



Extended virtual pipes for the stable and real-time simulation of small-scale shallow water

François Dagenais^{a,d}, Valentin Vervondel^a, Julián E. Guzmán^a, Alexander Hay^b, Sébastien Delorme^b, David Mould^c, Eric Paquette^{a,*}

^aÉcole de technologie supérieure, 1100 Notre-Dame Street West, Montreal, Quebec, Canada

^bOSSimTech, 930 Wellington Street, Suite 500, Montreal, Quebec, Canada

^cCarleton University, 1125 Colonel By Drive, Ottawa, Ontario, Canada

^dDigital District Canada, 1711 Saint-Patrick Street, Montreal, Quebec, Canada

ARTICLE INFO

Article history:

Received September 26, 2018

Keywords: Fluid simulation, Viscosity, Virtual pipes, Surface optimization, Meniscus, Shallow water

ABSTRACT

We propose an approach for real-time shallow water simulation, building upon the virtual pipes model with multi-layered heightmaps. Our approach introduces the use of extended pipes that are capable of resolving flows through fully flooded passages, which is not possible using current multi-layered techniques. We extend the virtual pipe method with a physically-based viscosity model that is both fast and stable. Our viscosity model is integrated implicitly without the expense of solving a large linear system. Despite the few simplifications necessary to achieve a real-time viscosity model, we show that our new viscosity model produces results that match the behavior of an offline fluid-implicit particle (FLIP) simulation for various viscosity values. The liquid is rendered as a triangular mesh surface built from a heightmap. We propose a novel surface optimization approach that prevents interpenetrations of the liquid surface with the underlying terrain geometry. To improve the realism of small-scale scenarios, we present a meniscus shading approach with a view-dependent adjustment of the liquid surface normals based on a distance field. Our implementation runs in real time on various scenarios of roughly 10×10 cm at a resolution of 0.5 mm, with up to five layers.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

In this paper, we focus on a real-time simulation of shallow water at small scales, such as in scenarios of spilled coffee or bleeding during surgery (Fig. 1). In such situations, thin layers of liquid flow on a surface and may also eventually fill up small cavities. At this scale, effects such as the viscous drag force exerted on the liquid by the surface of the obstacles, as well as the meniscus at the wet-dry boundary, are much more prominent.

In many real-time contexts, such as those involving medical applications and games, it is necessary to have a very efficient simulation since other systems are running on the same

resources (e.g., haptic feedback, other physics, AI). A full 3D simulation is often too expensive; only very coarse resolutions can be achieved in real time. However, a coarse 3D resolution fails to represent thin films of liquids, such as blood or paint flowing over a surface. To reduce computation times, several methods focus on performing a 2D simulation on a heightmap; these methods include those of Dagenais et al. [1], Chentanez and Müller [2], and Mei et al. [3]. Such methods can simulate liquids that are arbitrarily deep or shallow, with no impact on the resolution of the simulation.

Our approach builds upon the virtual pipes (VP) method [4]. With the exception of the work of Dagenais et al. [1], previous papers do not use a physical model to handle the viscous drag force from the terrain. This force is non-negligible for various liquids such as blood and paint. To allow more complex ter-

*Corresponding author: Tel.: +1-514-396-8587

e-mail: eric.paquette@etsmtl.ca (Eric Paquette)




Fig. 1: Frames from an animation where blood flows from a vertebra. Note the multiple overhangs and holes.

rain geometries that contain overhangs and holes, previous researchers extended virtual pipe methods to use a multi-layered heightmap [5] and created interconnections between the layers to allow the liquid to flow between them. Such configurations are important for different scenarios, such as when blood should flow below organs in a surgery simulation. Nevertheless, only the method of Dagenais et al. [1] handles the flow through fully flooded passages below obstacles, but is limited to axis-aligned passages. When passages get completely filled, previous methods stop the flow for these cells, preventing future flow while the passages remain filled. Additionally, current multi-layered VP methods have a limited surface representation, with some leading to discontinuities in the surface mesh, while others are unable to accommodate multiple overlapping layers. Finally, most related work typically aims for large-scale simulations and ignores the surface meniscus shading.

Our approach extends the work of Dagenais et al. [1], enhancing both the behavior and the shading. The behavior is improved using a physically-based viscosity model and by considering the flow below obstacles. Furthermore, our improved multi-layered surface reconstruction does not suffer from discontinuity issues. Our surface optimization approach provides a correction to the mesh surface, preventing interpenetrations with the underlying terrain geometry. We work on the simulation of moderate amounts of liquid, roughly in the 10 ml to slightly over 1 liter range. At this scale, our new meniscus shading approach significantly improves the visual results of our simulations. To summarize, our contributions are as follows:

- We present a physically-based stable viscosity model and analyze its behavior compared to an offline simulation.
- We propose two approaches allowing flow through passages of any orientation and shape, including around corners.
- We perform independent multi-layered surface reconstruction with smooth boundaries.
- We use a surface optimization to prevent unwanted interpenetrations between the heightmap and an arbitrary surface.
- We propose a view-dependent correction of the normals to enhance the specular shading of the meniscus.

2. Related work

In the computer graphics field, most fluid animation efforts concentrate on offline simulations using either an Eulerian [6], particle-based [7], or hybrid [8] simulation. Some work focuses computation on areas with more details using adaptive grid structures [9], narrow band surfaces [10], or adaptive particle radii [7]. These methods are capable of generating astonishing visual results, but are unfortunately too slow for real-time applications. Macklin and Müller [11] were able to achieve impressive visual results by simulating and rendering more than 100,000 particles in real time using their position-based dynamics framework. While this method is real-time, its ability to reproduce smooth thin layers of liquid is limited.

To reduce memory and computational costs, it is more efficient to use a heightmap to represent the fluid and perform a 2D simulation to update the liquid's height. While methods using the heightmap as such cannot exhibit some more complex behaviors such as splashes and wave crests, they are adequate for a broad range of scenarios. For example, intricate wave patterns in shallow water can be encoded as height displacements by particles [12] or packets of similar wavelengths [13]. Furthermore, heightmap methods can simulate arbitrarily thin films of liquids with no impact on the simulation resolution. Lee and O'Sullivan [14] allowed some compressibility in their 2D simulation based on the Navier-Stokes equations and adjusted the liquid's height based on its density. While simple and efficient, this technique does not account for the underlying terrain elevation. By assuming a vertical anisotropy of the liquid's velocity, the Navier-Stokes equations can be simplified, resulting in the shallow water equations, which can be further simplified to the shallow wave equations [15]. Several papers focus on implicitly solving these equations on a regular grid [15, 16] or on triangular mesh surfaces [17, 18]. Methods with an implicit integration maintain stability at larger timesteps, but are prone to a lot of diffusion as well as volume gain when using a large timestep. Furthermore, faster-moving boundaries require a smaller timestep, which can considerably increase the computation time. On the other hand, Chentanez and Müller [2] showed that their explicit integration of the shallow water equations is able to simulate large-scale scenarios in real time. For their part, our experiments show that the explicit integration requires a considerably smaller timestep for small-scale examples because of the larger ratio between the liquid's velocity and the simulation cell size, which limits its use for real-time applications in that context.

A simpler model for simulating shallow water, the VP

method, was introduced by O'Brien and Hodgins [19]. It is based on the hydrostatic pressure difference between neighbor cells of a uniform grid. Liquid is transferred between them through virtual pipes connected at their bottom, and the simulation uses an explicit integration. This method has been extended to support multi-layered heightmaps in order to allow simulations above partially submerged floating obstacles [5] and on more complex terrains with overhangs [20]. Furthermore, our experiments show that the VP method allows a considerably larger timestep than with the shallow water method of Chentanez and Müller [2] for small-scale scenarios. As such, our approach builds upon the VP method to simulate real-time, small-scale shallow waters. Dagenais et al. [1] introduced a physically-based viscosity model to the VP method. We show that this viscosity model produces results that are similar to those of a full 3D offline simulation based on the Navier-Stokes equations. Dagenais et al. [1] also introduced another novelty to the VP method: the extended pipes. These allowed flow below obstacles, but were restricted to axis-aligned passages. We improve the method of Dagenais et al. [1] by allowing extended pipes with any shape, eliminating the axis-aligned limitation.

Our goal is to simulate small amounts of liquids, with a millimeter to sub-millimeter resolution. Another feature is important at such scales: the meniscus, which is the effect of the capillary action at the fluid–solid boundary. Kerwin et al. [21] propose a real-time meniscus shading method relying on a pixel-based edge detection. Although fast, this method delivers limited realism, as it does not take into account the size of the meniscus, nor does it differentiate between a concave and a convex meniscus. The method of Dagenais et al. [1] greatly improved the meniscus shading, allowing control of the size of the meniscus, and included a physically-based curvature incorporating the contact angle. Nevertheless, their method did not account for the view direction, which then often produced spurious dark areas for convex menisci. Our approach accounts for the view direction by correctly masking these areas.

Prior VP methods often ignore effects that are visible at a smaller scale, such as the viscous drag force from the terrain and the meniscus near boundaries. Furthermore, multi-layered frameworks neglect the flow underneath fully-flooded passages below obstacles. We thus extend the VP method with a physically-based model for the viscous drag force, and propose an extended pipe model that handles the flow underneath obstacles. Furthermore, we improve the surface to account for its evolving multi-layer topology and develop a new meniscus shading approach.

3. Liquid simulation

Our approach targets the real-time simulation of small amounts of liquid based on the VP method (Sec. 3.1). Our contribution breaks down into two categories: improving fluid behavior and improving the rendering. On the behavior side, we first propose a new formulation for the viscosity (Sec. 3.2), allowing a varying amount of viscosity and widening the range of behavior from water to liquids such as blood and paint. In a further improvement to the behavior, we add the ability to handle

simulation domains containing overhangs and holes, extending the multi-layered VP method to account for the flow under obstacles through *extended pipes* (Sec. 3.3). Our rendering improvements include optimizing the surface to prevent various forms of artifacts (Sec. 4), and adjusting the surface normals to account for the meniscus (Sec. 5).

3.1. Virtual pipes simulation

This section explains how we combined different VP methods to derive our specific VP model. Our base liquid simulation follows the VP method formulation introduced by O'Brien and Hodgins [19], with the multi-layer structure of Kellomäki [5]. The VP method uses a 2D simulation grid divided into 2D cells. In turn, each cell has one or multiple columns when we extend to the multi-layered VP model. Each column corresponds to a range $[\min_i, \max_i]$ along the vertical axis, with base height b_i , liquid height h_i , and maximum range \max_i (Fig. 2). Our terrain data comes from a 3D signed distance field representing solid objects. We choose Δx for our simulation grid to be the same size as the terrain cell size so that we capture all of the details from the terrain. At initialization and as needed when the terrain changes, the column properties concerning the terrain (\min_i , b_i , and \max_i) are set by ray casting the terrain distance field upward through the center of a cell.




Fig. 2: The simulation grid is divided into cells, which in turn contain one or more columns. Here we see three linked columns i , j , and k , and column i 's properties: its range $[\min_i, \max_i]$ along the vertical axis, its base height b_i , and its liquid height h_i .

Adjacent columns are connected using virtual pipes through which they exchange liquid. In the multi-layer case, a pipe connects two adjacent columns i and j when their liquid ranges $[b_i, \max_i]$ overlap. Note that one column can be connected to multiple columns from the same adjacent cell. As in the work of Mei et al. [3], pipe connections are created only for the 4-neighborhood. These connections are created at the beginning of the simulation and adjusted when there is a change in the geometry of the obstacles. Throughout this paper, we often use the term *terrain* to refer to any obstacle (soft or rigid) lying below the columns of liquid. At each simulation step, the flux $f_{i,j}$ between a column i and its neighbor column j is updated using their difference in hydrostatic pressure and an explicit Euler integration:

$$f_{i,j}^{t+\Delta t} = \zeta f_{i,j}^t + \Delta t A \frac{g(h_i^t - h_j^t)}{l}, \quad (1)$$

where ζ is the friction parameter introduced by Mould and Yang [22], Δt is the simulation timestep, g is the gravitational

acceleration, A is the cross-section area of the virtual pipe, and l is the length of the virtual pipe. As in the work of Štáva et al. [4], we set l to be equal to the grid cell size Δx , and $A = \Delta x^2$. Note that damping is added in the simulation by scaling the amount of flux from the previous timestep that is preserved using the friction parameter ζ , instead of adding a friction force in the simulation. This approach has the advantage of being stable for $\zeta \in [0, 1]$. We want ζ to have the same effect irrespective of the time step. As such, we propose to set $\zeta = \omega \Delta t$, where $\omega = [0, 1]$ is the fraction of flux retained per unit time. In our examples, we set $\omega = 0.5$. Note that $f_{i,j} = -f_{j,i}$; thus we compute the flux once per pipe. In order to prevent the liquid height from dropping below the base height, the outflow fluxes are scaled as described by Mei et al. [3]. Similarly, to prevent the liquid height from going above the maximum height, we scale the inflow fluxes as explained by Kellomäki [5]. The scaled fluxes are then used to compute the new liquid heights, again using an explicit Euler integration:

$$h_i^{t+\Delta t} = h_i^t + \frac{\Delta t}{\Delta x^2} \sum_{j \in \text{neighbor}(i)} f_{j,i}^{t+\Delta t}. \quad (2)$$

- 1 The main steps of our simulation loop are shown in Alg. 1.
 2 Lines in bold are the steps that we added to the VP method
 3 to handle viscosity and the flow under fully-flooded passages.

Algorithm 1: Simulation loop.

- 1 **Update extended pipes connections (Sec. 3.3)**
 - 2 Update liquid fluxes
 - 3 **Apply viscosity (Sec. 3.2)**
 - 4 Scale liquid fluxes
 - 5 Update liquid heights
 - 6 Source liquid in the simulation
-

3
 4 **3.2. Viscosity model**

5 When a thin layer of liquid flows on a surface, it is slowed
 6 down by the shear stress forces from its interaction with the ter-
 7 rain. This effect is propagated further away through the liquid
 8 by the viscous shear forces. As the depth of the liquid layer
 9 increases, the impact of this interaction on the overall liquid ve-
 10 locity decreases, at a rate determined by its viscosity. Handling
 11 a wide range of liquid depths is important in many applications,
 12 including virtual surgery, for example. Hence, our goal is to
 13 handle a wide range of depths in a physically-based manner,
 14 which cannot be done by the current VP method.

Our proposed viscosity model is based on the Navier-Stokes equations, which are based on the velocity. The VP method works with the flux $f_{i,j}$ instead of the velocity $u_{i,j}$; These are related as follows:

$$f_{i,j} = \frac{u_{i,j}}{C}, \quad (3)$$

15 where $C = \Delta x(h - b)$ is the cross-sectional area of the liquid's
 16 flow from one column to its neighbor. Note that $u_{i,j}$ is a scalar
 17 that represents the speed of the flow along the pipe orientation.
 18 For simplicity we assume that the vertical axis is z ; pipes are

19 thus aligned with either the x - or y - axis. Because $u_{i,j}$ lies be-
 20 tween adjacent columns, it can be interpreted as a staggered
 21 grid, where $u_{i,j} = u_x$ or u_y , depending on its orientation.

The VP method assumes that the liquid's properties are constant along the vertical axis. However, because viscous forces are computed from the spatial differences of the velocity, this assumption would result in no shear viscous forces along the vertical axis. For that reason, we define a velocity profile that varies vertically $u_{i,j}^p(z)$, and interpret $u_{i,j}$ as the average vertical velocity. We derive the relationship between $u_{i,j}^p(z)$ and $u_{i,j}$ in Appendix A, which leads to:

$$u_{i,j}^p(z) = -\frac{3u_{i,j}}{H^2} \left(\frac{z^2}{2} - Hz \right), \quad (4)$$

where H is the liquid's depth ($h_i - b_i$) of the column from which the flux originates. Using this vertically varying velocity, we derive (Appendix B) the effect of the viscosity on the average velocity $u_{i,j}$:

$$u_{i,j}^{n+1} = \left(\frac{H^2}{H^2 + 3\Delta t\nu} \right) u_{i,j}^n,$$

where ν is the kinematic viscosity of the liquid. Using Eq. 3, we can convert from velocity to flux:

$$f_{i,j}^{n+1} C^{n+1} = \left(\frac{H^2}{H^2 + 3\Delta t\nu} \right) f_{i,j}^n C^n.$$

By assuming that the liquid depth remains the same during the timestep, we get $C^n = C^{n+1}$:

$$f_{i,j}^{n+1} = \left(\frac{H^2}{H^2 + 3\Delta t\nu} \right) f_{i,j}^n. \quad (5)$$

At each step, this equation is used to enforce viscosity on the fluxes computed from Eq. 1.

22 In Eq. 5, the right-hand side coefficient is guaranteed to be
 23 in the $[0, 1]$ range for $0 \leq 3\Delta t\nu$ and $0 < H$. This means that
 24 no energy is added in the simulation, and thus it remains stable
 25 even for arbitrarily small or large kinematic viscosity values.
 26 Fig. 3 shows how the coefficient of Eq. 5 varies for a repre-
 27 sentative range of values of H and ν . The behavior can be roughly
 28 characterized as follows. As the liquid depth H increases, the
 29 coefficient of the right-hand side gets closer to 1, lessening the
 30 impact of viscosity. Similarly, the viscosity ν damps the fluxes
 31 more severely as H becomes smaller. This shows that our vis-
 32 cosity model is stable and produces the intended behavior. Fur-
 33 thermore, it is simple and fast to compute, making it ideal for
 34 real-time purposes.
 35
 36

37 **3.3. Multi-layer and extended pipes**

38 The multi-layer VP method can only exchange liquid be-
 39 tween neighbor columns. If there is a passage through or be-
 40 low obstacles, then as soon as the level of one column of the
 41 passage (Fig. 4 left, p_{in}) reaches its maximum, the VP method
 42 will block the flow of liquid for that column. To allow liquid
 43 flow inside a passage, we identify fully-flooded passages and
 44 change the connections using extended pipes to link both ends
 45 of the passage (Fig. 4, right). Thus, instead of only having con-
 46 nnections between neighboring columns, we will also connect




Fig. 3: A 3D graph showing how the viscosity coefficient $\gamma = \frac{H^2}{H^2 + 3\Delta v}$ varies for $v \in [0, 0.1]$ and $H \in [0, 0.005]$




Fig. 4: Left: The VP method prevents flow between b_{in} and b_{out} (it is blocked at p_{in}). Right: Our approach permits flow between b_{in} and b_{out} using an extended pipe.

the ends of fully-flooded passages. A fully-flooded passage consists of a group of consecutive connected columns having a level of liquid equal to their maximum limit. For efficiency, we restrict the search for consecutive cells along two orientations, namely, alignment with the local x - and y - axes of the grid. Once a fully-flooded passage is identified, we use an extended pipe to make a connection between the columns at the boundary (b_{in} and b_{out}). The flux between these columns is initialized to zero. Once we have adjusted the pipe connections of the columns on both sides of the passage, the standard VP method is used with the new connections. At each timestep, extended pipes are connected and disconnected as needed based on the identified fully-flooded passages. In Sec. 7, we investigated a different approach which removes the axis-aligned limitation of the extended pipes.

4. Liquid surface

There are several requirements for the liquid surface: real-time generation and rendering, support for multi-layer surfaces, and flexibility to use an arbitrary geometry for the obstacles' surface. To render the liquid surface, Borgeat et al. [20] displace the terrain using the liquid height in the nearest column, and adjust the vertex colors and normals. This is problematic as the meshes of the obstacles are unlikely to have the appropriate resolution and uniformity to correctly represent the liquid.

Furthermore, discontinuities are introduced at overhangs where the obstacle meshes of both levels are not connected to each other. Fig. 5 and the accompanying video show such problems on the surface near fully submerged overhangs. Kel-




Fig. 5: Surface seam problems from the method of Borgeat et al. [20].

lomäki [5] displaces the vertices of a regular mesh grid based on the height of the topmost columns, but this method assumes a single continuous surface over the whole simulation domain, which is not the case in multi-layer scenarios. Considering the limitations of current methods, we devised a new surface creation approach which detects links between adjacent columns (Sec. 4.1), and connects them using an optimal triangle configuration (Sec. 4.2). Furthermore, we handle boundaries of the liquid to ensure a coherent surface (Sec. 4.3). Finally, we propose an optimization approach to prevent conflicting intersections between the liquid surface and the surface of the obstacles (Sec. 4.4). Alg. 2 presents the steps for the preparation of the surface.

Algorithm 2: Steps related to the preparation of the surface.

- 1 Surface links (Sec. 4.1)
 - 2 Surface creation (Sec. 4.2)
 - 3 Surface boundaries (Sec. 4.3)
 - 4 Surface optimization (Sec. 4.4)
-

4.1. Multi-layered surface links

When handling multiple layers, the surface of the liquid gets more complex: adjacent cells could have a different number of columns, and parts of the obstacle geometry could exhibit overhangs. This increases the complexity of the surface creation. It is thus important to derive a robust approach that can handle all cases. In our approach, neighbor columns that will be linked by the liquid mesh surface are identified using a multi-layered link test (Fig. 6). With this test, columns are linked if their re-




Fig. 6: Neighbor columns are linked if their liquid height lies within one another's range $[\min, \max]$.

spective liquid heights fit between each other’s minimum and maximum heights. As such, we define two neighbor columns i and j as linked if: $\min_j < h_i < \max_j$, and $\min_i < h_j < \max_i$, as shown in Fig. 6. In contrast to the pipe connections, the surface links are computed for the 8-neighborhood of each cell. They are computed for each pair of adjacent wet or wet-dry columns. A dry column is defined as a column where $h_i \leq b_i$. This approach works in the general case, while a few special cases occurring at the boundary are handled differently (Sec. 4.3). The surface links change over the course of the simulation and are verified at each frame (Fig. 7). In the next section, we explain




Fig. 7: Multi-layered linking. The thick lines show linked columns. On the left: columns a and b are initially linked. On the right: as the liquid height of b increases, it becomes linked with c .

our approach to generating the surface from the surface links.

4.2. Surface creation

Our work extends the idea of displacing the vertices of a regular mesh grid to multi-layered simulations. At the beginning of the simulation, a 3D vertex is allocated for each column, and matches the 2D coordinates of its column. Then, at each frame, its vertical position is set to match the liquid’s height. Afterward, these vertices are used to create triangular faces across adjacent linked columns (based on the link test from Sec. 4.1). To that end, quartets of neighbor cells $\{(i, j), (i + 1, j), (i, j + 1), (i + 1, j + 1)\}$ are iterated. For each quartet of cells, we analyze the links between their columns to find their corresponding triangle configuration, as shown in Fig. 8. First, we find the groups of




Fig. 8: Possible triangle configurations between 4 or 3 linked columns inside a quartet. The black points correspond to column centers.

four mutually interlinked columns (these will form two triangles). From the remaining columns, we then identify the groups of three mutually interlinked columns (these will form a single triangle). Finally, the remaining groups of two mutually interlinked columns are discarded as they do not correspond to a triangle. For groups of four linked neighbor columns, two triangle configurations are possible. In such cases, we pick the configuration for which the sum of the liquid height of the diagonal endpoints is the greatest, as described by Chentanez and Müller [2], in order to align the diagonal edge of the triangles with the curvature of the surface.

Before rendering, the normals are computed from the liquid

heights, and the vertex opacity o_i is adjusted to the liquid depth:

$$o_i = \min\left(\frac{h_i - b_i}{\text{depth}_{\max}}, 1\right),$$

where depth_{\max} is the depth from which the liquid starts being opaque. This enhances realism by better matching the opacity of the simulated liquid, and it improves the look of the wet-dry boundary on flat and convex surfaces.

4.3. Boundaries

In order to render a smooth liquid boundary and avoid having cracks at the junction with the terrain, the surface requires special treatment at the boundary of the liquid. Wet columns can be linked with dry columns having a base height (and thus liquid height) significantly higher than the liquid height of the wet column. Blindly creating triangles to the liquid height of dry columns would result in unrealistic slopes on the liquid surface near dry columns (Fig. 9, left). The vertex height of dry




Fig. 9: Bowl filled with liquid without adjusting boundaries (left), and with boundary adjustments (right)

columns is thus adjusted to the average height of their linked wet neighbors, and the surface normals of dry columns are adjusted in a similar fashion. With these corrected vertices and normals, the surface is smoother and appears as expected (Fig. 9, right). Also, during the triangle generation step, the triangle configuration that allows the ends of the diagonal edge to be both inside or both outside the boundary is prioritized.

In some cases, generally below overhangs, the liquid surface in a wet column i might be next to a wall boundary belonging to a neighbor column j to which it is not linked, i.e., $\min_j < h_i < \max_j$, but $\max_i < b_j$. This results in a crack between the wet column and the wall boundary. To prevent such cracks, we create an extra vertex positioned at the center of column j . This new vertex is linked to column i , as well as to its wet neighbor columns that are linked with column i . Its height and normal are then set to the average of these linked neighbor columns. The triangle generation step is triggered anew for all such new boundary links.

4.4. Surface optimization for rendering

The terrain distance field we get as input can have any orientation, which means that in most cases it is not aligned with the simulation grid. Terrain vertices, resulting from applying the marching-cubes algorithm on the terrain distance field, can thus appear anywhere within a simulation cell, which often leads to protrusions through the surface of the simulated liquid




Fig. 10: A comparison without (left) and with (right) our minimum height correction of the liquid surface heights.

(Fig. 10, left). Such issues are avoided by correcting columns with thin layers of liquid; to prevent surface interpenetrations and z-fighting, any non-zero liquid height is forced to be above a computed minimum height (Fig. 10 right). This minimum height h_k^{\min} is a combination of a global parameter to avoid z-fighting and a locally-computed height to avoid interpenetrations. The global parameter is set to $0.05\Delta x$ in our examples. By contrast, the locally-computed height is different for each column: it is precomputed from the terrain mesh, and adjusted when the terrain changes.

The goal of the local height optimization is to compute a minimum height for each column to ensure that all triangles of the liquid surface will be above the terrain mesh vertices, even when the amount of liquid in a column is very small. We iterate through each vertex of the terrain mesh surface, identifying the quartet of four columns surrounding the vertex, and optimizing the local heights to make them as small as possible, under the constraint that the triangles formed by these columns are above the terrain vertex. We express the height of the liquid surface, h_l , along the vertical axis through a terrain vertex as the bilinear interpolation of the heights of the related quartet. The height of each column is the sum of the local height l_k and the base height b_k . We optimize the local heights l_k under the constraint that the bilinear interpolation h_l should be slightly above the terrain vertex h_v :

$$\min_k \sum_k (w_k l_k^2), \quad \text{subject to } h_v + \epsilon = \underset{k \in \text{quartet}}{\text{bilinear}}(b_k + l_k), \quad (6)$$

where w_k are weights assigned to each column, which will be discussed later in this section. This constrained equation is solved analytically using the Lagrange multipliers. As such, the approach is efficient as it does not require that a system of equations be solved for each terrain vertex. The final l_k for a column is the maximum over the l_k computed for all terrain vertices. Once the local heights are computed, the local minimum height h_k^{\min} of each of the four columns can be computed as $h_k^{\min} = b_k + l_k$. The accompanying video shows a visualization of the impact of the position of a terrain vertex on the local minimum heights of a quartet.

The solution of Eq. 6 can create unnatural bumps on the surface, generally in regions near a steep slope on the terrain, where the differences between base heights in the quartet are quite large. We prevented such issues by designing weights w_k based on the distance between the column base height b_k and

the vertex height h_v :

$$w_k = \begin{cases} \frac{\Delta x}{h_v - b_k}, & \text{if } b_k \leq h_v \\ 1E10, & \text{otherwise} \end{cases} \quad (7)$$

These weights enforce a greater correction to the columns with a base height at the bottom of a large slope, and limit corrections to those already above the terrain vertex height.

Terrain vertices whose normal is facing downward are skipped. Furthermore, we discard terrain vertices that create a height h_k^{\min} higher than a cell size Δx above the highest base height of the quartet. Such cases usually generate unrealistic bumps near edges of large slopes. Finally, large corrections can be induced by terrain vertices that are extremely close to one of the four columns. To prevent such large corrections, we take the global minimum into account during the calculations of the local minimum height, constraining the l_k to be greater than or equal to the global minimum height.

5. Meniscus shading

The capillary action at the fluid–solid interface causes the fluid surface to curve near the borders, which affects the surface normals and results in specular highlights at the boundary. We rely on an inexpensive correction of the normal to provide the desired visual effect, similar to bump mapping (Fig. 11). First, we identify the columns of liquid affected by the meniscus.




Fig. 11: Example of a rendering without meniscus (on the left), and with our approach (on the right)

This is done by computing a distance field to the boundary on the surface of the liquid, within the given meniscus range. For each meniscus column, we also obtain the direction to the boundary, as well as an estimate of the terrain inclination at the boundary. This information is needed for our last step, where we tilt the normals of the meniscus columns to emulate the expected curvature, as depicted in Fig. 12.




Fig. 12: Correction of the normals on the liquid surface based on the distance to the boundary (left: original normals, middle: expected curvature, right: corrected normals)

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

5.1. Distance and direction to boundary

To identify the region containing the meniscus, we use the 2D distance and direction to the Nearest Boundary Column (NBC), i.e., the closest column on the “dry” side of the fluid–solid boundary. We proceed in a flood-fill manner from the NBC inward to the wet columns. We start from the *first ring* of wet columns, i.e., wet columns linked to a dry column or an extra vertex below an overhang (Sec. 4.3). We use these links to initialize the NBCs and distances of the first ring. We then do several iterations to propagate the information from the boundary. The number of iterations is automatically set to cover the meniscus region based on the meniscus length and the cell size. During the iterations, for the current wet column, we consider its linked columns and use the Dead Reckoning algorithm [23], a variant of the Chamfer distance transform algorithm: if the NBC distance of the current column is greater than that of a linked column summed with the distance to that column, we update the current NBC and distance (with the Euclidean distance to the NBC).

We also need the direction to the NBC for each column in the meniscus region in order to orient the meniscus accordingly. To that end, we can either compute the directions to the NBCs and then filter to attenuate any aliasing from the discrete grid distance field or use the finite central difference. We choose the latter, as it gives a smoother approximation than the raw directions, while limiting the amount of computation required.

5.2. Normal correction

The contact angle α between a liquid and a solid is given by Young’s equation [24]. The normal correction angle at the contact point (i.e., the fluid–solid interface) varies with the slope of the terrain at that point, as illustrated in Fig. 13. We use the




Fig. 13: Contact angle α affects the surface differently based on the inclination of the solid, resulting in a concave (left) or convex (right) meniscus, or no meniscus (middle)

base heights of our column and of its NBC to obtain an estimate of the tilt angle of the solid β . This allows us to compute the normal correction angle $\psi = \beta - \alpha$ at the contact point (Fig. 14).




Fig. 14: From contact angle α and solid tilt angle β , we obtain the correction angle ψ to apply to the normal at the contact point.

We adjust the normals of the liquid surface for all the columns within the meniscus region. In our examples, we use a meniscus length of 2.8 mm, based on the densities of water and

air [25]. Using the distance to the NBC and the correction angle at the contact point, we find the tilting angle for each affected column, linearly interpolated from 0 at the maximal meniscus distance to the angle ψ at the interface. Finally, we rotate the normal vector accordingly, as shown in Fig. 12, using the horizontal vector perpendicular to the NBC direction.

5.3. Convex meniscus

Altering the surface normals (and keeping the same geometry) is sufficient to emulate the presence of a concave meniscus (Fig. 12), but can lead to issues for a convex meniscus. Depending on the view angle, some parts of the convex meniscus could be hidden by the surface of the meniscus itself (Fig. 15(b), from A to B). If we simply compute the normal at the contact point and perform our interpolation for the normal correction, we may obtain normals facing away from the viewer (Fig. 15(b)-(c)). These back-facing normals result in black outlines in the meniscus region (Fig. 15(a)).

Consider Fig. 15(b): the last normal the viewpoint sees on the surface is the one at point A, while the normal at the contact point B is occluded by the surface of the meniscus. To emulate this on our geometrically flat surface, we cap the normal at the contact point such that it is at most orthogonal to the view vector, corresponding to the normal at the silhouette of the meniscus (Fig. 15(b), point A). Fig. 15(d) shows the normals interpolated as in Sec. 5.2, but this time using the capped normal. This approach results in a convincing shading effect for convex menisci (Fig. 15(e)).

6. Results

We tested our approach with a wide range of scenarios to show its behavior, as well as the impact of its parameters. Results are also available in the accompanying video. Parameter values and timings can be found in Table 1. We found that dynamic viscosity values in the $[1.0E-6, 1.0E-5]$ range provide realistic results for many liquids such as water, blood, and paint. As with typical fluid simulation methods with an explicit integration, preserving the stability of the simulation requires the timestep to be adjusted according to the CFL condition. This can be seen in Fig. 16: the timestep is inversely proportional to the number of cells dividing the simulation domain.

An important contribution of our work is the viscosity model for small-scale liquid simulation (Sec. 3.2). We tested our approach with varying viscosity values (see the video and stills in Fig. 17). The images show the simulation at time $t = 3.0s$, with $\nu = 0, 4.0E-6, 4.0E-5$, and $4.0E-1 \text{ m}^2/\text{s}$. As the viscosity increases, the liquid flows more slowly on the inclined terrain. Even with a very large viscosity (Fig. 17(d)), our approach is stable.

We derive our viscosity model by applying some simplifications to the Navier-Stokes equations. One notable simplification is that we neglect the contribution from neighbor cells to the viscosity (Appendix B). This means that our model ignores the effect of the viscosity related to the lateral friction within the liquid. Nevertheless, our model considers the friction at the interface between the liquid and the terrain. While




Fig. 15: (a) With a high contact angle (i.e., a convex meniscus), interpolating to the normal at the contact point results in a noticeable black outline. (b) For this viewing angle, the part of the convex meniscus from A to B should be hidden. (c) When interpolating to the contact point normal, some normals are facing away from the view point. We cap the contact point normal to the normal at the silhouette (point A), as it is the last part of the meniscus visible from the camera. (d) Interpolating to the capped normal avoids back-facing normals. (e) The erroneous black outline is no longer visible when using the capped normals.

Example	Resolution	Δt (ms)	Δx (m)	ν (m ² /s)	Simulation		Surface		Rendering		Total	
					Max	Avg	Max	Avg	Max	Avg	Max	Avg
Surgery (Fig. 1)	200 × 200 × 5	3.00	0.0005	4.0E−6	5.43	4.86	2.11	1.10	0.66	0.34	7.04	6.30
Viscosity (Fig. 17(a))	200 × 200 × 1	3.00	0.0005	0	1.05	0.61	0.49	0.43	0.19	0.16	1.63	1.21
Viscosity (Fig. 17(b))	200 × 200 × 1	3.00	0.0005	4.0E−6	1.03	0.62	0.52	0.46	0.21	0.16	1.63	1.22
Viscosity (Fig. 17(c))	200 × 200 × 1	3.00	0.0005	4.0E−5	1.02	0.60	0.52	0.44	0.25	0.16	2.41	1.21
Viscosity (Fig. 17(d))	200 × 200 × 1	3.00	0.0005	4.0E−1	0.93	0.60	0.52	0.41	0.22	0.16	2.40	1.20
Three layers (Fig. 19)	100 × 100 × 3	9.00	0.001	4.0E−6	0.68	0.32	0.85	0.37	0.89	0.17	1.53	0.86
Passages (Fig. 20)	100 × 100 × 3	9.00	0.001	4.0E−6	0.50	0.26	0.48	0.34	1.10	0.15	2.36	0.78

Table 1: Statistics for all the examples shown in this paper: resolution ($N \times N \times \text{Layers}$), timestep Δt , cell size Δx , and kinematic viscosity ν as well as maximum/average timings (in ms) per frame for the simulation, surface generation, and rendering




Fig. 16: Given the same simulation domain, the timestep resulting in a stable simulation is inversely proportional to the number of cells




Fig. 17: A liquid flows on an inclined terrain with varying viscosity, ν , expressed in m²/s

difference results from the lack of contribution from neighbor cells. Nonetheless, with a slight adjustment of our viscosity values, our simulation can achieve an even closer match with respect to the FLIP simulation, as shown in Fig 18(c) and in the video.

We make several contributions to the optimization of the liquid surface. As shown in Sec. 4, we improve the smoothness of the boundaries, and prevent incorrect interpenetrations. To demonstrate the surface construction with multiple layers, we show a three-layer scenario with cavities and overhangs in the accompanying video and in Fig. 19. Through simulation, the surfaces of the multiple layers link the appropriate columns, even with this configuration exhibiting overhangs and a hole in the middle platform. We also validate extended pipes with mul-

our model does not capture all the features of the complete viscosity model, we can see in our animations that the model produces the behavior expected from a viscous fluid. To further demonstrate that our model captures the important effects of the viscosity, we compared our simulation's behavior with that of a full 3D FLIP simulation [26]. Fig. 18(a)-(b) and the accompanying video show results for three viscosity values. For matching viscosity values, our approach produces a similar behavior, with a slightly less viscous look. We believe that this

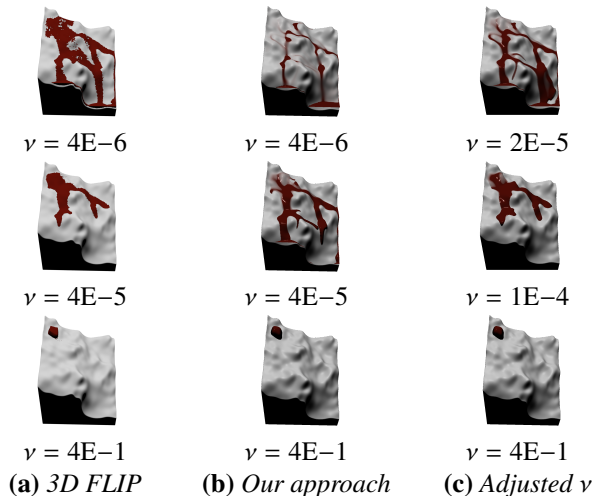


Fig. 18: Comparison between (a) a 3D FLIP simulation and (b) our approach with the same viscosity values. The behavior is quite similar, and as can be seen in (c), with slight adjustments of the viscosity, our approach can achieve a behavior even closer to that of the FLIP simulation. All images are captured at time $t = 10$ s. In this scenario, a volume of 1 cm^3 of liquid flows on an inclined surface. The 3D simulation is computed using Houdini's state-of-the-art offline FLIP solver with the same grid spacing as our simulation and an average of 8 particles per cell.




Fig. 19: This scenario shows how our surface linking approach correctly handles multiple changes in the surface topology.

multiple scenarios. The scenario with four passages presented in Fig. 20 shows how, after the passages are flooded, the liquid transfer does not stop, and is able to fill the other parts of the container. We also validate our approach with a surgery simulation scenario, including multiple overhangs and holes. In Fig. 1, blood originates on top of a vertebra during surgery. As the simulation domain is filled, our surface correctly handles the evolving topology of the liquid surface, and no holes or discontinuities are visible.

We compare our surface with that of Borgeat et al. [20] in Fig. 21 and in the video. In their technique, discontinuities in the geometry create seams on the surface near fully submerged overhangs. Furthermore, temporal artifacts can be seen in their video: the fluid surface suffers from popping near the boundary as it spreads. Our fluid surface is free of seams near the overhangs, and our approach preserves a sharp liquid/dry boundary without temporal popping.

We executed the tests presented in this paper on an Intel Core i5-6600 3.30 GHz with 16 GB of RAM and a GeForce GTX 970. To take advantage of GPU parallelism, our approach was implemented using CUDA. The timings (Table 1) show that all





Fig. 20: The liquid flows from the center and through four passages (left). When the passages are fully flooded, the extended pipes continue the transfer of liquid (right).



(a) Borgeat et al. [20]

(b) Our approach

Fig. 21: Comparison between the surface from Borgeat et al. [20] (left), and our surface (right).

of our examples run in real time. The most computationally expensive part is the simulation step, which takes around 50–80 % of the total computation time in our examples. As can be seen in Table 1 and in Fig. 22, the simulation times grow quadratically with respect to the number of cells. All of our examples take




Fig. 22: The average simulation time per frame (excluding rendering) is quadratic with respect to the number of cells.

considerably less than the 16.6 ms needed to achieve real-time rates, leaving time for other computations such as soft body simulation, haptic feedback, rendering, and collision detection.

7. Discussion

We show in the accompanying video that our viscosity model, surface reconstruction, and meniscus modeling are independent of the underlying height map simulation; it works just as well with the Shallow Water Equation (SWE) simulation [2]. While the SWE method can take advantage of our contributions, the VP method proved to be a better choice since it had

1 results of equivalent quality with significantly smaller compu-
2 tation times.

The simulation behavior is improved by the use of extended pipes, which allow the liquid to flow through fully-flooded passages. However, the extended pipes approach sometimes introduces some ripples when these pipes get connected and disconnected, even when using the scaling method of Mei et al. [3]. Furthermore, the axis-aligned extended pipes described in Sec. 3.3 block the flow of liquid for each column that does not have an axis-aligned line-of-sight to another side of the passage. As such, the axis-aligned extended pipes behave very poorly in scenarios exhibiting more complex passages such as those shown in Fig. 23. In contrast with the axis-aligned extended




Fig. 23: Examples of more complex passages that the axis-aligned extended pipes cannot handle. The liquid flow is drastically slowed down when it should turn from the axis-aligned direction. See the accompanying video for a comparison between the axis-aligned extended pipes and the unified pipe for the cross-shaped (left) scenario.

pipes that establish multiple independent connections between the columns facing each other on both sides of the passage, we experimented with another approach which consists in having a *unified extended pipe* connecting each and every boundary column b around a passage. We change the one-to-one pipe connections with a many-to-many connection for each passage. We identify the set b_P of boundary columns connected to the passage (Fig. 24), and we adapt Eq. 1 to compute the flux update




Fig. 24: This figure presents the top view of a complex passage with several branches, some of different size. Each gateway to the passage cannot establish axis-aligned connections to each of the other gateways; multiple configurations of source of liquid will result in many gateways not getting their share of liquid. We thus use a many-to-many connection to ensure that all of the boundary columns at each gateway can exchange liquid.

between each boundary column b and the passage P :

$$f_{b,P}^{t+\Delta t} = \zeta f_{b,P}^t + \Delta t A \frac{g(h_b^t - \bar{h}_P^t)}{l} \quad (8)$$

Using the unified extended pipe, we need an extra step to ensure that the fluxes all around the passage sum up to zero.

If the fluxes do not sum up to zero, which happens in most cases, we want to adjust them so that their sum is zero. This is equivalent to subtracting the average flux to each boundary flux around the passage:

$$f_{b,P}^{*t+\Delta t} = f_{b,P}^{t+\Delta t} - \bar{f}_P, \quad (9)$$

where $\bar{f}_P = \frac{1}{|b_P|} \sum_{i \in b_P} f_{i,P}$ is the average flux. These new fluxes are used in the same way with Eq. 2 to compute the new liquid heights. We validated this new approach with several scenarios presented in the accompanying video.

The surface constructed by our approach can handle the evolving topology of the liquid throughout the simulation. However, we do not create any geometry to close the gap near edges between unlinked cells (e.g., the edge of an overhang). Nonetheless, these are in areas where the liquid falls down; thus, the water height is generally low, and the gap is not very apparent. Another downside is that generally, the grid does not exactly follow the silhouette at the edge of these overhangs, which sometimes results in having a part of the geometry that will never be covered by the liquid surface. Additionally, the dry/wet boundary moves on a cell-by-cell basis, which can result in some popping. While the depth-based opacity significantly reduces this issue, it is still visible in our examples. This behavior can be improved by increasing the resolution of the grid, or by tracking the surface as suggested by Thuerey and Hess [27, Chapter 11].

In our approach, we modify only the normals and not the mesh. This may not trigger enough fragments to give an acceptable highlight, especially when viewing the surface of the fluid from the side. However, modifying the mesh would be more demanding as we would need to handle problems involving cracks and interpenetrations between the surface of the meniscus and the solid.

8. Conclusion

In this paper, we present an extension of the method by Dagenais et al. [1] for real-time fluid simulation in small-scale scenarios. We demonstrate that the viscosity model by Dagenais et al. [1] is comparable to state-of-the-art 3D offline models, and that it can be used for any heightmap simulation method, such as the VP or the shallow water equations. Our approach improves the behavior of multi-layered simulations by handling the flow inside passages. Moreover, it constructs a triangular mesh surface that accounts for the interlinks between the multiple layers of liquid. We also introduce a new surface optimization approach that computes minimum heights of the liquid surface in order to prevent interpenetration of the underlying terrain surface. Finally, we improved the shading with a novel real-time meniscus approach that proved to be useful for the VP, but can also be used for other heightmap-based simulation methods. All these contributions improve the realism of small-scale simulations. Furthermore, the approach is computationally inexpensive, which allows it to co-exist with other real-time systems normally required in the context of real applications such as surgery training and games. We have demonstrated that

the approach works in real time for various scenarios, such as blood simulation for virtual surgery.

Although we improved the surface boundaries while maintaining real-time performance, the former could be further improved by closing the gap near overhanging edges. In those regions, the liquid falling from one layer to another could be simulated using particles to improve realism, such as in the work of Chentanez and Müller [2]. Finally, we believe the computation time for the meniscus shading could be improved by using a precomputed distance field for static solids, such as in the work of Morgenroth et al. [25].

Acknowledgements

This work was supported by OSSimTech, Digital District, NSERC, Mitacs, and Prompt. We would like to thank Side Effects Software for providing Houdini licenses. The collaboration between ÉTS and Carleton University resulted from the Bellairs Multidisciplinary Computer Science Workshop held in 2017 at the Bellairs Research Institute of McGill University.

Appendix A. Velocity profile

In this section, we derive the vertically varying velocity $u_{i,j}^p(z)$. For simplicity, we assume that the base of the liquid is at $z = 0$ and that the top is at $z = H$. The function $u_{i,j}^p(z)$ is thus defined inside the liquid, i.e., for $0 \leq z \leq H$. The velocity is derived from the incompressible Navier-Stokes equations:

$$\frac{\partial \vec{u}}{\partial t} + \vec{u} \cdot \nabla \vec{u} - \nu \nabla \cdot \nabla \vec{u} = -\frac{1}{\rho} \nabla p + \vec{g}, \quad (\text{A.1})$$

$$\frac{\partial u_x}{\partial x} + \frac{\partial u_y}{\partial y} + \frac{\partial u_z}{\partial z} = 0, \quad (\text{A.2})$$

where p is the pressure, $\vec{g} = (0, 0, g_z)$ is the gravity, and $\vec{u} = (u_x, u_y, u_z)$ is the velocity. To make the velocity profile a function of only the height z inside the liquid, we show that the Navier-Stokes momentum equation (Eq. A.1) can be simplified. To begin, the first term can be removed by considering a steady flow:

$$\frac{\partial \vec{u}}{\partial t} = 0. \quad (\text{A.3})$$

We also assume the velocity is constant along the horizontal plane:

$$\frac{\partial \vec{u}}{\partial x} = \frac{\partial \vec{u}}{\partial y} = \vec{0}, \quad (\text{A.4})$$

Using this assumption, Eq. A.2 yields:

$$\frac{\partial u_z}{\partial z} = 0. \quad (\text{A.5})$$

At the terrain level, we use a no-slip boundary condition, and at the surface level, a traction-free boundary condition:

$$\vec{u}|_{z=0} = \vec{0}, \quad \frac{\partial \vec{u}}{\partial z}|_{z=H} = \vec{0}. \quad (\text{A.6})$$

The terrain boundary condition and Eq. A.5 tell us that $u_z = 0$ everywhere, and can thus be ignored. Using this result and Eq. A.4, the second term of the Navier-stokes momentum equation becomes zero:

$$\vec{u} \cdot \nabla \vec{u} = u_x \frac{\partial \vec{u}}{\partial x} + u_y \frac{\partial \vec{u}}{\partial y} + u_z \frac{\partial \vec{u}}{\partial z} \quad (\text{A.7})$$

$$= u_x(0) + u_y(0) + (0) \frac{\partial \vec{u}}{\partial z} = \vec{0}. \quad (\text{A.8})$$

Finally, the third term can be simplified similarly:

$$-\nu \nabla \cdot \nabla \vec{u} = -\nu \left(\frac{\partial^2 \vec{u}}{\partial x^2} + \frac{\partial^2 \vec{u}}{\partial y^2} + \frac{\partial^2 \vec{u}}{\partial z^2} \right). \quad (\text{A.9})$$

Using Eq. A.4, the first two terms become zero, yielding:

$$-\nu \nabla \cdot \nabla \vec{u} = -\nu \frac{\partial^2 \vec{u}}{\partial z^2}. \quad (\text{A.10})$$

After these simplifications, Eq. A.1 now becomes:

$$-\nu \frac{\partial^2 \vec{u}}{\partial z^2} = -\frac{1}{\rho} \nabla p + \vec{g}. \quad (\text{A.11})$$

Now that the Navier-Stokes equations have been simplified, we can derive the horizontal velocity as a function of z . As we use fluxes between cells, we derive the velocity along the u_x and u_y axes. We will show the derivations for u_x , and u_y can be obtained in the same way. From Eq. A.11, considering the x component, we get:

$$\nu \frac{\partial^2 u_x}{\partial z^2} = \frac{1}{\rho} \frac{\partial p}{\partial x}. \quad (\text{A.12})$$

We assume the horizontal pressure variation to be constant, and replace it by a constant w :

$$\nu \frac{\partial^2 u_x}{\partial z^2} = w, \quad (\text{A.13})$$

Integrating both sides of the equality with respect to z , and also dividing both sides by ν yields:

$$\frac{\partial u_x}{\partial z} = \frac{wz + D}{\nu}. \quad (\text{A.14})$$

Using the surface traction-free boundary condition from Eq. A.6 to find the value of D yields:

$$\frac{\partial u_x}{\partial z} = \frac{wz - wH}{\nu}. \quad (\text{A.15})$$

Integrating both sides of the equality with respect to z a second time, and rearranging the terms yields:

$$u_x = \frac{w}{\nu} \left(\frac{z^2}{2} - Hz \right) + E. \quad (\text{A.16})$$

Using the ground no-slip boundary condition from Eq. A.6, we find that $E = 0$:

$$u_x = \frac{w}{\nu} \left(\frac{z^2}{2} - Hz \right). \quad (\text{A.17})$$

To link the velocity function to the flux velocity $u_{i,j}$, we find the expression of the pressure gradient w that will make the average velocity of u_x inside the liquid equal to $u_{i,j}$:

$$u_{i,j} = \frac{1}{H} \int_0^H \frac{w}{v} \left(\frac{z^2}{2} - Hz \right) dz = -\frac{wH^2}{3v} \quad (\text{A.18})$$

$$w = -\frac{3vu_{i,j}}{H^2} \quad (\text{A.19})$$

Using that result in Eq. A.17 yields:

$$u_x = -\frac{3u_{i,j}}{H^2} \left(\frac{z^2}{2} - Hz \right) \quad (\text{A.20})$$

Performing the same derivation on u_y yields the same result. Knowing that a flux is always aligned with the x - or y - axis, u_x can be substituted by $u_{i,j}^p(z)$:

$$u_{i,j}^p(z) = -\frac{3u_{i,j}}{H^2} \left(\frac{z^2}{2} - Hz \right) \quad (\text{A.21})$$

Appendix B. Viscosity model

In this section, we derive the viscosity model, using the vertically varying velocity $u_{i,j}^p(z)$ (Appendix A). Our goal is to determine how to update the velocity, i.e., $\frac{\partial u_{i,j}}{\partial t}$, based on the viscosity. Our model accounts for the effect of viscosity from the terrain to the liquid surface, ignoring the effect of neighbor cells. While this is less accurate, neglecting the relation between neighbor cells allows us to implicitly integrate each column independently, instead of having to solve a large system of equations. We derive our viscosity model by considering only the viscous term of the Navier-Stokes momentum equation. Here, we show the derivation for $\frac{\partial u_x}{\partial t}$, but the same process can be applied to $\frac{\partial u_y}{\partial t}$:

$$\frac{\partial u_x}{\partial t} = v \left(\frac{\partial^2 u_x}{\partial x^2} + \frac{\partial^2 u_x}{\partial y^2} + \frac{\partial^2 u_x}{\partial z^2} \right). \quad (\text{B.1})$$

With the assumption of Eq. A.4, Eq. B.1 is simplified to:

$$\frac{\partial u_x}{\partial t} = v \frac{\partial^2 u_x}{\partial z^2} \quad (\text{B.2})$$

Using the velocity profile defined by Eq. A.21, this equation becomes:

$$\frac{\partial u_x}{\partial t} = v \frac{\partial^2}{\partial z^2} \left[-\frac{3u_{i,j}}{H^2} \left(\frac{z^2}{2} - Hz \right) \right] = -\frac{3vu_{i,j}}{H^2} \quad (\text{B.3})$$

This result is used to compute how the velocity $u_{i,j}$ varies over time:

$$\frac{\partial u_{i,j}}{\partial t} = \frac{\partial}{\partial t} \left(\frac{1}{H} \int_0^H u_x dz \right). \quad (\text{B.4})$$

Using the Leibniz integral rule, and simplifying the problem by assuming that the liquid depth H does not vary over time,

yields:

$$\frac{\partial u_{i,j}}{\partial t} = \frac{1}{H} \int_0^H \frac{\partial u_x}{\partial t} dz \quad (\text{B.5})$$

$$= \frac{1}{H} \int_0^H \left(-\frac{3vu_{i,j}}{H^2} \right) dz \quad (\text{B.6})$$

$$= -\frac{3vu_{i,j}}{H^2}. \quad (\text{B.7})$$

Finally, we integrate the average velocity implicitly using the backward Euler method to get the temporal update of the velocity, considering the viscosity:

$$u_{i,j}^{n+1} = u_{i,j}^n - \Delta t \frac{\partial u_{i,j}}{\partial t} \quad (\text{B.8})$$

$$= u_{i,j}^n - \frac{3\Delta tv}{H^2} u_{i,j}^{n+1}. \quad (\text{B.9})$$

Rearranging the terms yields:

$$u_{i,j}^{n+1} = \left(\frac{H^2}{H^2 + 3\Delta tv} \right) u_{i,j}^n. \quad (\text{B.10})$$

References

- [1] Dagenais, F, Guzmán, J, Vervondel, V, Hay, A, Delorme, S, Mould, D, et al. Real-time virtual pipes simulation and modeling for small-scale shallow water. In: Proceedings of VRIPHYS 2018 Workshop on Virtual Reality Interaction and Physical Simulation. The Eurographics Association; 2018.
- [2] Chentanez, N, Müller, M. Real-time simulation of large bodies of water with small scale details. In: Proc. of ACM SIGGRAPH/Eurographics Symp. on Comp. Animation. SCA '10; 2010, p. 197–206.
- [3] Mei, X, Decaudin, P, Hu, BG. Fast hydraulic erosion simulation and visualization on gpu. In: Pacific Conference on Computer Graphics and Applications, PG '07. 2007, p. 47–56.
- [4] Št'ava, O, Beneš, B, Brisbin, M, Křivánek, J. Interactive terrain modeling using hydraulic erosion. In: Proc. of ACM SIGGRAPH/Eurographics Symp. on Comp. Anim. SCA '08; Eurographics Association; 2008, p. 201–210.
- [5] Kellomäki, T. Rigid body interaction for large-scale real-time water simulation. Int J Comput Games Technol 2014;.
- [6] Enright, D, Marschner, S, Fedkiw, R. Animation and rendering of complex water surfaces. ACM Trans Graph 2002;21(3):736–744.
- [7] Ihmsen, M, Orthmann, J, Solenthaler, B, Kolb, A, Teschner, M. SPH Fluids in Computer Graphics. In: Lefebvre, S, Spagnuolo, M, editors. Eurographics 2014 - State of the Art Reports. The Eurographics Association; 2014.
- [8] Jiang, C, Schroeder, C, Selle, A, Teran, J, Stomakhin, A. The affine particle-in-cell method. ACM Trans Graph 2015;34(4):51:1–51:10.
- [9] Aanjaneya, M, Gao, M, Liu, H, Batty, C, Sifakis, E. Power diagrams and sparse pagged grids for high resolution adaptive liquids. ACM Trans Graph 2017;36(4):140:1–140:12.
- [10] Ferstl, F, Ando, R, Wojtan, C, Westermann, R, Thurey, N. Narrow band flip for liquid simulations. Computer Graphics Forum 2016;35(2):225–232.
- [11] Macklin, M, Müller, M. Position based fluids. ACM Trans Graph 2013;32(4):104:1–104:12.
- [12] Yuksel, C, House, DH, Keyser, J. Wave particles. ACM Trans Graph 2007;26(3).
- [13] Jeschke, S, Wojtan, C. Water wave packets. ACM Trans Graph 2017;36(4):103:1–103:12.
- [14] Lee, R, O'Sullivan, C. A Fast and Compact Solver for the Shallow Water Equations. In: Workshop in Virtual Reality Interactions and Physical Simulation (VRIPHYS). ISBN 978-3-905673-65-4; 2007, p. 51–57.
- [15] Kass, M, Miller, G. Rapid, stable fluid dynamics for computer graphics. Comput Graph 1990;24(4):49–57.

- 1 [16] Layton, AT, van de Panne, M. A numerically efficient and stable algo-
2 rithm for animating water waves. *The Visual Computer* 2002;18(1):41–
3 53.
- 4 [17] Wang, H, Miller, G, Turk, G. Solving general shallow wave equations on
5 surfaces. In: *Proc. of ACM SIGGRAPH/Eurographics Symp. on Comp.*
6 *Anim. SCA '07; Eurographics Association; 2007*, p. 229–238.
- 7 [18] Angst, R, Thuerey, N, Botsch, M, Gross, M. Robust and effi-
8 cient wave simulations on deforming meshes. *Computer Graphics Forum*
9 2008;27(7):1895–1900.
- 10 [19] O'Brien, JF, Hodgins, JK. Dynamic simulation of splashing fluids. In:
11 *Proc. of Computer Animation '95*. 1995, p. 198–205, 220.
- 12 [20] Borgeat, L, Massicotte, P, Poirier, G, Godin, G. Layered surface
13 fluid simulation for surgical training. In: *Proc. of Medical Image Com-*
14 *puting and Computer-Assisted Intervention – MICCAI 2011 Conference.*
15 Springer; 2011, p. 323–330.
- 16 [21] Kerwin, T, Shen, HW, Stredney, D. Enhancing Realism of Wet Surfaces
17 in Temporal Bone Surgical Simulation. *IEEE Trans on Visualization and*
18 *Computer Graphics* 2009;15(5):747–758.
- 19 [22] Mould, D, Yang, YH. Modeling water for computer graphics. *Computers*
20 *& Graphics* 1997;21(6):801–814.
- 21 [23] Grevera, GJ. The “dead reckoning” signed distance transform. *Computer*
22 *Vision and Image Understanding* 2004;95(3):317–333.
- 23 [24] Young, T. An essay on the cohesion of fluids. *Philosophical Trans of the*
24 *Royal Society of London* 1805;95:65–87.
- 25 [25] Morgenroth, D, Weiskopf, D, Eberhardt, B. Direct raytracing of a
26 closed-form fluid meniscus. *The Visual Computer* 2016;32(6-8):791–800.
- 27 [26] Zhu, Y, Bridson, R. Animating sand as a fluid. *ACM Trans Graph*
28 2005;24(3):965–972.
- 29 [27] Müller, M, Stam, J, James, D, Thuerey, N. Real time physics. In:
30 *SIGGRAPH Class Notes*. ACM; 2008, p. 88:1–88:90.