forces is simply as follows.

$$\mathbf{f}_i = \tilde{\mathbf{f}}_i - \frac{\Phi_{\tilde{\mathbf{f}}}}{m} \mathbf{n}_i, i = 1, \dots, m. \quad (7)$$

The Complete Shape Feedback Force Field. Once the flux of the forces on the liquid boundary is zero, it becomes feasible to solve a divergence-free force field throughout the liquid volume. We formulate the complete shape feedback force field as the gradient field of a scalar function, H . Thus, $\mathbf{f}_{shape} = \nabla H$.

To ensure zero divergence everywhere, $\nabla \cdot \nabla H$ has to be zero everywhere in the liquid. Thus, H is a harmonic function [Axler et al. 2001] which is the solution to the following Laplace equation with a boundary condition specified by the gradients on the boundary,

$$\nabla^2 H = 0, \quad \nabla H|_{\partial\Omega} = \mathbf{f}^*|_{\partial\Omega}, \quad (8)$$

where $\partial\Omega$ denotes the liquid boundary, and \mathbf{f}^* represents the boundary control forces obtained from the previous section. The discretization of (8) on a volume grid gives rise to a sparse linear system which can be efficiently solved by a preconditioned conjugate gradient method [Golub and Loan 1996].

4 Adaptive Geometric Potential

Controlling smoke simulation with the help of a potential field defined by the geometric shape of a target object is first investigated in [Hong and Kim 2004]. In their method, the interior of the target object is assigned with a uniform minimal potential value, and points outside the target have higher potential values. The negative gradient of the potential field, which is zero inside and nonzero outside the target, served as the driving force. The smoke converges to the target shape when the equilibrium with minimal potential energy has been reached.

In this paper, we generalize this method to liquids and design an effective potential field for them. The general form of the potential field we adopt is simply a monotonically increasing function of the signed distance of the target shape,

$$\phi(\mathbf{x}) = C \operatorname{sgn}(d_T(\mathbf{x})) |d_T(\mathbf{x})|^\gamma \quad (9)$$

where C is a constant factor representing the overall strength of the potential field. $C = 1300$ and $\gamma = 2$ in most of our experiments. The reason for this formulation is that the gradient of this new potential field has increasing magnitude when the distance to the target surface becomes larger. That means the potential field generates a




Figure 5: Two frames from a two-dimensional controlled fluid simulation with a rotating star shape as the target motion. The 2D grid resolution is 1000^2 , and the computational time is 2.5 minutes per frame.

larger force to pull the liquid back when it is further away, but the force remains small around the target surface to produce more natural fluid motion. This is in contrast to the original signed distance which only has a constant gradient. This new potential still has the same equilibrium as before. When the target object is moving, we simply calculate its potential field at every frame.

The negative gradient of the above potential field points towards the central skeleton of the target shape. However, using a function of the signed distance as the potential field creates problems when the target has a complex shape, such as the shape of a human character whose limbs and torso have very different thickness. Thick regions can tolerate more deviation in shape than thin regions which thus demand larger gradient forces to produce faster response. For this purpose, we define an adaptive signed distance based on two isosurfaces assuming that a connected skeleton of the target shape is known. In this definition, all the points on the surface of the target shape are required to have the same distance value, d_1 , and all the points on the skeleton of the target shape are required to have another distance value, d_0 ($< d_1$). As shown in Fig. 3, the signed distance at a point \mathbf{x} inside the liquid volume is defined to be

$$d_T^a(\mathbf{x}) = d_0 + (d_1 - d_0) \frac{\|\mathbf{ox}\|}{\|\mathbf{os}\|}, \quad (10)$$

where \mathbf{o} is the closest point of \mathbf{x} on the skeleton, and \mathbf{s} is the intersection between the boundary of the target shape and the line defined by \mathbf{ox} . Thus, the potential field in (9) should be updated by replacing $d_T(\mathbf{x})$ with $d_T^a(\mathbf{x})$. With such an adaptive potential field, thin regions have a larger gradient force because the Euclidean distance between the skeleton and the surface is smaller. Another advantage of considering the skeleton as an isosurface is that it is easier for the liquid to move around and less likely to be trapped in a local region. In practice, we always set $d_0 = -1$ and $d_1 = 0$. In general, such a gradient force field does not guarantee zero divergence.

5 Implementation and Analysis

The input to our system is a continuous animation of single or multiple target shapes, which can be created from key-frames, physics-based simulations, or motion capture. The sequence of animated target shapes should be automatically converted to a sequence of implicit functions because the liquid surface is represented as an implicit function as well. The output from our system is a liquid simulation that approximately follows the target animation while still maintaining natural liquid dynamics.

Our overall control scheme exploits both types of forces introduced in the previous two sections. The feedback control force




Figure 6: Left: One liquid dolphin in the air. The resolution is $234 \times 180 \times 108$, and computational time is around 4.8 minutes per frame. Right: Three liquid dolphins jumping back to the water. The resolution is $180 \times 78 \times 180$, and computational time is 4.2 minutes per frame.

field is very powerful in making the liquid follow the target shape and motion. Our experiments have confirmed that a strong feedback force field with a reasonably large gain for velocity feedback is able to completely control the shape of the liquid so that it tightly matches the target shape without oscillations. However, such a strong control is not desirable because the liquid motion would not look natural anymore.

On the other hand, the force generated by the geometric potential field resembles the gravity in the real world. If we look at the earth from the space, the gravity actually points roughly towards the center of the earth. Due to the similarity between the geometric potential and the gravity field, when the target does not have large accelerations, the controlled fluid motion under the geometric potential is expected to be similar to that under the gravity, and therefore, appears natural. However, the control capability of the geometric potential weakens and produces undesirable oscillations when it is changing quickly from frame to frame. Therefore, to achieve both desirable control and natural fluid-like motion, we apply both force fields simultaneously, but use a relatively small feedback force to avoid artifacts. The typical overall strengths of these force fields have been given in the previous sections.

A comparison is given in Fig. 4 where the target animation is a dumbbell shape bouncing around in a box. When the target hits one of the walls of the box, it receives a large and sudden acceleration from the wall to change its moving direction. In addition, its elongated shape facilitates rotation. Thus, this target animation represents a challenge to liquid control methods. We tested our method on this sequence against the original potential-based method in [Hong and Kim 2004]. A pair of rendered frames from these two methods are shown in Fig. 4. Our result is shown in Fig. 4(b). There are obvious differences between these two results. First, the liquid in our result has a better match with the target object especially in the middle of the shape because we define the geometric potential using a central skeleton. In comparison, the shape of the liquid in Fig. 4(c) is broken in the middle. Second, more importantly, our simulation produces a much more stable animation of the liquid than the method in [Hong and Kim 2004] because our method also applies a feedback force field introduced in Section 3. Strong oscillations can be observed in the animation produced by the method in [Hong and Kim 2004]. Both animations can be found in the accompanying video. We did not render shadows to avoid any additional distraction.

According to Section 3 and 4, our control technique computes spatially varying force fields automatically. Such computations are fast, and do not incur much extra cost beyond a regular fluid simulation. As a result, we can adopt any grid resolution that a regular




Figure 4: (a) The target shape and orientation at one of the frames of the target animation. (b) A rendered image from our liquid control method. (c) A rendered image from the method in [Hong and Kim 2004]. The resolution is 300^3 , and the computational time is 3.2 minutes per frame.




Figure 7: First row: a liquid simulation following a dancing sequence. Second row: the simulated liquid surfaces rendered as meshes. The simulation grid resolution is 300^3 , and the computational time is 7 minutes per frame.

fluid simulation can use. There are only three parameters that an animator needs to interactively adjust. Namely, the gain of the shape feedback, the gain of the velocity feedback and the overall strength of the geometric potential field. Given a new target animation, the user only needs to slightly tune these parameters around the typical values to produce a desirable fluid simulation. It should be noted that a small amount of user intervention is necessary for an authoring tool since an animator should be able to choose preferable results by tuning a small number of "control knobs". Experiments indicate that our force fields are insensitive to the resolution of the simulation grid. Similar effects can be obtained at various different resolutions. Thus, a coarser grid can be used during the tuning stage to let the animator quickly see intermediate results. Therefore, it typically only takes a few hours to author a reasonable fluid animation with controlled behavior.

6 Additional Experimental Results

We have successfully applied our liquid control method to a few examples in addition to the one shown in Fig. 4. Fig. 5 shows two frames from a two-dimensional controlled fluid simulation which uses a rotating star shape as the target animation. The overall shape of the fluid boundary resembles the target object. However, there are also many intricate details indicating this is a fluid.

The left image in Fig. 6 shows one frame from a controlled fluid simulation with a dolphin jumping sequence as the target animation. When the target jumps out of the water, it also brings a portion of the water with it. The isolated portion of the water evolves into a dolphin shape. The right image in Fig. 6 shows another related simulation with three dolphin-shaped targets.

Fig. 1 and 8 show a liquid horse emerging from a water, running on the water and then collapsing back into the water. Both the simulated liquid surface (blue meshes) and the underlying target shape (yellow meshes) are shown as well. The target shape has two components, the horse shape and a horizontal plane under the horse. The horizontal plane defines the target shape for the free moving water which is supposed to have a flat surface when it stabilizes. The original horse animation is from [Sumner and Popović 2004]. Fig. 8(a)-(c) show interesting liquid dynamics when the shape of the horse is being formed. Fig. 8(d) shows a frame where the liquid has stabilized and matched the target shape. Most small-scale details on the target shape are preserved by the matching liquid interface. To generate the collapse in Fig. 8(e)-(f), we simply remove the horse shape and only leave the horizontal plane.

Fig. 7 shows another controlled liquid simulation using a MO-CAP sequence as the target animation. The MO-CAP data have very large, rapidly changing accelerations. Nevertheless, our control method was able to keep the liquid in pace with the target. Meanwhile, the simulated liquid surface (blue meshes) exhibits natural waves and other fluid motion.

In our experiments, we use an effective grid resolution up to 300^3 using a grid windowing and resizing technique similar to that in [Rasmussen et al. 2004]. The computational time is almost the same as a regular liquid simulation which costs 5-10 minutes a frame on a high-resolution grid and only around 30 seconds a frame on a coarse grid ($< 100^3$) using an AMD 3200+ processor. Computing the control force fields only costs less than 10% of the total simulation time. It should be noted that the finest grid resolution used in [McNamara et al. 2004] is $50 \times 50 \times 50$ which only has 125,000 elements. Thus, the grid in their simulations has two orders of magnitude less elements than ours.

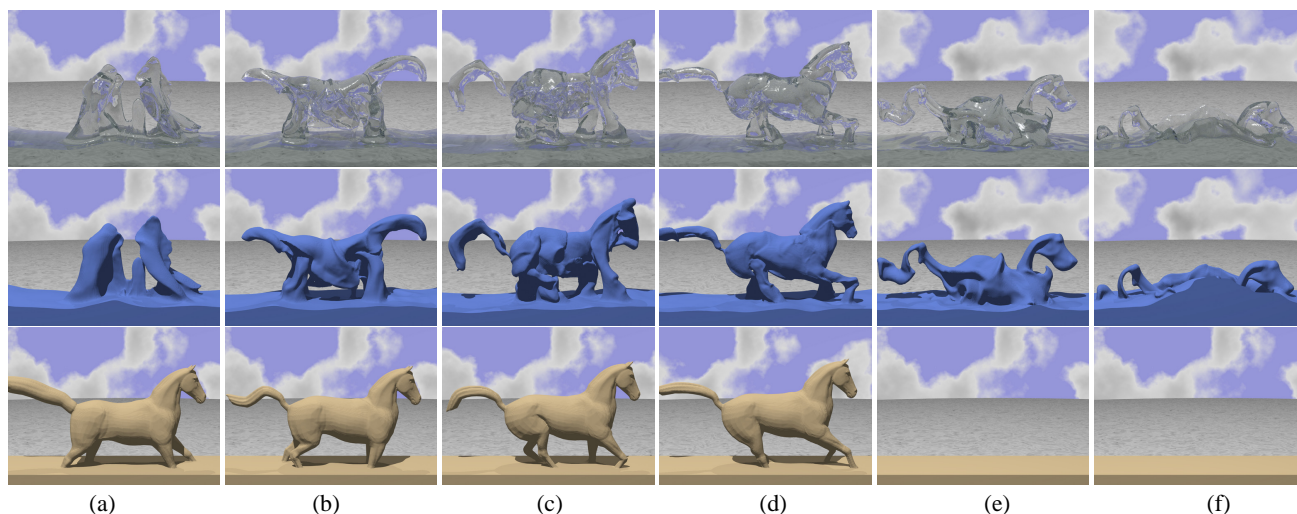


Figure 8: First row: a liquid simulation representing a water horse emerging (a & b & c), running and collapsing (e & f). Second row: the simulated liquid surfaces rendered as meshes. Third row: the underlying target animation. The resolution is 275x250x75, and the computational time is 4.4 minutes per frame.

7 Conclusions and Future Work

Following rapidly changing target objects is a challenging problem in fluid control. We introduced a simple but very effective solution by applying two external force fields: a carefully designed feedback force field and the negative gradient field of a geometric potential function. Experiments indicate that our method can achieve pleasing results. In future, we would like to extend this work to include interactions between the controlled liquid and other more rigid objects. During such interaction, interesting phenomena such as splashing should be simulated. A hybrid system based on both target shapes and control particles might be necessary for such simulations. The technique in [Carlson et al. 2004] might also be helpful.

References

- AXLER, S., BOURDON, P., AND RAMEY, W. 2001. *Harmonic Function Theory*. Springer.
- BULLO, F. 2000. Stabilization of relative equilibria for under-actuated systems on riemannian manifolds. *Automatica* 36, 12, 1819–1834.
- CARLSON, M., MUCHA, P., AND TURK, G. 2004. Rigid fluid: Animating the interplay between rigid bodies and fluid. *ACM Trans. Graphics* 23, 3, 375–382.
- ENRIGHT, D., MARSCHNER, S., AND FEDKIW, R. 2002. Animation and rendering of complex water surfaces. *ACM Transactions on Graphics* 21, 3, 736–744.
- ENRIGHT, D. 2002. *Use Of The Particle Level Set Method For Enhanced Resolution of Free Surface Flows*. PhD thesis, Stanford University.
- FATTAL, R., AND LISCHINSKI, D. 2004. Target-driven smoke animation. *ACM Transactions on Graphics* 23, 3, 439–446.
- FEDKIW, R., STAM, J., AND JENSEN, H. 2001. Visual simulation of smoke. In *SIGGRAPH 01 Conference Proceedings*, 15–22.
- FOSTER, N., AND FEDKIW, R. 2001. Practical animation of liquids. In *SIGGRAPH 2001 Conference Proceedings*, 23–30.
- FOSTER, N., AND METAXAS, D. 1997. Controlling fluid animation. In *Proceedings of Computer Graphics International*, 178–188.
- GOKTEKIN, T., BARGTEIL, A. W., AND O'BRIEN, J. 2004. A method for animating viscoelastic fluids. *ACM Trans. Graphics* 23, 3, 461–466.
- GOLUB, G., AND LOAN, C. V. 1996. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, third edition.
- HAUGEN, F. 2004. *PID Control of Dynamic Systems*. Tapir Forlag.
- HONG, J.-M., AND KIM, C.-H. 2004. Controlling fluid animation with geometric potential. *Computer Animation and Virtual Worlds* 15, 3-4, 147–157.
- LAMORLETTE, A., AND FOSTER, N. 2002. Structural modeling of flames for a production environment. In *SIGGRAPH 02 Conference Proceedings*, 729–735.
- MCMANARA, A., TREUILLE, A., POPOVIĆ, Z., AND STAM, J. 2004. Fluid control using the adjoint method. *ACM Trans. Graphics* 23, 3, 447–454.
- OSHER, S., AND FEDKIW, R. 2001. Level set methods: an overview and some recent results. *Computational Physics* 169, 2.
- RASMUSSEN, N., ENRIGHT, E., NGUYEN, D., MARINO, S., SUMNER, N., GEIGER, W., HOON, S., AND FEDKIW, R. 2004. Directable photorealistic liquids. In *Eurographics/ACM SIGGRAPH Symposium on Computer Animation*.
- SETHIAN, J. 1999. *Level Set Methods and Fast Marching Methods*. Cambridge University Press.
- SHI, L., AND YU, Y. 2005. Controllable smoke animation with guiding objects. *ACM Transactions on Graphics* 24, 1, 140–164.
- STAM, J. 1999. Stable fluids. In *SIGGRAPH 99 Conference Proceedings*, 121–128.
- SUMNER, R., AND POPOVIĆ, J. 2004. Deformation transfer for triangle meshes. *ACM Transactions on Graphics* 23, 3, 397–403.
- TREUILLE, A., MCMANARA, A., POPOVIĆ, Z., AND STAM, J. 2003. Keyframe control of smoke simulations. *ACM Trans. Graphics* 22, 3, 716–723.