Jan Bender bender@cs.rwth-aachen.de RWTH Aachen University

Marcel Weiler weiler@cs.rwth-aachen.de RWTH Aachen University Tassilo Kugelstadt kugelstadt@cs.rwth-aachen.de RWTH Aachen University

Dan Koschier d.koschier@ucl.ac.uk University College London



Figure 1: Our novel implicit boundary representation based on volume maps is able to handle scenarios with complex static and dynamic boundaries. Left: 8 million turbulent fluid particles interact with a large-scale canyon boundary. Right: Four emitters generate 12 million fluid particles with high velocities that interact with static dragons and dynamic ducks.

ABSTRACT

In this paper, we present a novel method for the robust handling of static and dynamic rigid boundaries in Smoothed Particle Hydrodynamics (SPH) simulations. We build upon the ideas of the density maps approach which has been introduced recently by Koschier and Bender. They precompute the density contributions of solid boundaries and store them on a spatial grid which can be efficiently queried during runtime. This alleviates the problems of commonly used boundary particles, like bumpy surfaces and inaccurate pressure forces near boundaries. Our method is based on a similar concept but we precompute the volume contribution of the boundary geometry and store it on a grid. This maintains all benefits of density maps but offers a variety of advantages which are demonstrated in several experiments. Firstly, in contrast to the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MIG '19, October 28-30, 2019, Newcastle upon Tyne, United Kingdom

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6994-7/19/10...\$15.00

https://doi.org/10.1145/3359566.3360077

density maps method we can compute derivatives in the standard SPH manner by differentiating the kernel function. This results in smooth pressure forces, even for lower map resolutions, such that precomputation times and memory requirements are reduced by more than two orders of magnitude compared to density maps. Furthermore, this directly fits into the SPH concept so that volume maps can be seamlessly combined with existing SPH methods. Finally, the kernel function is not baked into the map such that the same volume map can be used with different kernels. This is especially useful when we want to incorporate common surface tension or viscosity methods that use different kernels than the fluid simulation.

CCS CONCEPTS

• Computing methodologies \rightarrow Physical simulation.

KEYWORDS

Smoothed Particle Hydrodynamics, fluid simulation, boundary handling

ACM Reference Format:

Jan Bender, Tassilo Kugelstadt, Marcel Weiler, and Dan Koschier. 2019. Volume Maps: An Implicit Boundary Representation for SPH. In *Motion, Interaction and Games (MIG '19), October 28–30, 2019, Newcastle upon Tyne,* United Kingdom. ACM, New York, NY, USA, 10 pages. https://doi.org/10. 1145/3359566.3360077

1 INTRODUCTION

Smoothed Particle Hydrodynamics (SPH) has become a powerful tool in computer graphics, where it is used for visual effects in movies, and in interactive applications like games or virtual training simulators. It can be used to simulate lots of different materials ranging from incompressible fluids and highly viscous liquids to granular materials and deformable solids.

Using particles as the primary material representation is appealing because it simplifies implementations. Moreover, it provides a straightforward way for two-way coupling of many different materials and for two-way coupling with rigid bodies. However, the accurate application of boundary conditions remains challenging, but is especially important when materials interact with complex static and dynamic boundaries as depicted in Figure 1. Over the last years several boundary representations based on particles, meshes or implicit representations, like signed distance fields, have been proposed. But most of them are either inaccurate or have high computational demands. For instance, sampling the boundary with particles results in bumpy surfaces which lead to inaccurate forces and introduce artificial friction or even jumping particle artifacts (cf. [Band et al. 2017; Koschier and Bender 2017]). A detailed discussion of previous methods and the associated problems can be found in Section 2.

Recently, Koschier and Bender [2017] proposed an implicit boundary representation called *density maps*. They use a spatial grid to discretize the density contribution function of static and dynamic rigid obstacles which is computed using adaptive Gauss-Kronrod quadrature in a preprocessing step. To reduce the memory footprint, the density maps are only stored for a narrow band around the solid surfaces and approximated with higher-order polynomials. During runtime they can be efficiently queried and accurate forces can be computed without the aforementioned problems of particle based approaches and without significant computational overhead. Moreover, friction between particles and solids, even in complex scenarios, can be handled.

Our method is inspired by the density maps approach and is also based on an implicit boundary representation. We use the same data structures as described in [Koschier and Bender 2017] to store and query function values. However, instead of precomputing the density contribution of the boundary, we determine the boundary volume that overlaps with the support domain of the smoothing kernel. This has multiple advantages which are demonstrated in a series of experiments. When using the density maps approach, the derivatives of density values are determined by differentiating the shape functions of the cubic elements. However, this is prone to errors because continuity at cell interfaces is not guaranteed [Koschier et al. 2017]. Our experiments show that this is not problematic for high resolution grids but it results in visual artifacts for lower resolutions due to force discontinuities. Compared to that, we compute all derivatives by differentiating the SPH kernel which is the standard SPH formulation. This results in smooth forces even for low resolution volume maps which enables significantly faster precomputations and reduces the memory requirements drastically.

Another advantage is that it directly fits into the SPH concept such that volume maps can be seamlessly integrated into existing solvers and combined easily with nearly all existing SPH methods. Finally, in contrast to the density maps approach, volume maps are independent of the smoothing kernel. Hence, a single instance of a volume map can be used in conjunction with various kernel functions. This is especially advantageous as some methods employ individual smoothing kernels for specialized purposes, e.g. surface tension or viscosity approaches, where multiple instances of the kernel-dependent density maps would be required.

2 RELATED WORK

During the last decades there has been a large variety of research on SPH methods. They have become a popular tool for simulating incompressible fluids, highly viscous liquids, fluid structure interactions and deformable solids. In this section, we briefly discuss prior works that are related to the presented method. A more general overview of recent developments in SPH can be found in the state-of-the-art report by Ihmsen et al. [2014b] and in the course notes of Koschier et al. [2019].

SPH was introduced to the computer graphics community by Stam and Fiume [1995] to simulate fire and gaseous phenomena. Later, it was used to simulate deformable solids by Desbrun and Gascuel [1996] and to simulate fluids by Müller et al. [2003]. To achieve nearly incompressible fluids, Becker and Teschner [2007] proposed an equation of state based approach. By choosing an appropriate stiffness constant for the pressure forces they can guarantee a small maximal compression. To alleviate the small time step requirements and stability issues of explicitly integrated pressure forces, many implicit solvers have been developed. One way of enforcing incompressibility is to apply a constant density constraint as proposed by Solenthaler and Pajarola [2009] who solve it with a predictive-corrective scheme. Similar approaches based on position based and projective dynamics have been proposed by Macklin and Müller [2013] and Weiler et al. [2016], respectively. Ihmsen et al. [2014a] showed that constant density can also be enforced by solving a discrete pressure Poisson equation. Another way to guarantee incompressibility is to solve a Poisson equation in order to make the velocity field divergence-free [Cornelis et al. 2019]. This has been done in a hybrid SPH and grid based approach by Raveendran et al. [2011]. Bender and Koschier [2017] proposed to solve for both, a divergence-free velocity field and a constant density, which increases stability and computational performance. In our experiments we use the DFSPH pressure solver of Bender and Koschier although the proposed boundary handling method can be combined with any of the aforementioned solvers.

Due to the particle nature of SPH it is not straightforward to enforce boundary conditions at interfaces between fluids and solid surfaces. One problem is that particles which are located close to the boundary typically suffer from neighborhood particle deficiencies leading to inaccurate density estimates and pressure forces. Another problem is that implicit pressure solvers also require pressure values inside the boundary domain which are unknown. These problems are orthogonal to each other and several methods to solve them have been proposed in recent years. Most approaches use specialized boundary representations, including particles, triangle meshes, and implicit representations such as signed distance fields. In the following we will give an overview of these methods.

A popular approach is to sample the solid boundaries with particles which exert penalty forces onto nearby fluid particles such that penetrations are avoided [Becker and Teschner 2007; Monaghan 1994]. Instead of integrating these forces with explicit schemes, Becker et al. [2009] proposed a predictor-corrector scheme. A problem is that their method suffers from particle stacking artifacts near the boundary. Alternatively, the boundary particles can be treated as fluid particles such that they contribute to the density and pressure computations as suggested by Solenthaler et al. [2007]. This solves the particle deficiency problem and results in smooth density distributions near solid boundaries. However, small time steps are necessary to guarantee non-penetration. Ihmsen et al. [2010] combine the approaches of Becker et al. [2009] and Solenthaler et al. [2007] to get smooth particle distributions without requiring small time steps. Their method was further improved by Akinci et al. [2012] such that two-way coupling of fluids with rigid bodies can be achieved. Moreover, they introduced normalized pseudo masses to account for non-uniform boundary particle distributions. He et al. [2012] proposed to add auxiliary staggered particles to discretize a Poisson equation with appropriate boundary conditions for projecting the velocity field onto a divergence-free state. For two-way coupling of fluids and deformable solids several adaptive particle sampling approaches have been proposed [Akinci et al. 2013; Müller et al. 2004; Yang et al. 2012].

One problem of the particle based boundary representations is that the surfaces are typically irregular. This causes non-smooth surface normals and therefore non-penetration forces that are not completely orthogonal to the surface which leads to undesired artificial friction or even jumping particle artifacts [Koschier and Bender 2017]. Band et al. [2017] proposed a method to alleviate this problem. They use the moving least squares (MLS) method to fit local planes to the boundary particles. This results in correct normals and eliminates most artifacts. However, it can be only applied in planar surface regions while curved surfaces cannot be handled. Another disadvantage of boundary particles is that they introduce a substantial computational overhead. They need to be considered during neighborhood searches and for each fluid particle it is required to iterate over all neighboring boundary particles during force and pressure computations.

Alternatively, solid boundaries can be represented as triangle meshes. Bodin et al. [2012] use unilateral constraints to prevent the fluid particles from penetrating triangulated surfaces of static and dynamic rigid bodies. The two-way coupling of fluids with cloth was addressed by Huber et al. [2015]. They use continuous collision detection and apply correction impulses such that the fluid particles cannot penetrate the cloth. However, neither of the approaches addresses the problem of inaccurate density estimates close to the boundary. Fujisawa and Miura [2015] also use meshes as a boundary representation and adopt a solution for the density problem from the engineering community [Kulasegaram et al. 2004]. In order to correct the SPH approximations, they compute a renormalization factor which takes the overlapping volume of the SPH kernel and the boundary into account. However, they compute these factors at runtime which results in significantly higher computation times compared to the particle based approaches.

Solid boundaries can be alternatively represented with implicit functions, e.g. signed distance fields (SDFs) as proposed by Harada et al. [2007a; 2007b]. To virtually extend the fluid density into the solid they sample the density contribution of a planar surface depending on the distance of a prototype fluid particle. The resulting values are discretely sampled prior to the simulation and are then queried during runtime based on the signed distance value of the fluid particle. However, this is only correct for planar surfaces and results in incorrect density values for curved geometries. Koschier and Bender [2017] proposed to precompute the density contributions of rigid boundaries and store them in a so called density map which can be efficiently queried during runtime. This results in accurate density values and avoids the artifacts of bumpy particle samplings. Since our work is inspired by the density maps approach, it will be discussed in more detail in Section 4.2.

A very popular approach to solve the problem of required but unknown pressure values at the boundary is a technique called pressure mirroring (cf. [Akinci et al. 2012]). While processing a fluid particle that requires to determine the pressure value at a boundary sample, pressure mirroring assumes that this value is equal to the fluid's pressure value. This is easy to implement and computationally cheap, but has the problem that the pressure values at the boundary are not unique. This happens when several fluid particles interact with the same boundary sample and each particle mirrors its own individual pressure value. As shown by Band et al. [2018a] this reduces the stability and convergence of the pressure solver. Therefore, they suggested to include the boundary particles into the pressure projection step. In a follow-up work, Band et al. [2018b] point out that it is even more beneficial to compute unique boundary pressures by extrapolating the values from the fluid particles using MLS. In our work we focus on removing artifacts induced by the irregularity of the typically particle-sampled surface and to enhance accuracy in the evaluation of discrete field samples and their derivatives in close proximity to the boundary. However, we would like to stress the fact that the discussed boundary pressure computation approaches are orthogonal to the proposed method and we are planning to adopt them in future work.

3 FOUNDATIONS

In this section we briefly introduce the governing equations to simulate incompressible fluids. Moreover, we derive the standard SPH discretization which is then extended in the next section by boundary handling.

Governing Equations. In our work we simulate fluids by using the Navier-Stokes equations for incompressible flow in Lagrangian coordinates

$$\frac{D\rho}{Dt} = 0 \quad \Leftrightarrow \quad \nabla \cdot \mathbf{v} = 0 \tag{1}$$

$$\frac{D\mathbf{v}}{Dt} = -\frac{1}{\rho}\nabla p + \nu\nabla^2 \mathbf{v} + \frac{\mathbf{f}}{\rho},\tag{2}$$

where ρ , *t*, **v**, *p*, *v* and **f** denote density, time, velocity, pressure, kinematic viscosity and body forces, respectively. To solve the Navier-Stokes equations numerically a discretization is required.

Spatial Discretization. In the following we introduce the SPH formalism which is a common approach for spatial discretization.

MIG '19, October 28-30, 2019, Newcastle upon Tyne, United Kingdom



Figure 2: Definition of the fluid domain \mathcal{F} (blue), the boundary domain \mathcal{B} (gray) and the support domain $\mathcal{N}(\mathbf{x})$ of the compact kernel function (green) with radius *r* at position **x**. The fluid domain is discretized by particles.

To derive the standard SPH discretization for a function $g(\mathbf{x})$, we first rewrite it using the Dirac- δ identity (cf. [Koschier et al. 2019])

$$g(\mathbf{x}) = \int_{\Omega} g(\mathbf{x}') \, \delta(\mathbf{x} - \mathbf{x}') d\mathbf{x}',$$

where Ω defines the problem domain. Then, we approximate the integral by replacing δ with a kernel function W with compact support and finally approximate the resulting integral using a sum over a set of sampling points \mathbf{x}_i

$$g(\mathbf{x}) \approx \int_{\mathcal{N}(\mathbf{x})} g(\mathbf{x}') W(\|\mathbf{x} - \mathbf{x}'\|, h) d\mathbf{x}'$$
(3)

$$\approx \sum_{j} V_{j} g(\mathbf{x}_{j}) W(\|\mathbf{x} - \mathbf{x}_{j}\|, h), \tag{4}$$

where $N(\mathbf{x})$ defines the support domain of the kernel function (see Figure 2), *h* is the smoothing length of the kernel and V_j is the volume represented by a sampling point.

In an SPH simulation the fluid is discretized by particles. Each particle *i* represents a sampling point and its volume is computed as $\frac{m_i}{\rho_i}$. The mass m_i is determined by the size of the particle and the rest density ρ_0 of the fluid. The density ρ_i of particle *i* is computed using the SPH formalism

$$\rho_i = \sum_j m_j W_{ij},\tag{5}$$

where $W_{ij} = W(||\mathbf{x}_i - \mathbf{x}_j||, h)$.

4 BOUNDARY HANDLING

Most graphics-related techniques to handle boundary conditions in the context of SPH solvers are based on the concept of virtually extending field quantities into the boundary domain, e.g. density, pressure, velocity etc. [Koschier et al. 2019]. Moreover, the vast majority of proposed methods consider incompressible continua and, therefore, ensure that the incompressibility constraint (see Equation (1)) is satisfied. By extending the density field into the boundary domain, the typically employed pressure solvers implicitly resolve boundary penetrations as the virtual density contributions lead to local compressions that violate the incompressibility constraint. Recent work has demonstrated the generality and versatility of the concept (cf. [Gissler et al. 2019]). Building on this concept, we will assume that (at least) the mass density ρ is extended into the boundary such that $\rho > 0$ on \mathcal{B} .

Given the (extended) density field, a smoothed density sample $\rho(\mathbf{x})$ located in close proximity to the boundary domain is determined as follows. The density function is plugged into Equation (3). We can then partition the integral into fluid and boundary portion (see Figure 2)

$$\rho(\mathbf{x}) \approx \int_{\mathcal{N}(\mathbf{x})} \rho(\mathbf{x}') W(\|\mathbf{x} - \mathbf{x}'\|, h) d\mathbf{x}'$$

$$= \int_{\mathcal{N}(\mathbf{x})\cap\mathcal{F}} \rho(\mathbf{x}') W(\|\mathbf{x} - \mathbf{x}'\|, h) d\mathbf{x}' +$$

$$\int_{\mathcal{N}(\mathbf{x})\cap\mathcal{B}} \rho(\mathbf{x}') W(\|\mathbf{x} - \mathbf{x}'\|, h) d\mathbf{x}'$$

$$= \rho_{\mathcal{F}}(\mathbf{x}) + \rho_{\mathcal{B}}(\mathbf{x}).$$
(6)

While the fluid integral $\rho_{\mathcal{F}}$ is typically discretized into SPH particles and numerically approximated using Equation (5), there exist various approaches to numerically approximate the boundary portion $\rho_{\mathcal{B}}$. The most widely used class of approaches is based on particle sampling. Recently, an alternative concept, referred to as density maps, has been proposed, where samples of $\rho_{\mathcal{B}}$ are computed using adaptive numerical quadrature to discretize the function onto a spatial grid (cf. [Koschier and Bender 2017]). In the following, we first briefly describe the core concepts of particle-based approaches and density maps and then introduce our novel volume maps approach. We further discuss the differences between all approaches and review the advantages of our novel method in the next section.

4.1 Particle-Based Approaches

The core idea of particle-based approaches is to discretize the boundary geometry with particles [Akinci et al. 2012; Ihmsen et al. 2010; Monaghan 1994]. These boundary particles typically have the same size as the fluid particles and serve as additional sampling points. In this way the boundary portion $\rho_{\mathcal{B}}(\mathbf{x})$ can be discretized analogously to the fluid portion using a sum (see Equation (5))

$$\rho_i = \rho_{\mathcal{F}}(\mathbf{x}_i) + \rho_{\mathcal{B}}(\mathbf{x}_i) \approx \sum_j m_j W_{ij} + \sum_k \tilde{m}_k \tilde{W}_{ik},$$

where *j* and *k* denote the indices of particle *i*'s fluid and boundary particle neighbors, respectively. Some methods employ a specialized kernel function \tilde{W} for boundary handling, e.g. [Becker and Teschner 2007], although the boundary kernels are usually chosen as $\tilde{W} \equiv W$. Please note that the boundary particle masses \tilde{m}_k are independent of the boundary object's material properties and solely fulfill the purpose of extending the fluid's mass density field into the boundary. The values \tilde{m}_k are therefore often called *pseudo masses*.

One of the most popular approaches for particle-based boundary handling was proposed by Akinci et al. [2012]. In this approach the pseudo masses are chosen such that $\rho_{\mathcal{B}} \approx \rho_0$ inside the boundary domain. Therefore, Akinci et al.'s method uses a constant function to extend the mass density field into the boundary. They compute

MIG '19, October 28-30, 2019, Newcastle upon Tyne, United Kingdom

the pseudo mass of a boundary particle k as

$$\tilde{m}_k = \frac{\rho_0}{\sum_l W_{kl}},$$

where l denotes the indices of the neighboring boundary particles of k, to account for non-uniform boundary particle distributions.

For a more detailed discussion of particle-based approaches we would like to refer the reader to the work of Koschier et al. [2019].

4.2 Density Maps

Instead of sampling the boundary with particles one could also solve the integral over the boundary part in Equation (6) by numerical integration methods such as Gauss quadrature while setting the density in the boundary to ρ_0 . However, the problem with this approach is twofold. First, the numerical integration is expensive and significantly decreases the performance of the simulation. Second, we have to integrate a discontinuous, piecewise constant function as the density field is zero outside of the boundary and takes on ρ_0 inside the boundary. This leads to staircase artifacts in the resulting function due to the limited accuracy of fixed pattern quadrature schemes (see Figure 3).

The density maps concept of Koschier and Bender [2017], which is introduced in this section, solves both problems. The boundary density function $\rho_{\mathcal{B}}$ is discretized and precomputed on a spatial grid to solve the performance problem. The discontinuities are removed by extending the boundary density to the outside of the boundary to get a smooth function. Since our novel boundary handling approach is inspired by the density maps method, we will discuss this concept here in more detail.

In contrast to the particle-based approaches, the density maps concept is based on an implicit representation of the boundary. Building on a signed distance function $\Phi : \mathbb{R}^3 \to \mathbb{R}$ whose zero iso-surface describes the boundary geometry, the authors extend the density field into the boundary using the extension function $\gamma : \mathbb{R} \to \mathbb{R}^+$. They further convolve the density extension function with the kernel in order to ensure compliance with the SPH formalism (cf. Equation (6)), i.e.

$$\rho_{\mathcal{B}}(\mathbf{x}) = \int_{\mathcal{N}(\mathbf{x})} \gamma(\Phi(\mathbf{x}')) W(\|\mathbf{x} - \mathbf{x}'\|, h) d\mathbf{x}', \tag{7}$$

where the signed distance function is defined as

$$\Phi(\mathbf{x}) = \begin{cases} -d(\mathbf{x}, \partial \mathcal{B}) & \text{if } \mathbf{x} \in \mathcal{B} \\ d(\mathbf{x}, \partial \mathcal{B}) & \text{otherwise,} \end{cases}$$
(8)

and where

$$d(\mathbf{x}, \partial \mathcal{B}) = \inf_{\tilde{\mathbf{x}} \in \partial \mathcal{B}} \|\mathbf{x} - \tilde{\mathbf{x}}\|$$

defines the shortest (unsigned) distance from **x** to the surface of the boundary $\partial \mathcal{B}$. They further model the extension function as a linear polynomial that takes the rest density value on the boundary surface

$$\gamma(x) = \begin{cases} \rho_0 \left(1 - \frac{x}{r}\right) & \text{if } x < r\\ 0 & \text{otherwise.} \end{cases}$$
(9)

This means instead of having zero density outside of the boundary and ρ_0 inside, we have a linear function increasing from 0 to ρ_0 within the support radius *r*. In this way the extension function has no discontinuity within the support radius of the fluid



Figure 3: Left: As the black fluid particle with support radius r enters the boundary, the boundary part of the density $\rho_{\mathcal{B}}$ increases. Right: Computing $\rho_{\mathcal{B}}$ by numerical integration leads to a discontinuity each time a sampling point (blue) enters the boundary.

particle. In order to evaluate the boundary density function ρ_B at a certain sample position **x**, Koschier and Bender numerically approximate the integral using adaptive Gauss-Kronrod quadrature. As this procedure is computationally expensive for arbitrarily complex boundary shapes and therefore not bearable at runtime, they discretize ρ_B on a regular grid with cubic shape functions of Serendipity type. The discretized function can then be efficiently accessed at runtime such that boundary density values and gradients can be computed in O(1). They justify the usage of the higher-order polynomial discretization with the argument that the convolution with typically piecewise polynomial kernels results in a sufficiently smooth field. Therefore, the cubic discretization usually requires fewer degrees of freedom to capture the field with high accuracy than linear discretizations.

4.3 Volume Maps

Our novel volume maps approach is inspired by the density maps concept and is also based on an implicit boundary representation. Moreover, we also define an extension function to avoid discontinuities and use a spatial grid to improve performance. However, instead of computing the density part of the boundary directly, we determine the intersection volume of the sphere around a fluid particle that is defined by the support radius r and the boundary (see Figure 2). This offers several advantages as shown in the following.

The most important advantage is that the boundary volume can be directly used in the general SPH formulation (see Equation (4)) which enables a consistent computation of the density gradient. Volume maps only determine the intersection volume between kernel support domain and boundary and use the SPH formulation to compute the density gradient. In contrast, density maps determine the gradient by differentiating the cubic shape functions of Serendipity type of their spatial grid discretization. The latter way to compute the gradient has a significant drawback. It is not guaranteed that a density map is continuous over the interfaces of neighboring grid cells [Koschier et al. 2017]. This leads to 'kinks' in the discretized representation of the density function when using a coarse resolution. At the position of such a kink the described computation



Figure 4: (a) The linear extension function γ (Eq. (9)) used for density maps is not smooth at x = r. (b) In contrast our cubic function γ^* (Eq. (4.3)) is continuously differentiable. Note that the plot in (a) shows the normalized function γ/ρ_0 .

of the density gradient leads to pressure forces that are not consistently oriented and therefore to noticeable visual artifacts (see Section 5). This can be improved by using a high-resolution density map which is computationally expensive to generate and has significantly higher memory requirements. In our method we use the map only to determine the boundary volume. The gradient is computed with the classical SPH formulation. Therefore, our method yields smooth pressure forces that do not cause visual artifacts even for low resolution maps which is shown in Section 5.

In order to generate a volume map, we determine the boundary volume as

$$V_{\mathcal{B}}(\mathbf{x}) = \int_{\mathcal{N}(\mathbf{x})} \gamma^*(\Phi(\mathbf{x}')) d\mathbf{x}', \qquad (10)$$

where Φ is the signed distance function (see Equation (8)). Note that this formulation is similar to the integral for the boundary density in Equation (7). However, the computation of the volume does not require a convolution with the kernel function. Therefore, in contrast to density maps, volume maps are independent of the kernel and a single map can be used in conjunction with various kernel functions. Moreover, in our formulation we use a cubic extension function instead of a linear one (cf. Equation (9))

$$\gamma^*(x) = \begin{cases} \frac{C(x)}{C(0)} & \text{if } 0 < x < r\\ 1 & \text{if } x \le 0\\ 0 & \text{otherwise,} \end{cases}$$

where *C* denotes the cubic spline kernel [Monaghan 2005]. Note that this kernel is only used to define the cubic function γ^* while an arbitrary kernel can be used in the SPH formulation. Using a cubic extension function has the advantage that in contrast to the linear function of the density maps approach, we get a smooth transition at distance x = r (see Figure 4). Moreover, the fact that our cubic extension function is continuously differentiable significantly improves the Gauss quadrature approximation of the integral in Equation (10) in comparison to using the non-smooth linear function [Müller et al. 2012]. Finally, the integral in Equation (10) is solved using Gauss-Legendre quadrature and the volume function is spatially discretized using a regular grid with cubic shape functions of Serendipity type to get a volume map.

To compute the density of a particle we use the SPH formulation in Equation (4). Here a quantity at a position **x** is determined by a sum over a set of sampling points in the neighborhood of **x**. Each sampling point represents a volume and its quantity is weighted by a kernel function depending on the distance between the sampling point and **x**. To determine the boundary portion of the density $\rho_{\mathcal{B}}$ using our volume maps approach, we add a sampling point \mathbf{x}^* on the boundary which represents the boundary volume $V_{\mathcal{B}}$. Moreover, we extend the density field to the boundary and set the density at this new sampling point to ρ_0 . This yields the following equation for the boundary part of the density

$$\rho_{\mathcal{B}}(\mathbf{x}) = V_{\mathcal{B}}(\mathbf{x})\rho_0 W(\|\mathbf{x} - \mathbf{x}^*\|, h).$$

In our work we define \mathbf{x}^* as the closest point to \mathbf{x} on the boundary. This point can be easily determined by the signed distance function Φ that we also use to generate the volume map.

The volume map is generated in a precomputation step before the simulation. To compute the density contribution of the boundary during runtime, in each simulation step we have to determine the closest point \mathbf{x}^* on the surface of the boundary and the boundary volume $V_{\mathcal{B}}$ for each fluid particle \mathbf{x} that is close to boundary. Note that both quantities can be obtained by simple lookups in the map which require constant time. The resulting density contribution can be easily used in any pressure solver. The required density gradient can be determined analogously by taken the gradient of the kernel. Finally, the symmetric pressure force [Monaghan 2005] for a particle *i* is computed as

$$\begin{split} \mathbf{F}_{i}^{p} &= -m_{i}\sum_{j}m_{j}\left(\frac{p_{i}}{\rho_{i}^{2}} + \frac{p_{j}}{\rho_{j}^{2}}\right)\nabla W(\|\mathbf{x} - \mathbf{x}_{j}\|, h) \\ &-m_{i}V_{j}\rho_{0}\left(\frac{p_{i}}{\rho_{i}^{2}} + \frac{\tilde{p}}{\tilde{\rho}^{2}}\right)\nabla W(\|\mathbf{x} - \mathbf{x}^{*}\|, h) \\ &= \mathbf{F}_{i\leftarrow\mathcal{F}}^{p} + \mathbf{F}_{i\leftarrow\mathcal{B}}^{p}, \end{split}$$

where $\mathbf{F}_{i\leftarrow\mathcal{F}}^{p}$ is the part of the pressure force that is exerted by the neighboring fluid particles and $\mathbf{F}_{i\leftarrow\mathcal{B}}^{p}$ is the part that is exerted by the boundary. \tilde{p} and $\tilde{\rho}$ are the pressure and the density at the boundary, respectively. In our work we used pressure mirroring [Akinci et al. 2012] and therefore set $\tilde{p} = p_i$ and $\tilde{\rho} = \rho_0$. However, in future we want to compute the boundary pressure value by extrapolation as proposed by Band et al. [2018b]. Their pressure extrapolation can be combined easily with our approach by setting \tilde{p} to the extrapolated value.

In case of dynamic boundaries we apply the negative force $F_{\mathcal{B}\leftarrow i}^{p} = -F_{i\leftarrow\mathcal{B}}^{p}$ at the position \mathbf{x}^{*} to the boundary object in order to realize two-way coupling. This two-way coupling is shown in the simulations in Figures 1 and 10.

5 RESULTS AND DISCUSSION

In this section, we present results and compare our volume maps concept to the particle-based method of Akinci et al. [2012] and the density maps approach [Koschier and Bender 2017]. For the comparison we implemented all methods in the open-source SPH library SPlisHSPlasH [Bender 2019]. All simulations were performed using the implicit pressure solver DFSPH with an adaptive time-stepping scheme based on the CFL condition [Bender and Koschier 2017]. In fluid simulations we further employed a micropolar model to

MIG '19, October 28-30, 2019, Newcastle upon Tyne, United Kingdom



Figure 5: A regular grid of fluid particles is dropped on an inclined plane. (a) The particle-based sampling causes a chaotic motion. (b) Our approach shows the expected smooth motion.



Figure 6: Sequence of five steps in two simulations of a sliding cube on a frictionless inclined plane comparing the method of Akinci et al. (left) and our method (right). While our method shows the expected behavior (right), the cube simulated using the particle-based approach of Akinci et al. (left) slightly drifts to the left, starts to rotate and gets slower due to artificial friction.

simulate vorticity [Bender et al. 2018]. Dynamic rigid bodies are simulated using a position-based method [Bender et al. 2017; Deul et al. 2014]. The performance was measured on two Intel Xeon E5-2697 processors with 2.3GHz.

Comparisons. To compare our method to the particle-based approach of Akinci et al., we performed several experiments. In the first simulation we dropped a regular grid of fluid particles on an inclined plane (see Figure 5). The space between the particles is larger than the support radius so that the particles do not influence each other. The particle-based sampling of the boundary causes a chaotic movement of the particles (see Figure 5a) while the simulation with our volume maps approach shows the expected smooth motion (see Figure 5b).

In the second experiment we simulated the motion of a deformable cube on a frictionless inclined plane (see Figure 6). The elastic behavior is simulated using the implicit SPH formulation for deformable solids of Peer et al. [2017]. The results show that the non-smooth surface of the particle-based sampling leads to an undesired drift, artificial friction and a rotational motion. Since the



Figure 7: Simulation of a deformable cube that rotates on a plane without friction. Left: In the simulation with the method of Akinci et al. the cube looses about 90% of its angular velocity and it drifts to the right. Right: Using our method no undesired drift is noticeable and the angular velocity is almost maintained.

volume maps are better suited to represent smooth surfaces, we did not experience any artifacts when using our novel method.

Finally, we simulated another deformable cube that rotates without friction for 100 seconds on a plane (see Figure 7). Again the particle-based sampling causes an undesired drift and artificial damping. The cubes started with an angular velocity of 5.1 rad/s. At the end of the simulation the cube which was simulated with the boundary handling of Akinci et al. lost almost 90 percent of its angular velocity and the cube's center shifted substantially. Using our novel method the loss of angular velocity was less than one percent and no drift occurs.

To compare our method to the density maps approach of Koschier et al., we computed the pressure force for a single fluid particle which is moved along the normal direction to a planar boundary. The plots in Figure 8 show the normal component of the resulting force (the other components are zero) for density maps and volume maps with different resolutions. For the density map construction we solved the integral in Equation (7) for each grid point which is close to the boundary using a high-resolution Gauss-Legendre quadrature with 17.5k sampling points. Since the extension function of density maps is not continuously differentiable, this high resolution of sampling points is required to get a sufficient approximation when using Gauss quadrature [Müller et al. 2012]. Figure 8 shows that low resolution density maps lead to discontinuities in the pressure force which in turn cause visual artifacts in a simulation as demonstrated in Figure 9. The volume map construction was also performed by solving an integral (see Equation (10)) using Gauss-Legendre quadrature. However, due to our smooth cubic extension function 4k sampling points were already sufficient. As shown in Figure 8 the resulting force was continuous even for low resolution maps and no visual artifacts occurred during the dam break simulation in Figure 9.

We believe that the kinks in the pressure force derived using density maps (see Figure 8), are caused by small discontinuities between neighboring cells of the spatial grid (cf. [Koschier et al. 2017]). Since the density maps approach uses the shape functions of the spatial discretization to compute the density gradient, these



Figure 8: Pressure force. A single fluid particle is moved in normal direction towards a boundary plane. The plots show the pressure force in normal direction (y-axis) for different map resolutions in relation to the distance to the boundary (x-axis). When using density maps, the pressure force increases slowly and is not smooth for low resolution maps (a,b,c,d). The force increases faster when using volume maps and is smooth even for low resolutions (e,f).

discontinuities have a large influence on the resulting pressure force and cause noticeable artifacts (see Figure 9). In contrast our volume maps method only determines the volume using a map while the gradient is computed using the standard SPH formulation. Therefore, the influence on the pressure force is only small and neither noticeable in the graph (see Figure 8) nor in the simulation (see Figure 9).

Finally, we compared the construction times and the memory requirements of the volume and density maps. We compared the 3D maps of the unit cube that were used to generate the graphs in Figure 8. At the same resolution the requirements are comparable. However, the density map required a resolution of $60 \times 60 \times 60$ while a volume map only needed a resolution of $10 \times 10 \times 10$ to obtain a smooth result. We observed the same resolution requirements

Table 1: Average times per step required for the neighborhood search, the pressure solve and the computation of the non-pressure forces in a double dam break simulation (see Figure 11).

	Akinci et al.	Density Maps	Volume Maps
neigh. search	411 ms	379 ms	369 ms
pressure	1090 ms	1054 ms	1078 ms
non-pressure	357 ms	339 ms	342 ms
total	1858 ms	1715 ms	1789 ms

in our 2D dam break simulation. The corresponding density map required 78.2 MB while the volume map only needed 392 KB due to its lower resolution. Hence, the memory consumption of the volume map was 200 times lower. The density map construction took 175 s while the generation of the volume map only needed 0.3 s. This yields a speedup factor of about 580. Naturally the speedup and memory requirements depend on the geometry that is discretized. Considering a complex geometry with small features, the discrete map naturally requires a higher resolution. However, the density map still relies on a far higher resolution to avoid the stability issues. At this point we would like to mention that the current implementation uses a regular grid to store the map. In the future, we plan to improve our implementation by using an adaptive sparse grid which will further reduce the memory requirements.

Complex Boundaries. We performed two simulations to show that our novel method is able to handle scenarios with complex static and dynamic boundaries. First, we simulated a breaking dam in a canyon (see Figure 1, left). To capture the detailed canyon surface we used a map resolution of $128 \times 64 \times 768$ in this scenario. We want to emphasize that this high resolution map is not required to guarantee a stable simulation. A realistic fluid-air interaction is simulated using the method of Gissler et al. [2017]. The simulation shows that our method is able to handle 8 million turbulent fluid particles that interact with a large-scale complex boundary.

In the second simulation 12M fluid particles, which were generated by four emitters, interact with three static dragons and ten dynamic ducks (see Figure 1, right). This demonstrates that our method robustly handles coupling between a turbulent fluid and dynamic bodies.

Combination with Existing SPH Methods. Our novel boundary handling method can be easily combined with other existing SPH methods. To demonstrate this we created a scenario consisting of several emitters, static dragons and dynamic ducks (see Figure 10). This time we combined our method with the multiphase approach of Solenthaler and Pajarola [2008] and the implicit viscosity solver of Weiler et al. [2018] to simulate the interaction of water, highly viscous material and dynamic and static boundaries. We used 5 million particles for the water phase and 1.3 million particles for the highly viscous material. The simulation shows that our novel boundary handling method can be seamlessly combined in other existing SPH methods.

Performance. In our last experiment we simulated a double dam break consisting of 3 million fluid particles (see Figure 11). This



(a) Density Maps 10×10

(b) Density Maps 60×60

(c) Volume Maps 10×10

Figure 9: Comparison of the density and volume maps using a 2D dam break simulation. (a) Low resolution density maps cause visual artifacts. (b) We can alleviate this problem by using a high-resolution density map. (c) Volume maps avoid these artifacts by a consistent gradient computation and do not suffer from visual artifacts even when using a low-resolution map.



Figure 10: Multiphase simulation where 5 million fluid particles interact with 1.3 million particles of a highly viscous material and multiple dynamic rigid bodies. The simulation demonstrates that our volume maps boundary handling can be easily combined with an implicit viscosity solver and an SPH multiphase method.

scenario was simulated once using the boundary handling of Akinci et al., once using density maps and once using our novel volume maps formulation in order to compare the performance. The average times per simulation step for the neighborhood search and the pressure and non-pressure forces are given in Table 1. The results show that the neighborhood search is faster when using density or volume maps. The reason for this is that no boundary particles need to be considered in the neighborhood search. The computation of the pressure and non-pressure forces using a map is also slightly faster than for the particle-based method since instead of a sum over all boundary neighbors only a lookup in a map is required. The comparison shows that our novel method has a performance that is comparable to the other methods and is even slightly faster than the particle-based approach.

6 CONCLUSION

In this paper we introduced a novel implicit boundary representation called volume maps. The most popular boundary handling

methods in SPH simulations are based on a particle representation of the boundary. However, the boundary particles lead to bumpy surfaces and inaccurate pressure forces. Density maps provide an implicit representation of the boundary surfaces which avoids these problems. However, high-resolution maps are required since the density gradient computation suffers from discontinuities in the map. In contrast to the density maps approach our method only determines the overlapping volume of the boundary and the support domain of the smoothing kernel. The gradient computation uses the classical SPH formulation which provides a consistent pressure force computation even for low resolution maps and does not suffer from map discontinuities. Therefore, in comparison to density maps we need about two orders of magnitude less precomputation time and memory. Since our method directly fits into the classical SPH formulation, it can be easily combined with any existing SPH method. Finally, in contrast to density maps our representation is independent of a specific kernel function and a single map can be combined with methods that use different kernels.

MIG '19, October 28-30, 2019, Newcastle upon Tyne, United Kingdom



Figure 11: Double dam break simulation with 3 million water particles.

Our method has the same limitations as the density maps approach. A volume map cannot be precomputed for a deformable solid and it is computationally expensive to compute the map at runtime. Moreover, the cubic shape functions that we use for our map smooth out sharp features in the boundary model. However, this problem can be alleviated by using an adaptive discretization for the map which is a goal for our future work. Another future goal is to combine our volume maps with pressure extrapolation to further improve the boundary handling.

REFERENCES

- N. Akinci, J. Cornelis, G. Akinci, and M. Teschner. 2013. Coupling elastic solids with smoothed particle hydrodynamics fluids. *Computer Animation and Virtual Worlds* 24, 3-4 (2013), 195–203.
- N. Akinci, M. Ihmsen, G. Akinci, B. Solenthaler, and M. Teschner. 2012. Versatile rigid-fluid coupling for incompressible SPH. ACM Transactions on Graphics 31, 4 (2012), 62:1–62:8.
- Stefan Band, Christoph Gissler, Markus Ihmsen, Jens Cornelis, Andreas Peer, and Matthias Teschner. 2018a. Pressure Boundaries for Implicit Incompressible SPH. ACM Transactions on Graphics 37, 2, Article 14 (2018), 14:1–14:11 pages.
- Stefan Band, Christoph Gissler, Andreas Peer, and Matthias Teschner. 2018b. MLS pressure boundaries for divergence-free and viscous SPH fluids. Computers & Graphics 76 (2018), 37–46.
- Stefan Band, Christoph Gissler, and Matthias Teschner. 2017. Moving least squares boundaries for SPH fluids. In Proceedings of the 13th Workshop on Virtual Reality Interactions and Physical Simulations. Eurographics Association, 21–28.
- M. Becker and M. Teschner. 2007. Weakly compressible SPH for free surface flows. In ACM SIGGRAPH/Eurographics Symposium on Computer Animation. 1–8.
- M. Becker, H. Tessendorf, and M. Teschner. 2009. Direct Forcing for Lagrangian Rigid-Fluid Coupling. *IEEE Transactions on Visualization and Computer Graphics* 15, 3 (2009), 493–503.
- Jan Bender. 2019. SPlisHSPlasH Library.
- https://github.com/InteractiveComputerGraphics/SPlisHSPlasH.
- Jan Bender and Dan Koschier. 2017. Divergence-Free SPH for Incompressible and Viscous Fluids. IEEE Transactions on Visualization and Computer Graphics 23, 3 (2017), 1193–1206.
- Jan Bender, Dan Koschier, Tassilo Kugelstadt, and Marcel Weiler. 2018. Turbulent Micropolar SPH Fluids with Foam. *IEEE Transactions on Visualization and Computer Graphics* 25, 6 (2018), 2284–2295.
- Jan Bender, Matthias Müller, and Miles Macklin. 2017. Position-Based Simulation Methods in Computer Graphics. In EUROGRAPHICS 2017 Tutorials. Eurographics Association.
- K. Bodin, C. Lacoursière, and M. Servin. 2012. Constraint fluids. IEEE Transactions on Visualization and Computer Graphics 18 (2012), 516–526.
- Jens Cornelis, Jan Bender, Christoph Gissler, Markus Ihmsen, and Matthias Teschner. 2019. An optimized source term formulation for incompressible SPH. *The Visual Computer* 35, 4 (2019), 579–590.

- M. Desbrun and M.-P. Gascuel. 1996. Smoothed Particles: A new paradigm for animat-
- ing highly deformable bodies. In Eurographics Workshop on Computer Animation and Simulation. 61–76.
- Crispin Deul, Patrick Charrier, and Jan Bender. 2014. Position-based rigid-body dynamics. Computer Animation and Virtual Worlds 27, 2 (2014), 103–112.
- M. Fujisawa and K. Miura. 2015. An Efficient Boundary Handling with a Modified Density Calculation for SPH. Computer Graphics Forum 34, 7 (2015), 155–162.
- Christoph Gissler, Stefan Band, Andreas Peer, Markus Ihmsen, and Matthias Teschner. 2017. Generalized drag force for particle-based simulations. *Computers & Graphics* 69 (2017), 1–11.
- Christoph Gissler, Andreas Peer, Stefan Band, Jan Bender, and Matthias Teschner. 2019. Interlinked SPH Pressure Solvers for Strong Fluid-Rigid Coupling. ACM Transactions on Graphics 38, 1 (2019), 5:1–5:13.
- T. Harada, S. Koshizuka, and Y. Kawaguchi. 2007a. Smoothed particle hydrodynamics in complex shapes. In Spring Conf. on Computer Graph. 191–197.
- T. Harada, S. Koshizuka, and Y. Kawaguchi. 2007b. Smoothed Particle Hydrodynamics on GPUs. In Computer Graphics International. 63–70.
- X. He, N. Liu, S. Li, H. Wang, and G. Wang. 2012. Local Poisson SPH for Viscous Incompressible Fluids. Computer Graphics Forum 31 (2012), 1948–1958.
- M. Huber, B. Eberhardt, and D. Weiskopf. 2015. Boundary Handling at Cloth-Fluid Contact. Computer Graphics Forum 34, 1 (2015), 14–25.
- M. Ihmsen, N. Akinci, M. Gissler, and M. Teschner. 2010. Boundary handling and adaptive time-stepping for PCISPH. In Virtual Reality Interactions and Physical Simulations. 79–88.
- M. Ihmsen, J. Cornelis, B. Solenthaler, C. Horvath, and M. Teschner. 2014a. Implicit incompressible SPH. *IEEE Transactions on Visualization and Computer Graphics* 20, 3 (2014), 426–435.
- M. Ihmsen, J. Orthmann, B. Solenthaler, A. Kolb, and M. Teschner. 2014b. SPH Fluids in Computer Graphics. In *Eurographics (State of the Art Reports)*. 21–42. Dan Koschier and Jan Bender. 2017. Density Maps for Improved SPH Boundary
- Dan Koschier and Jan Bender. 2017. Density Maps for Improved SPH Boundary Handling. In ACM SIGGRAPH/Eurographics Symposium on Computer Animation. ACM, 1–10.
- Dan Koschier, Jan Bender, Barbara Solenthaler, and Matthias Teschner. 2019. Smoothed Particle Hydrodynamics for Physically-Based Simulation of Fluids and Solids. In *EUROGRAPHICS 2019 Tutorials*. Eurographics Association.
- Dan Koschier, Crispin Deul, Magnus Brand, and Jan Bender. 2017. An hp-Adaptive Discretization Algorithm for Signed Distance Field Generation. IEEE Transactions on Visualization and Computer Graphics 23, 10 (2017), 2208–2221.
- S. Kulasegaram, J. Bonet, R. W. Lewis, and M. Profit. 2004. A variational formulation based contact algorithm for rigid boundaries in two-dimensional SPH applications. *Computational Mechanics* 33, 4 (2004), 316–325.
- M. Macklin and M. Müller. 2013. Position Based Fluids. ACM Transactions on Graphics 32, 4 (2013), 1–5.
- J.J. Monaghan. 1994. Simulating Free Surface Flows with SPH. J. Comput. Phys. 110, 2 (1994), 399–406.
- J J Monaghan. 2005. Smoothed Particle Hydrodynamics. Reports on Progress in Physics 68, 8 (2005), 1703–1759.
- Björn Müller, Florian Kummer, Martin Oberlack, and Yongqi Wang. 2012. Simple multidimensional integration of discontinuous functions with application to level set methods. *Internat. J. Numer. Methods Engrg.* 92, 7 (2012), 637–651.
 M. Müller, D. Charypar, and M. Gross. 2003. Particle-Based Fluid Simulation for
- M. Müller, D. Charypar, and M. Gross. 2003. Particle-Based Fluid Simulation for Interactive Applications. In ACM SIGGRAPH/Eurographics Symposium on Computer Animation. 154–159.
- M. Müller, S. Schirm, M. Teschner, B. Heidelberger, and M. Gross. 2004. Interaction of fluids with deformable solids. *Computer Animation and Virtual Worlds* 15, 34 (2004), 159–171.
- Andreas Peer, Christoph Gissler, Stefan Band, and Matthias Teschner. 2017. An Implicit SPH Formulation for Incompressible Linearly Elastic Solids. *Computer Graphics Forum* 37, 6 (2017), 135–148.
- K. Raveendran, C. Wojtan, and G. Turk. 2011. Hybrid smoothed particle hydrodynamics. In ACM SIGGRAPH/Eurographics Symposium on Computer Animation. 33–42.
- B. Solenthaler and R. Pajarola. 2008. Density Contrast SPH Interfaces. In ACM SIG-GRAPH/Eurographics Symposium on Computer Animation. 211–218.
- B. Solenthaler and R. Pajarola. 2009. Predictive-corrective incompressible SPH. ACM Transactions on Graphics 28, 3 (2009), 40:1-40:6.
- B. Solenthaler, J. Schläfli, and R. Pajarola. 2007. A unified particle model for fluid-solid interactions. *Computer Animation and Virtual Worlds* 18, 1 (2007), 69–82.
- Jos Stam and Eugene Fiume. 1995. Depicting Fire and Other Gaseous Phenomena Using Diffusion Processes. In Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '95).
- M. Weiler, D. Koschier, and J. Bender. 2016. Projective Fluids. In ACM Motion in Games. ACM, New York, NY, USA, 79–84.
- Marcel Weiler, Dan Koschier, Magnus Brand, and Jan Bender. 2018. A Physically Consistent Implicit Viscosity Solver for SPH Fluids. *Computer Graphics Forum* 37, 2 (2018).
- L. Yang, S. Li, A. Hao, and H. Qin. 2012. Realtime Two-Way Coupling of Meshless Fluids and Nonlinear FEM. *Computer Graphics Forum* 31, 7 (2012), 2037–2046.