



OXFORD CENTRE FOR COLLABORATIVE APPLIED MATHEMATICS

Report Number 11/56

**Real-Time Fluid Effects on Surfaces using the Closest Point
Method**

by

**S. Auer, C. B. Macdonald, M. Treib, J. Schneider, and R.
Westermann**



Oxford Centre for Collaborative Applied Mathematics
Mathematical Institute
24 - 29 St Giles'
Oxford
OX1 3LB
England

Real-Time Fluid Effects on Surfaces using the Closest Point Method

S. Auer^{†,1}, C. B. Macdonald^{‡,2}, M. Treib^{§,1}, J. Schneider^{¶,3}, R. Westermann^{||,1}

¹Technische Universität München

²Oxford Centre for Collaborative Applied Mathematics (OCCAM)

³King Abdullah University of Science and Technology

Abstract

The Closest Point Method (CPM) is a method for numerically solving partial differential equations (PDEs) on arbitrary surfaces, independent of the existence of a surface parametrization. The CPM uses a closest point representation of the surface, to solve the unmodified Cartesian version of a surface PDE in a 3D volume embedding, using simple and well-understood techniques. In this paper we present the numerical solution of the wave equation and the incompressible Navier-Stokes equations on surfaces via the CPM, and we demonstrate surface appearance and shape variations in real-time using this method. To fully exploit the potential of the CPM, we present a novel GPU realization of the entire CPM pipeline. We propose a surface-embedding adaptive 3D spatial grid for efficient representation of the surface, and present a high-performance approach using CUDA for converting surfaces given by triangulations into this representation. For real-time performance, CUDA is also used for the numerical procedures of the CPM. For rendering the surface (and the PDE solution) directly from the closest point representation without the need to reconstruct a triangulated surface, we present a GPU ray-casting method that works on the adaptive 3D grid.

Categories and Subject Descriptors (according to ACM CCS): I.6.8 [Simulation and Modeling]: Types of Simulation—Parallel I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing

1. Introduction

Techniques for solving partial differential equations (PDE) have numerous applications in physics-based simulation and modeling, and they are frequently employed in computer graphics for realistically simulating real-world phenomena such as fluids or deformable solids. Most commonly, partial differential equations over \mathbb{R}^2 or \mathbb{R}^3 are considered, and, since analytical solutions only rarely exist, numerical solutions using some form of spatial discretization are employed.

In computer graphics, numerical solutions of PDEs are also used to produce effects on manifolds such as surfaces

in \mathbb{R}^3 . Prominent examples include surface texture synthesis [Tur91], texture assignment [CLB*09], or flow simulation on surfaces [Sta03]. If an isometric surface parametrization exists, a PDE defined on a surface can be transformed into a PDE on the 2D parameter domain and solved using standard discretizations [Sta03, LWC05, LFW07]. PDEs on surfaces can also be solved directly using finite element discretizations on a surface triangulation [DE07]. These techniques, however, rely on changing the non-parametric form of the PDE to agree with the underlying discretization. An alternative are so-called embedding techniques [BCOS01, NNRW09, CLB*09, RM08, CRT04], which lift the PDE to a narrow band around the surface and solve a transformed PDE on this band.

Embedding techniques are particularly attractive because they do not rely on the existence of a low distortion parametrization. The computation of such a parametrization can be difficult to achieve, if not impossible, and it is not suitable in applications where the surface undergoes frequent

[†] e-mail:auer@in.tum.de

[‡] e-mail:macdonald@maths.ox.ac.uk

[§] e-mail:treib@in.tum.de

[¶] e-mail:jens.schneider@kaust.edu.sa

^{||} e-mail:westermann@tum.de

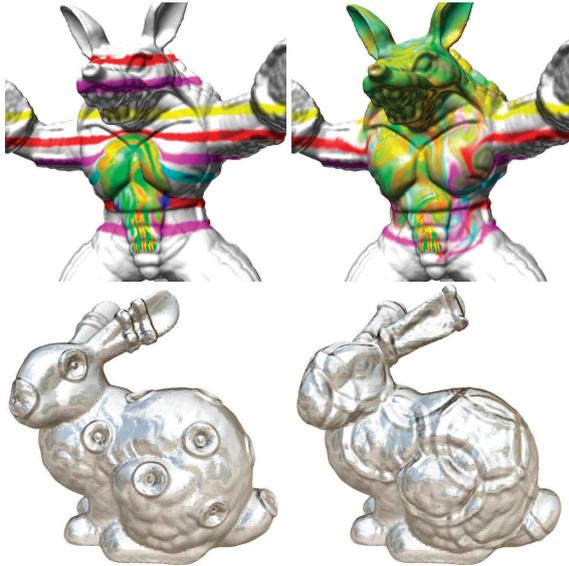


Figure 1: Numerical simulation of fluids on complicated surfaces via the Closest Point Method. Top: The simulation result is rendered via rasterization and pixel shading. Bottom: Raycasting the embedding computational grid allows simulating surface displacements. At a resolution that corresponds to a 512^3 Cartesian grid, simulation and rendering takes less than 100 ms per time step.

shape changes. Another advantage of embedding techniques is that once the PDE has been lifted to the embedding 3D space, standard numerical schemes can be used in this space to efficiently solve the PDE.

Among the embedding techniques, the Closest Point Method (CPM) is one of the less complicated methods in that it only requires the existence of a closest point representation of the surface (i.e., for any point in the embedding space, the surface point with the least Euclidean distance is known) and it solves completely standard PDEs—without any metric terms—in the embedding space using common numerical methods on a uniform Cartesian grid [RM08]. In addition, the CPM can handle general open surfaces and surfaces without orientation, and, thus, is applicable even in scenarios where the surface separates into parts.

Despite its simplicity and accuracy, however, a naïve implementation of the CPM is exhaustive in memory. Such an implementation would typically construct a full surface-embedding Cartesian grid and pre-compute at any grid point its closest surface point. It is clear that this severely limits the resolution of the computational grid to be used. Furthermore, the closest point representation remains fixed only under the assumption that the surface does not change. In scenarios where this is not the case, re-computing closest points can have a severe impact on performance.

Our contribution. The primary focus of this paper is the simulation of fluid effects on surfaces in real time using the CPM. To achieve this, we present a fast GPU method for realizing the CPM *and* rendering the simulation results. The method can efficiently solve non-linear PDEs on surfaces using a high-resolution embedding grid, and it can calculate and render the resulting fields independently of the input surface resolution. Due to the strict separation of the surface representation and the computational grid, our method supports initial surfaces of arbitrary topology and geometric detail. The method is implemented in CUDA and Direct3D10, and all parts have been optimized to make the best possible use of the massive parallel processing power and memory bandwidth on current GPUs.

The particular contributions of our paper are:

- A novel CUDA method for constructing an adaptive multiblock closest point grid representing a narrow spatial band around a given triangulated surface.
- A CPM for numerically solving the wave equation and the Navier-Stokes equations using Semi-Lagrangian advection.
- A GPU ray-caster that works directly on the adaptive closest point grid and can visualize surface appearance attributes and surface displacements.

Our paper is structured as follows: In Section 2 we discuss previous work that is related to ours. The CPM is reviewed in Section 3. Section 4 is dedicated to the efficient construction of an adaptive computational grid enclosing a triangulated surface representation, and the generation of the closest point representation. Section 5 sheds light on the numerical solution of the second-order linear wave equation and the incompressible Navier-Stokes equations via the CPM. Besides discretization aspects, the effects of the closest point extension on stability, accuracy and performance of the simulation are discussed. Section 6 introduces the volume raycaster that is used to visualize dynamic color and displacement effects independently of the surface triangulation. Section 7 presents a detailed performance analysis. The paper is concluded with a discussion of the advantages and limitations of the proposed method and some ideas for future work.

2. Related Work

An interesting surface-based problem is the simulation of fluid dynamics on surfaces. Stam applied a modified 2D version of his stable fluid dynamics simulation [Sta99] to the patches of Catmull–Clark surfaces and introduced a technique to transfer information in an overlap between the patches [Sta03]. While working well in regular areas of the surface, the approach leads to errors at non valence-four vertices. For techniques relying on some sort of global parametrization similar problems can be expected at the singularities.

Several other authors directly use the vertices and edges

of triangle meshes as a discretization and employ finite element, finite volume, Lattice Boltzmann or related numerical methods in their solvers [SY04, FZKH05, NMZ07, ATBG08]. This however takes away some advantages of a regular grid discretization as finite element methods are generally more complicated. Particularly on surfaces, the resulting algorithms may be quite difficult to implement in practice [RWP06]. Finally, if the input surface is not given as a triangulation a costly reconstruction procedure may be required.

The interest in simulation on surfaces also resulted in some particle based solutions. Turk [Tur91] simulated reaction diffusion systems in a grid formed by particle relaxation on a surface, with connectivity given by a Voronoi diagram. Later, the technique was adapted to general shallow water equations [WMT07].

Chuang and co-workers [CLB*09] proposed a volume embedding scheme based on B-splines as Ansatz-functions in a hexahedral simulation grid. By restricting the basis functions to a surface instead of a sub-volume, a weak form of a surface PDE can be solved independently of the 3D simulation domain, but at the expense of explicitly clipping the surface mesh against the simulation cells.

The Closest Point Method was introduced by Ruuth and Merriman [RM08]. Later works presented an implicit time stepping [MR09] and used the CPM for the evolution of level sets [MR08] or segmentation on surfaces [TMR09]. Hong et. al [HZQW10] applied the CPM to fire simulation on animated surfaces.

Due to the embedding used, the CPM relies on certain properties of distance transforms, such as smoothness of properties close to the surface. Jones [JBS06] provides an excellent survey of these properties.

An additional topic related to our method is the GPU voxelization of triangle meshes. While earlier work employed the fixed function, 2D rasterization and blending stages for this task [ED08], recently Schwarz and Seidel [SS10], and Pantaleoni [Pan11] demonstrated the superior efficiency of the CUDA parallel programming API for constructing a solid volume from a surface mesh. Building upon these works we will demonstrate the efficient construction of a multiblock closest point grid.

3. The Closest Point Method

As mentioned above, the CPM is an embedding method where computations are performed in an embedding space surrounding the surface (typically a 3D volume) in such a way that the results are consistent with the solution of a surface partial differential equation. The main principle needed for this to be true is that of “equivalence of gradients”: the intrinsic surface gradient $\nabla_s u$ of a surface function u agrees on the surface with the standard Cartesian gradient of a volume function v , provided v is the closest point extension of

u [RM08]. Intuitively this is because the closest point extension $v(\mathbf{x}) = u(cp(\mathbf{x}))$ is constant in the normal direction to the surface so the change in v must be tangent to the surface. Here, $cp(\mathbf{x})$ denotes the surface point closest to the point \mathbf{x} . A second principle applies in a similar fashion to surface divergence operators [RM08]. The two principles can be combined to handle many other differential operators including the Laplace–Beltrami operator.

To build a numerical method on these ideas, consider a prototype problem describing the evolution of some attribute u on a surface

$$u_t = f(\nabla_s u)$$

where f is a general nonlinear function and $\nabla_s u$ is the intrinsic surface gradient. Suppose at some fixed time t we have a solution v^t defined over the volume which is constant in the normal direction to the surface. We can move the solution forward in time by one step of size Δt using a forward Euler discretization. This *evolution phase* requires us to evaluate the right-hand side at time t and by the principles mentioned above, this is exactly the same (for points on the surface) as evaluating $f(\nabla v^t)$ in the volume:

$$\tilde{v}^{t+1} := v^t + \Delta t f(\nabla v^t).$$

The new solution \tilde{v}^{t+1} might not be constant in the normal direction so we then perform a second *extension phase*, a closest point extension which projects values off of the surface into the surrounding volume to obtain a solution v^{t+1} which is constant in the normal direction:

$$v^{t+1}(\mathbf{x}) := \tilde{v}^{t+1}(cp(\mathbf{x})), \quad \text{for each point } \mathbf{x}.$$

To obtain values at $cp(\mathbf{x})$, which does not have to coincide with a grid sample \mathbf{x}_G in the discrete case, higher-order interpolation is used. We then repeat these two phases over and over to advance the solution in time. Several properties make the Closest Point Method simple and effective:

- In the embedding volume, simple well-understood finite difference schemes can be applied to evaluate $f(\nabla v)$.
- Posing the PDE in the embedding volume is easy: in our example, the nonlinear function f is *unchanged* in the volume calculation and we simply replace intrinsic differential operators with their standard Cartesian counterparts. There are no metric terms to deal with.
- In the extension phase, we need the closest point for each grid point in our embedding volume. This is the only geometry representation of the surface that is needed and it is a very general representation: allowing arbitrary co-dimension and non-orientable surfaces for example.
- Without effects on the accuracy, the computation can be performed on a narrow band enveloping the surface. No artificial boundary conditions need to be applied at the edges thereby.
- The accuracy of the method is well-understood and depends on the accuracy of the time-stepping scheme, the finite difference scheme and the interpolation scheme.

4. CPM on the GPU

In principle, the discrete closest point surface representation required by the CPM is a volumetric Cartesian grid, where each grid cell stores the surface point closest to its center. For smooth objects, the closest point is unique in a small band around the surface. For non-smooth objects or farther away from the surface, a point can have multiple closest points in general. In this case we select one of them arbitrarily.

4.1. Adaptive Multiblock Grid

For high resolution simulations, the use of a uniform simulation grid is impractical. For the CPM, on the other hand, it is sufficient to work on two narrow computational bands around the surface: The *evolution band* and the *extension band*, which are used in the evolution and extension phase of the CPM, respectively. This suggests using a sparse multiblock grid, which stores only those blocks of the full grid which overlap these bands. Figure 2 illustrates such a multiblock grid in 2D. The widths we choose for the bands are the ones proposed by Macdonald and Ruuth [MR08] (Section 4.1.1).

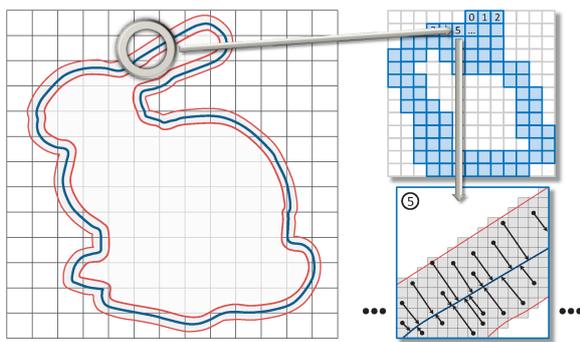


Figure 2: Sparse closest point grid. Left: A grid of blocks, some of them overlapping the narrow band around the surface. Top right: Only blocks overlapping the narrow band around the surface are stored in linear order. Bottom right: The closest points of the grid cells within these blocks.

4.2. Data Layout and Grid Generation

In the following discussions we assume a triangle mesh stored as a shared vertex list as the input surface representation. Since we are aiming to support both static and time-varying surfaces, for instance surfaces that are modified in turn by the CPM simulation, an important requirement is that the creation and update of the sparse closest point representation can be performed at high rates even for high-resolution surfaces and simulation grids. For this computationally and bandwidth intense task, we exploit the massively parallel architecture of the GPU.

Our proposed algorithm for constructing a closest point volume is inspired by recent CUDA voxelization approaches [SS10, Pan11]. Similar to the tile-based approaches proposed in these works, we use a sort-middle rasterization pipeline which first assigns triangles to coarse blocks and then performs a fine-grain voxelization and closest point calculation per block. There are, however, some specific differences of our GPU pipeline:

- It computes a voxelization of a computational band with arbitrary thickness around a surface, instead of a 6/26 separating triangle voxelization.
- The output is an adaptive multiblock grid instead of a full volume.
- It avoids scan-lines along the dominant triangle axis since a coarse raster to 3D blocks is used instead of 2D tiles.
- The result contains distances and closest points.
- It avoids blending via shared memory atomics and assigns threads in the fine raster to voxels instead of triangles.

For the sake of clarity, we will first describe the construction of the sparse multiblock grid, including the indexing schemes that are required to efficiently perform the CPM, before we go into the details of the actual closest point calculation.

To create the sparse simulation grid, the user first selects a simulation resolution by specifying the size Δx of one simulation cell. Then the bounding volume of the triangle mesh—enlarged to respect the width of the computational band—is subdivided into $k \times l \times m$ equally sized blocks. We assume that each block contains b^3 simulation cells, so that the values of k , l , and m can be computed.

The grid construction on the GPU consists of three stages: First, it is determined which and how many blocks are intersected by the computational band. The *block-index-buffer* is used to keep track of this information. Second, the *closest-point-buffer*—large enough to store the closest points of all simulation cells represented by these blocks—is allocated. To accommodate fast access to cells in adjacent blocks, an additional buffer—the *compacted-block-index-buffer*—is created. It has as many entries as there are intersected blocks, and it stores the positions of these blocks in *block-index-buffer*. From this, and by taking into account that in *block-index-buffer* all blocks are laid out in $x/y/z$ -order, adjacency information can be resolved. Third, the *expanded-tri-buffer* is constructed to support coalesced memory write operations in the closest point construction (see the following subsection). This buffer stores in contiguous sections the sets of triangles that are required by each block to compute the closest point grid. The different buffers used are illustrated in Figure 3. In the following, we will describe the parallel CUDA implementation of the three stages.

Stage 1: One CUDA-thread is launched per triangle, and each thread determines the blocks that are intersected by the computational band around this triangle. Therefore, each

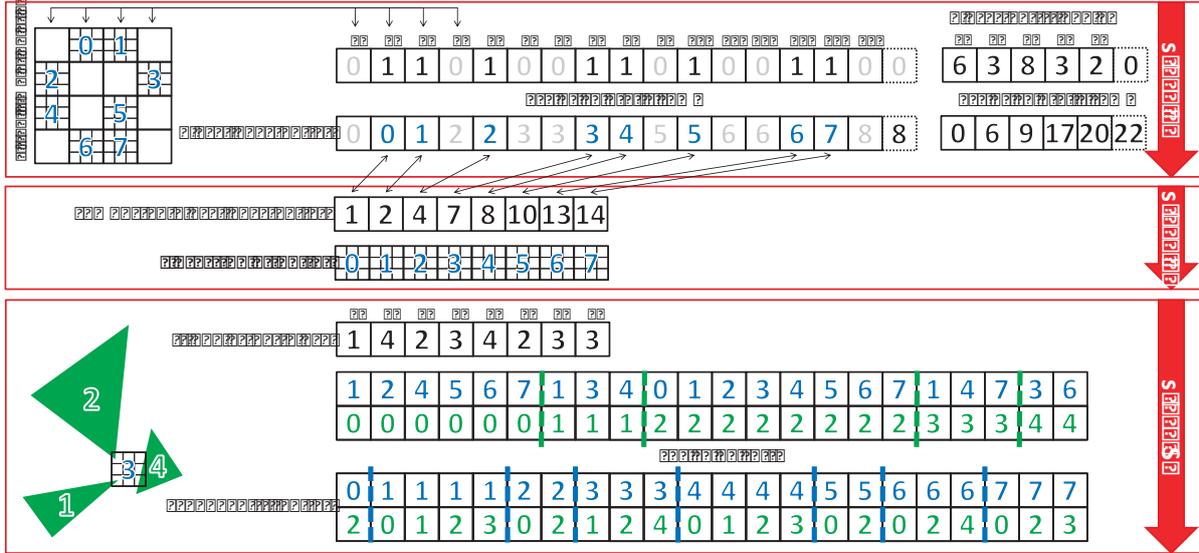


Figure 3: Overview of the different GPU buffers that are created to facilitate an efficient realization of the CPM.

thread computes the triangle’s axis aligned bounding box—enlarged about the width of the computational band—and computes the blocks that are intersected by this box. Since testing against the bounding box can result in false positives, i.e., blocks that are intersected by the bounding box but not the computational band, an additional test of the blocks is performed to prune as many of them as possible: For every intersected block, the thread calculates the distance from the block center to the triangle and skips the block if its bounding sphere does not intersect the triangle’s computational band. By writing a 1 into *block-index-buffer* at the positions of the remaining blocks, these blocks are marked as intersected. Because all threads write the same value, a synchronization of the write operations is not necessary in this case.

At the end of stage 1, each thread writes the number indicating how many blocks are intersected by the triangle to a temporary buffer—*blocks-per-triangle*—in global GPU memory. Finally, an exclusive parallel prefix sum [Har07] is computed in-place over *block-index-buffer* and *blocks-per-triangle*. At the positions that were marked, the *block-index-buffer* now contains the relative positions of the respective block in the sequence of marked blocks.

Stage 2: Since the prefix sum operation also calculates the total number of marked blocks, the *closest-point-buffer* can be allocated in GPU memory. Then, a CUDA kernel with one thread per entry in *block-index-buffer* is executed. The i -th thread reads the index value val from the buffer at position i and compares it to the value that is stored at position $i + 1$. If both values are different, the thread writes the value i into the *compacted-block-index-buffer* at position val .

Stage 3: We start by allocating the *expanded-tri-buffer* in GPU global memory using the size indicated by the prefix

sum over *blocks-per-triangle*. The elements of this buffer are pairs of triangle-block indices, and for each triangle these pairs are written in succession into the buffer. Since the parallel scan operation also gives the starting positions of each set of triangle-block pairs in *expanded-tri-buffer*, a CUDA kernel with one thread per triangle is executed to first build these sets in shared memory—including the computation of intersected blocks as described before—and then to write them into the buffer at the starting positions (see the first occurrence of *expanded-tri-buffer* in Figure 3). In the same kernel, we use CUDA’s *atomicAdd* operation to fill a buffer *triangles-per-block*, which stores for every marked block the number of triangles it intersects.

In a second pass, *expanded-tri-buffer* is sorted w.r.t. the block ID using the CUDA radix sort [Har07], and a prefix sum over *triangles-per-block* is calculated. From the content of the resulting buffers the indices of all triangles contributing to the closest point representation for a particular block can be determined.

4.3. Closest Point Computation

For each of the cells in *closest-point-buffer*, which are uniquely defined by the block they belong to and their relative position in the sequence of cells of this block, the point on the surface that is closest to the cell’s center has to be computed (we will call these centers the grid vertices from now on). The parallel CUDA implementation we propose is optimized to reduce global memory operations by performing a block-wise closest point computation in shared memory, and writing the results en block into global memory.

The CUDA kernel to compute the closest point representation executes one warp of CUDA threads per grid

block. Every thread computes the closest point of exactly one grid vertex, by sequentially going over the triangles that are stored in *expanded-tri-buffer* for that block. Thus, the threads of one warp process a number of grid vertices in parallel, in the same order they lie in memory, meaning that all required global memory buffers are accessed only by coalesced, fully utilized memory transactions. Furthermore, since the threads of a warp are executed in lock step and every grid vertex is processed by exactly one thread, all synchronization and atomic instructions can be omitted.

In shared memory, two buffers are allocated to store the closest point coordinates and the distances of the grid vertices to these points. All threads of one warp first perform a coalesced global memory read operation to fetch the triangles from *expanded-tri-buffer* that are needed to compute the closest points. The number of triangles is read from *triangles-per-block*.

Each thread computes the grid vertex for the cell it is working on. Then, it iterates over all triangles, and for each of them it computes the closest point and corresponding distance. If the distance is shorter than the one stored in the shared memory buffer, both the distance and the closest point coordinate are overwritten. The closest point computation is performed by first determining whether the grid vertex is closest to a corner, an edge, or the triangle's interior. By a coordinate system transformation that brings the triangle into one of the coordinate planes, this essentially breaks down to a 2D problem: According to the illustration in Figure 4, it has to be identified in which of the 7 zones of the triangle the projected vertex \mathbf{x}_0 lies. After all triangles have been processed, the shared memory buffers are copied in coalesced, fully utilized memory transactions to buffers in global memory.

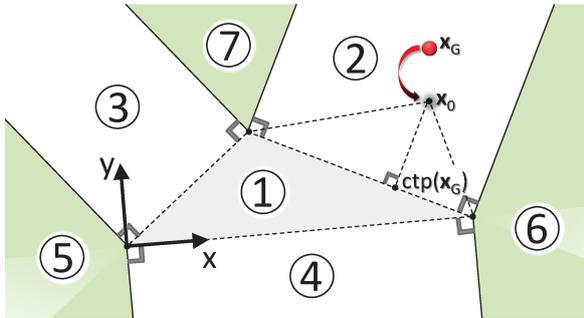


Figure 4: In each of the 7 zones, \mathbf{x}_G is closest to one region of the triangle (face, edge, corner).

5. Fluid Simulation

In the following we describe the simulation of fluid effects on surfaces using the CPM on the GPU. We focus on the numerical solution of two different PDEs describing such effects: the wave equation and the incompressible isothermal Navier-Stokes equations.

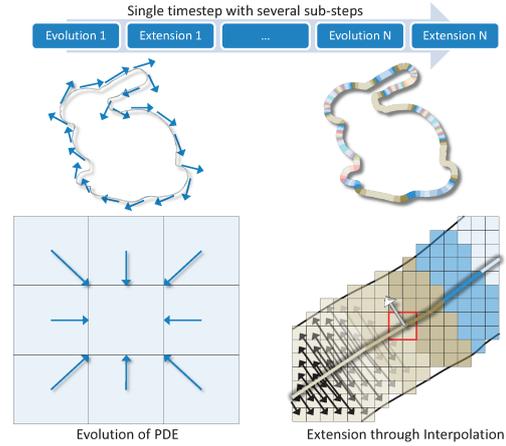


Figure 5: A time step may require several sub-steps, consisting of an evolution (left) and an extension phase (right).

Before a PDE solution can be computed via the CPM, appropriate initial conditions have to be specified (Section 5.1). Then, each solution step generally consists of two phases, both implemented on the GPU via CUDA: In the evolution phase, the solution of the embedding PDE (Section 5.2) in the next time step is calculated for all grid vertices in the evolution band. After the evolution, the closest point extension (Section 5.3) is applied to all vertices in the extension band. In this phase, the solution is propagated from the surface to the computational band by replacing the values at the grid vertices with values interpolated at the closest points on the surface.

For the wave equation (Section 5.4), the less involved of the two presented fluid models, just one evolution and one extension is necessary for each time step. In contrast, for the incompressible Navier-Stokes model (Section 5.5) a splitting method was used to handle the terms of the equations independently of each other. Consequently, in this model a time step requires several sub-steps and also more than one closest point extension (see Figure 5).

5.1. Initialization

The definition of the computational bands and the construction of the closest point representation (Section 4) are necessary to start the CPM. The initialization of the solution variable completes the initialization phase. In this step, initial data on the surface has to be extended to the simulation grid. When the surface is given as a triangle mesh, the initial data is typically given as per-vertex attributes or as a texture map. Every grid vertex in the extension band interpolates the surface data at the closest point and stores this value at its location in the simulation grid.

5.2. Embedding PDE

To compute the solution for the next time step, we run one CUDA thread per grid vertex inside the evolution band. For each term of the embedding PDE we choose an adequate evolution strategy. The embedding of terms of the PDE involving spatial differential operators up to order two is found by replacing surface gradients, surface divergences and Laplace–Beltrami operators with their standard Cartesian counterparts in \mathbb{R}^3 . The discretization can then be done via regular finite differences. The charming property of CPM in this case is that information is automatically propagated only in tangential directions along the surface. If we have an existing discretization for the problem of choice in \mathbb{R}^3 , we can reuse it without modification.

5.3. Closest Point Extension

The closest point extension is performed by one CUDA thread per grid point in the extension band. We retrieve the grid point’s closest point on the surface and resample the respective attribute at this position using an interpolation scheme. The value at the grid point is then replaced by the interpolated value, thereby creating a field that is constant in normal directions to the surface. The kind of interpolation is important and influences stability, quality and performance of the whole simulation.

A stable interpolation scheme is especially important when dealing with non-smooth surfaces. While the fundamental assumptions of the CPM are true only for smooth surfaces, it is certainly desirable to also handle C^0 surfaces (like in Figure 6 bottom). However, the closest point field for a C^0 surface can contain discontinuities even close to the surface. Through the closest point extension, this results in discontinuities also in the simulation attributes.

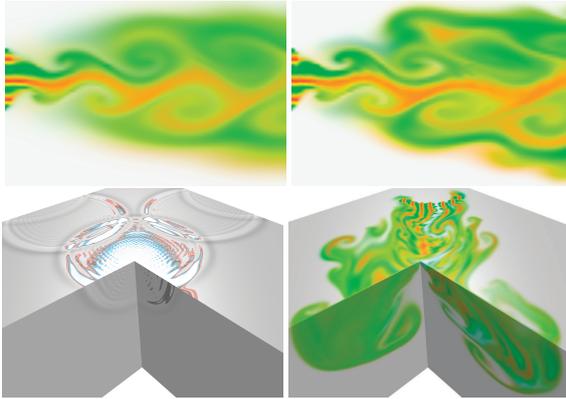


Figure 6: Top: With linear interpolation (left) the simulation shows artificial diffusion and loses energy. Both problems are not present with WENO4 interpolation (right). Bottom: While the CPM generally requires the surface to be smooth, with clamped WENO interpolation it also delivers reasonable results for C^0 surfaces.

submitted to COMPUTER GRAPHICS Forum (9/2011).

Linear interpolation is computationally cheap and fulfills the convex hull property, which guarantees stability of the interpolation even if the field is discontinuous. Unfortunately, at least for second-order PDE problems, linear interpolation is not accurate enough for the CPM to ensure consistency [RM08]. At any rate and in practice, linear interpolation introduces a high amount of artificial numerical diffusion, which manifests in a loss of energy in physics simulations (cp. Figure 6 top).

In the work of Macdonald and Ruuth [MR08], weighted essentially non-oscillatory (WENO) interpolation was presented. In smooth areas it is of higher order, which prevents numerical diffusion. As with most polynomial interpolations, stability can be critical near discontinuities where the polynomials overshoot the values in the stencil. To reduce this effect, the WENO algorithm calculates a weight for several candidate polynomial interpolants and takes a weighted sum of each interpolant. If a particular candidate is highly oscillatory in a given region it is assigned a very small weight. By this means oscillations are minimized in those areas, as is the formal order-of-accuracy. In the unlikely—yet unfortunately practically relevant—case that all polynomials oscillate, the WENO interpolation can still overshoot. To preserve stability in this event, we clamp the interpolation result to the minimum and maximum values in the stencil.

The desired interpolation order dictates the size of the stencil and thereby the bandwidth requirements. We use the WENO4 interpolation scheme which is built on quadratic candidate interpolants and recovers tri-cubic interpolation in smooth regions (tri-quadratic otherwise) using 64 entries in the stencil. It is sufficiently accurate for the CPM and we find it to be fast enough for an interactive application.

5.4. Wave equation

The wave equation is the classical example of a hyperbolic equation:

$$\frac{\partial^2 u}{\partial t^2} = c^2 \nabla_S^2 u$$

In our case u represents the height of an elastic surface over time t for a given wave speed c . Note that this is the surface-PDE version of the equation, which means that we can interpret u as the height above some original surface S . Because the right-hand side involves only the Laplace-Beltrami operator, we can simply replace it with the Cartesian Laplace operator to retrieve the embedding PDE for the CPM. We then obtain a discretization for a 3D Cartesian grid with spacing $\Delta \mathbf{x}$ and regular time intervals of length Δt by substituting the second order derivatives on both sides of the equation with finite differences. For each grid point (i, j, k) this leads to:

$$\frac{\tilde{u}_{i,j,k}^{t+1} - 2u_{i,j,k}^t + u_{i,j,k}^{t-1}}{\Delta t^2} = \frac{c^2}{\Delta \mathbf{x}^2} (u_{i+1,j,k}^t + u_{i-1,j,k}^t + u_{i,j+1,k}^t + u_{i,j-1,k}^t + u_{i,j,k+1}^t + u_{i,j,k-1}^t - 6u_{i,j,k}^t)$$

Solving for the unknown $\tilde{u}_{i,j,k}^{t+1}$ leads to the explicit solution

$$\begin{aligned} \tilde{u}_{i,j,k}^{t+1} = & \alpha \cdot (u_{i+1,j,k}^t + u_{i-1,j,k}^t + u_{i,j+1,k}^t + u_{i,j-1,k}^t \\ & + u_{i,j,k+1}^t + u_{i,j,k-1}^t) + (2 - 6\alpha) \cdot u_{i,j,k}^t - u_{i,j,k}^{t-1} \\ & \text{with } \alpha = \frac{\Delta t^2 \cdot c^2}{\Delta \mathbf{x}^2}, \end{aligned}$$

followed by a closest point extension

$$u_{i,j,k}^{t+1} := \tilde{u}^{t+1}(cp(\mathbf{x}_{i,j,k})).$$

Note that this is essentially a Verlet integration [Ver67] since the second derivatives in space as well as in time are approximated with central differences. Compared to an explicit Euler scheme, this method improves the accuracy of the integrator to 2nd order. Since the method is not self-starting we copy the initial conditions forward from u^0 to u^1 .

5.5. Incompressible Navier–Stokes equations

An incompressible, Newtonian fluid with constant temperature can be described by two continuity equations: One for momentum

$$\frac{\partial \mathbf{u}}{\partial t} = -(\mathbf{u} \cdot \nabla_S) \mathbf{u} - \frac{1}{\rho} \nabla_S p + \nu \nabla_S^2 \mathbf{u} + \mathbf{f}$$

and one for volume and mass

$$\nabla_S \cdot \mathbf{u} = 0.$$

In this variant of the Navier–Stokes equations—formulated as surface PDE— \mathbf{u} denotes the fluid’s velocity, ρ the constant density, p the pressure and ν the constant kinematic viscosity. Consequentially, the terms in the first equation represent accelerations resulting from self-advection $(-\mathbf{u} \cdot \nabla_S) \mathbf{u}$, pressure $(-\frac{1}{\rho} \nabla_S p)$, viscosity $(\nu \nabla_S^2 \mathbf{u})$ and external forces (\mathbf{f}), while the second equation states that the velocity field must be free of divergence. Instead of interpreting this flow problem as conservation laws with constraints, we split up the equations and treat each term individually.

The last two terms of the momentum equation, $\nu \nabla_S^2 \mathbf{u}$ and \mathbf{f} , are currently unrepresented in our solver. The viscosity term is skipped because we try to model nearly inviscid fluids. Simulation of viscous flows would be straight-forward, however, because one can simply replace the surface Laplacian ∇_S^2 with the standard Cartesian operator to obtain the embedding PDE for this term. The external forces \mathbf{f} would also be trivial to add but are not considered in the following.

For the nonlinear convective acceleration $(-\mathbf{u} \cdot \nabla_S) \mathbf{u}$, we use a semi-Lagrange back-trace as, e.g., described by Stam [Sta99]. By this means a computationally more expensive implicit solution can be avoided. To update the value of \mathbf{u} at each computational grid vertex, the back-trace follows the motion of a particle starting at the vertex backwards in time through the velocity field to find the previous location of the particle. At this position, the old velocity field is sampled and the retrieved value replaces \mathbf{u} at the grid vertex.

This back-trace could naïvely be implemented in a full simulation grid using its 3D equivalent. In this case, only one change is necessary in order to discretize the embedding of the surface-advection with the back-trace: as the velocity vector must stay in tangent space also at its new position, it must be projected onto the local tangent plane.

For the sparse grid discussed here, however, it must be ensured that such a back-trace does not leave the computational band. This results in additional constraints on the maximal velocity, the integration time step, and the minimal size of the extension band. To minimize these limitations we perform the back-trace only for points on the surface. Figure 7 illustrates this method.

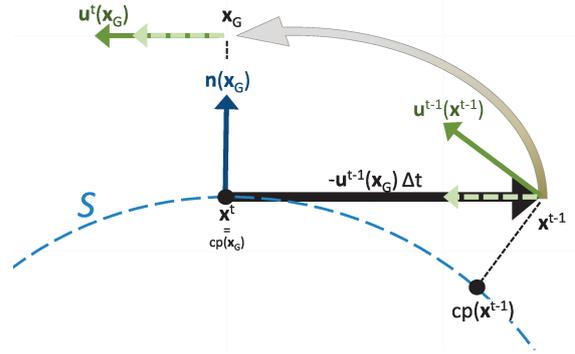


Figure 7: Velocity self-advection through a semi-Lagrange back-trace in the simulation grid.

The back-trace starts at the closest point of the current grid vertex $cp(\mathbf{x}_G)$ and uses a single step of the explicit Euler method. Because the current velocity values are a closest point extension, they are constant in the normal direction and for the start of the back-trace we can directly use the velocity value stored at the grid vertex \mathbf{x}_G . At the old particle location \mathbf{x}^{t-1} we sample the velocity field using an interpolation. To avoid numerical diffusion due to the interpolation in the back-trace, we use an interpolation of higher order like the one we described in the context of the closest point extension. The interpolated velocity vector from the particle’s old position is projected onto the tangent plane at the new position and then rescaled to its old length in order to retrieve the updated value for the grid vertex.

To obtain the surface normal necessary for such a projection at a grid point \mathbf{x}_G , one can normalize the difference vector $\mathbf{d} = \mathbf{x}_G - cp(\mathbf{x}_G)$ from the closest surface point to the point itself. These normals always point away from the surface on both of its sides. For non-orientable surfaces, for which both directions are equally valid, this behavior is thoroughly intended. A real problem with this approach, however, is that it fails for grid points very close to or on the surface. To reliably obtain a normal as well for these points, we

therefore examine the grid points \mathbf{x}_i , $i = 0, 1, \dots$ in a small stencil around \mathbf{x}_G and calculate a vector $\mathbf{d}_i = \mathbf{x}_i - cp(\mathbf{x}_i)$ for each of them. We accumulate the differences in the vector \mathbf{d}_S^n as follows:

$$\begin{aligned} \mathbf{d}_S^0 &= \mathbf{d}_0 \\ \mathbf{d}_S^{i+1} &= \mathbf{d}_S^i + \begin{cases} \mathbf{d}_{i+1} & \text{if } \mathbf{d}_{i+1} \cdot \mathbf{d}_S^i >= 0 \\ -\mathbf{d}_{i+1} & \text{otherwise} \end{cases} \end{aligned}$$

The negation is necessary to account for grid points on different sides of the surface. A surface normal is then retrieved by normalization of \mathbf{d}_S . Note that the resulting normal points to the side of the surface on which the first off-surface vertex in the stencil lies. In practice, we find using 8 points around \mathbf{x}_G works well.

In fluid mechanics, the pressure term $-\frac{1}{\rho} \nabla_S p$ is often used within a projection method [CM00] to ensure that the resulting velocity field fulfills the incompressibility equation $\nabla_S \cdot \mathbf{u} = 0$. Conceptually, this is achieved by applying a Hodge decomposition to the velocity field, which splits it into a divergence-free field and a curl-free field. The latter one can either be dismissed or it can be used for vorticity confinement [Sta99].

We determine a p that lets $\frac{1}{\rho} \nabla_S p$ reproduce the curl-free part, such that after the subtraction only the divergence-free part is left. At every time step, this ultimately requires to calculate the divergence of the current velocity and solving a surface Poisson problem of the form $\nabla_S \cdot \nabla_S p = \nabla_S \cdot \mathbf{u}$ to retrieve the desired pressure. Treating this Poisson equation with an artificial time iteration and the explicit CPM would require a closest point extension after each step, which makes this approach quite costly. Another option would be applying the implicit CPM [MR09], either to an artificial time iteration or directly to the Poisson problem.

We choose a cheaper computational approach where we apply the closest point extension only to the right-hand side of the equation, and we replace the left-hand side with the simpler Cartesian Laplacian. Then, we solve the linear system with the conjugate gradients method. While this may not enforce incompressibility to a high order of accuracy, in practice it gives visually plausible results if we choose a Neumann boundary condition for the values outside the evolution band and reuse the pressure from the last time step as the initial guess. After the pressure update, a closest point extension must also be applied to the velocity field itself to prepare it for the next time step.

In order to visualize the fluid flow we introduce an additional, final sub-step in which the advection of a mass-less, colored dye through the velocity field is simulated. The dye is transported through the flow by the same advection operator $\mathbf{u} \cdot \nabla_S$ that is used for the velocity. The only difference is that the dye is represented by three scalar fields—one for each *rgb* color channel—as such the operator does not include a projection onto the tangent plane.

6. Rendering

For rendering the simulated fluid effects on the surface we use Direct3D 10, and we make use of CUDA's interoperability functions to access the respective GPU memory resources within both APIs. If the surface is given as a triangle mesh and the simulated quantities are used to modulate its color, the mesh is rendered via polygon rasterization. The coordinates of the triangle vertices in the embedding 3D domain are interpolated by the rasterizer and then a pixel shader is employed, to retrieve values from the simulation buffers using trilinear interpolation. Figure 8 (left) demonstrates such a rendering with flat shading to show the triangulation, and in Figure 1 (top) Phong shading was used to realize a smooth look of the surface.



Figure 8: Navier-Stokes simulation on a Möbius strip; a non-orientable manifold. Left: Rendering with rasterization and flat shading to emphasize the triangle mesh serving as the original surface description. Right: Same simulation as left, but now rendered via raycasting and using the red ink as a displacement on both sides of the surface.

6.1. Volume Raycasting

If the initial surface is not given as a polygon mesh, for instance if it is given by a point set representation or an implicit surface description, or if the simulated quantities should be used to modulate the geometry of the surface, rasterization cannot be used any more. To overcome this problem, we present a GPU method that renders the surface directly from the sparse closest point volume representation, regardless of the initial surface representation.

A first approach is to perform exact voxel ray-casting of the closest point volume in a DDA-like fashion [Bre65]. The voxels correspond to the closest point cells, and the constant voxel attributes are given by the simulation values at the cell vertices. The rays of sight first traverse the grid of blocks until a marked block is hit. Then, the block's location in the *closest-point-buffer* is retrieved and traversal is continued on the cells inside. Traversal is stopped if a voxel is hit that contains the closest surface point stored at this voxel, and the voxel attribute is used as pixel color.

While the proposed technique can visualize a piece-wise (per voxel) constant distribution of attributes across a surface, it does not allow rendering a smooth distribution. This is achieved by interpolating the distances of the cell vertices to their closest surface points and performing an exact intersection test between the rays and the zero crossings in the reconstructed field.

Since the CPM is based solely on the existence of an *unsigned* distance transform, the zero crossings cannot be calculated exactly, and the intersection test comes down to finding the points along the rays where the distance between the sample position and the trilinearly interpolated closest point is below a given ϵ . At this point, the simulation attributes are retrieved via trilinear interpolation, too. To simulate smooth shading, per-pixel normals are calculated at each intersection point by normalizing the vector between the point and the closest surface point.

6.2. Rendering Surface Displacements

In addition to using the closest point representation for rendering the original surface, we can also use it to for simulating surface displacements that are given, for instance, by the simulation values at the grid vertices (see Figures 1 bottom, and 8 right). After interpolating the closest point $cp(\mathbf{x}_R)$ at a sample position \mathbf{x}_R on a ray, the simulation value $h(cp(\mathbf{x}_R))$ at this point is interpolated and compared to the distance between $cp(\mathbf{x}_R)$ and \mathbf{x}_R . As Figure 9 shows, a hit is detected if the distance value is smaller than the interpolated attribute value. The figure also shows, that the interpolation of closest points from different parts of the surface can lead to off-surface points. We take care of this case by comparing the bounding box of the closest points to the interpolation cell. The normal of the displaced surface is calculated from the original surface normal and the gradient of the attribute field. Since the gradient is tangent to the surface because of the CPM, it can be used in turn to disturb the surface normal appropriately. The maximum displacement is supposed to not exceed the width of the computational band. This constraint can be abandoned, however, by using an enlarged displacement band, where the grid vertices outside the computational band store only their closest points but no attributes.

The proposed rendering approach displaces the surface equally on both of its sides. For orientable surfaces it is also possible to distinguish the two half-spaces in order to displace in just one direction or to use negative displacements. For non-orientable surfaces the view-direction can be used to introduce an artificial orientation. Despite of being view-dependent, however, this approach has the additional problem that it results in artifacts at the silhouettes, where the orientation flips. We therefore recommend to treat both sides equally in this case.

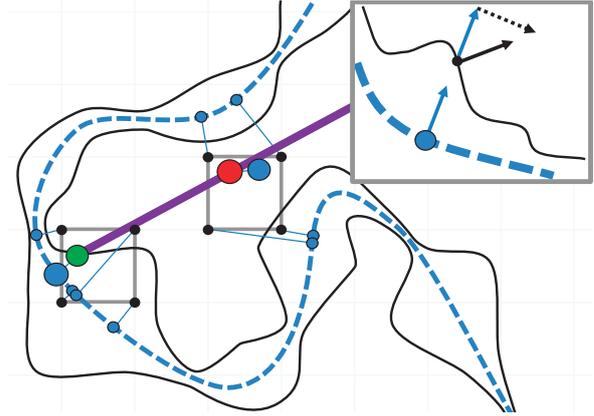


Figure 9: Bottom left: Intersection of a ray (purple) with a displaced surface (black). At the red ray sample, the interpolation of closest points on the original surface (blue) results in an off-surface point. At the green ray sample, an intersection with the displaced surface was found, as the distance to the closest point equals the interpolated height value. Top right: The normal of the original surface (blue arrows) combined with the gradient of the height value (dotted arrow) leads to the normal of the displaced surface (black arrow).

7. Performance Analysis

To validate the efficiency of the proposed method for simulating and rendering fluid effects on surfaces, we have performed a number of experiments using surface models and simulation grids at different resolutions. In all of our experiments, a block size of $b = 4$ was used. We found this size to give the best performance compared to smaller (tighter fit of the adaptive simulation grid to the surface but increasing number of indirections to access adjacent grid cells) and larger (increasing memory and computation requirements) blocks. All measurements were performed on a 2.4 GHz Core 2 Duo processor and an NVIDIA GeForce GTX 480 graphics card with 1536 MB local video memory. A detailed memory and performance statistic of the GPU CPM for fluid simulation is given in Table 1. Here, numbers separated by a slash refer to the simulation using the wave equation with linear interpolation and the Navier-Stokes equations with WENO4 interpolation, respectively. All timings are given in milliseconds and rendering was always performed on a $1K \times 1K$ viewport.

For each model, the first column lists the name and number of triangles, while the second column gives the resolution of the uniform Cartesian grid to which the simulation resolution corresponds. The third column lists the number of closest points within the computational band. The differences are due to differently sized computational bands that are dictated by the respective numerical stencils of the wave equation (bandwidth with linear interpolation: $2.4\Delta\mathbf{x}$) and the Navier-Stokes equations (bandwidth with WENO4

interpolation: $5.7\Delta x$). The fourth column gives the memory that is required for solving the equations on the GPU via the CPM. The higher memory consumption of the Navier-Stokes simulation is due to the larger computational band and the additional buffers that are required to store the simulation quantities. The times required to build the closest point representation are given in the fifth column. Here it is important to note that more than 75% of the time is always required by the closest points computation, meaning that the grid construction on the GPU consumes only a small portion of this time. It can further be seen that the GPU memory requirements mainly depend on the resolution of the simulation grid, while on the other hand, the time for constructing the grid is strongly dependent on the number of triangles.

The following four columns list the simulation times using linear and WENO4 interpolation in the closest point extension for the wave equation and the Navier-Stokes equations. As expected, WENO4 interpolation increases the overall simulation times significantly due to its larger computational stencil, and the enlarged computational band thereof. The Navier-Stokes simulation uses the respective interpolation type in three separate closest point extensions (divergence, velocity, ink) and two advection passes (velocity, ink). It can be observed that solving the pressure Poisson equation is the most expensive operation ($\sim 75\%$ of the total time) as long as linear interpolation is used. With WENO4 interpolation, however, these ratios turn into the opposite: Most of the time is now spent on interpolations in the advection ($\sim 50\%$) and extension kernels ($\sim 40\%$).

The last three columns show the rendering times for rasterization, voxel-based surface raycasting, and raycasting with surface displacements. The timings refer to the rendering of a WENO4 wave simulation. In the last column a computational band of width $8\Delta x$ was used to allow simulating large surface displacements. As expected, rasterizing the triangle mesh on the GPU leads to the highest frame rates. Direct surface raycasting is between 1.2 to 5 times slower, but still delivers highly interactive frame rates. Raycasting with smooth displacements, which allows simulating dynamic surface modifications, also delivers interactive frame rates even for the computational grids with the highest resolution. This is quite remarkable since much larger closest point and simulation buffers are used.

7.1. Surface Deformations

The timing statistics indicate that the creation of an adaptive closest point grid for the CPM is possible at high rates even for large triangular meshes. This is an important step towards a real-time CPM for simulating deformations of the surface itself. In the following we demonstrate mesh smoothing based on the Laplace operator [Tau95] via the GPU CPM as a first example (see Fig. 10).

Each smoothing step thereby consists of three operations: Firstly, we create an discrete, adaptive closest point

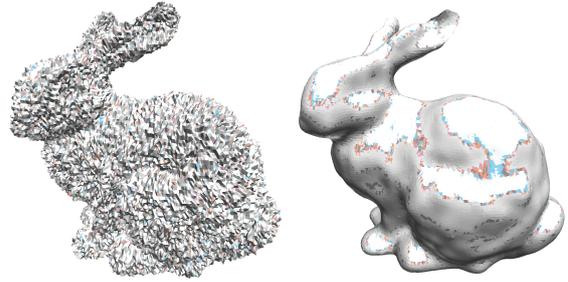


Figure 10: Laplacian mesh smoothing using the CPM. Left: Flat shaded, distorted Bunny model. Right: The model after 15 smoothing steps (64ms each) using a grid resolution of 256.

field as described. Secondly, we apply the Laplace-Beltrami operator—discretized with the standard Laplace operator on the embedding Cartesian grid—to the closest point field. At every point on the surface the resulting vector field is directed towards the mean of the point’s neighborhood on the surface. Thirdly, we interpolate the discrete field using WENO4 at every vertex position of the mesh and move each vertex into this direction.

8. Conclusion and Future Work

We have presented a real-time method for simulating fluid effects on surfaces using the CPM. To achieve this, we have developed a fast GPU method for realizing the CPM and rendering the simulation results as dynamic color or displacement effects. The method can efficiently solve non-linear PDEs on surfaces using a high-resolution embedding grid, and it can calculate and render the resulting fields independently of the resolution and topology of the input geometry.

The proposed method opens a number of future research directions. In particular, we would like to adopt the CPM for real-time effects on animated surfaces. So far, we are able to construct an adaptive embedding grid for the deforming surface in every frame. Further research is necessary, however, to define a CPM for PDEs on deforming surfaces and to develop an efficient solution for the frame-to-frame information flow. Due to the independence of the CPM and our rendering technique from the original surface representation, such a method would also allow simulating complex deformations of the surface itself.

From a numerical point of view we are most interested in exploring how multigrid schemes can be employed to speed up the simulation on the embedding grid. This requires pursuing research on the construction of a multigrid closest point representation that can accurately approximate the surface at ever coarser scales. By integrating such a construction into approaches that can create a coarse grid hierarchy independently of the complexity of the object’s shape [LPR*09,DGW11], good approximation quality and convergence rates can be expected even for complicated surfaces.

Table 1: Performance statistics for fluid simulation and rendering on the GPU via the CPM.

Mesh	Res.	Grid Generation			Wave Equation		Navier-Stokes		Rendering		Displ.
		# Closest Points	GPU Memory	Time	linear	WENO4	linear	WENO4	Rastization	Raycasting	
Möbius (Figure 8) 3,385Δ	64	34k / 48K	1MB / 4MB	6.1ms / 9.5ms	2.0ms	2.1ms	3.7ms	20ms	1.3ms	3.8ms	20ms
	128	77k / 184k	4MB / 18MB	6.4ms / 10ms	2.1ms	4.1ms	7.5ms	65ms	1.3ms	4.3ms	23ms
	256	303k / 721k	18MB / 73MB	7.4ms / 12ms	2.3ms	15ms	23ms	237ms	1.4ms	5.5ms	31ms
	512	1.2M / 2.9M	82MB / 297MB	12ms / 21ms	5.1ms	54ms	85ms	909ms	1.7ms	8.4ms	48ms
Bunny (Figure 1) 69,451Δ	64	41k / 80k	2MB / 7MB	29ms / 84ms	2.2ms	3.0ms	4.0ms	28ms	1.4ms	3.3ms	15ms
	128	177k / 374k	9MB / 35MB	32ms / 86ms	2.4ms	9.0ms	13ms	118ms	1.4ms	4.0ms	24ms
	256	733k / 1.6Mk	41MB / 159MB	34ms / 90ms	3.6ms	34ms	51ms	529ms	1.4ms	5.7ms	35ms
	512	3.0M / 6.8M	178MB / 679MB	49ms / 117ms	11ms	136ms	222ms	2174ms	1.7ms	8.6ms	52ms
Armadillo (Figure 1) 345,944Δ	64	28k / 58k	1MB / 5MB	117ms / 377ms	2.0ms	2.7ms	3.4ms	18ms	2.5ms	3.0ms	16ms
	128	123k / 258k	6MB / 25MB	118ms / 380ms	2.1ms	6.1ms	9.5ms	80ms	2.5ms	3.7ms	20ms
	256	513k / 1.1M	29MB / 110MB	120ms / 384ms	2.9ms	25ms	38ms	370ms	2.5ms	4.6ms	28ms
	512	2.0M / 4.8M	129MB / 477MB	135ms / 404ms	8.6ms	100ms	157ms	1667ms	2.6ms	5.9ms	45ms

Acknowledgements

This publication was based on work supported in part by the Munich Centre of Advanced Computing at the Technische Universität München (TUM) and by Awards No. KUK-C1-013-04 and UK-C0020, made by King Abdullah University of Science and Technology (KAUST).

References

- [ATBG08] ANGST R., THÜREY N., BOTSCH M., GROSS M.: Robust and efficient wave simulations on deforming meshes. *Computer Graphics Forum (Proc. Pacific Graphics)* 27 (2008), 1895–1900. 3
- [BCOS01] BERTALMIÓ M., CHENG L.-T., OSHER S., SAPIRO G.: Variational problems and partial differential equations on implicit surfaces. *J. Comput. Physics* 174, 2 (2001), 759–780. 1
- [Bre65] BRESENHAM J. E.: Algorithm for computer control of a digital plotter. *IBM Systems Journal* 4, 1 (1965), 25–30. 9
- [CLB*09] CHUANG M., LUO L., BROWN B. J., RUSINKIEWICZ S., KAZHDAN M.: Estimating the laplace-beltrami operator by restricting 3D functions. In *Proceedings of the Eurographics Symposium on Geometry Processing* (2009), pp. 1475–1484. 1, 3
- [CM00] CHORIN A., MARSDEN J.: *A Mathematical Introduction to Fluid Mechanics*, 4th ed. Springer Verlag, 2000, ch. 1.3, pp. 36–44. 9
- [CRT04] CLARENZ U., RUMPF M., TELEA A.: Surface processing methods for point sets using finite elements. *Computers and Graphics* 28, 6 (2004), 851–868. 1
- [DE07] DZIUK G., ELLIOTT C.: Surface finite elements for parabolic equations. *J. Comput. Math.* 25 (2007), 385–407. 1
- [DGW11] DICK C., GEORGII J., WESTERMANN R.: A hexahedral multigrid approach for simulating cuts in deformable objects. *IEEE Transactions on Visualization and Computer Graphics* 17 (2011), 1663–1675. 11
- [ED08] EISEMANN E., DÉCORET X.: Single-pass gpu solid voxelization for real-time applications. In *Proceedings of graphics interface 2008* (Toronto, Ont., Canada, Canada, 2008), GI '08, Canadian Information Processing Society, pp. 73–80. 3
- [FZKH05] FAN Z., ZHAO Y., KAUFMAN A., HE Y.: Adapted unstructured lbm for flow simulation on curved surfaces. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2005), SCA '05, ACM, pp. 245–254. 3
- [Har07] HARRIS M.: Parallel prefix sum (scan) with CUDA. NVIDIA Whitepaper, April 2007. <http://www.nvidia.com/developer.5>
- [HZQW10] HONG Y., ZHU D., QIU X., WANG Z.: Geometry-based control of fire simulation. *The Visual Computer* 26 (September 2010), 1217–1228. 3
- [JBS06] JONES M. W., BÄRENTZEN J. A., SRAMEK M.: 3d distance fields: A survey of techniques and applications. *IEEE Transactions on Visualization and Computer Graphics* 12 (2006), 581–599. 3
- [LFW07] LO K.-Y., FU C.-W., WONG T.-T.: Interactive reaction-diffusion on surface tiles. In *Proceedings of Pacific Graphics* (2007), pp. 65–74. 1
- [LPR*09] LIEHR F., PREUSSER T., RUMPF M., SAUTER S., SCHWEN L. O.: Composite finite elements for 3d image based computing. *Comput. Vis. Sci.* 12 (March 2009), 171–188. 11
- [LWC05] LUI L. M., WANG Y., CHAN T. F.: Solving PDEs on manifolds with global conformal parametrization. In *VLSM* (2005), pp. 307–319. 1
- [MR08] MACDONALD C. B., RUUTH S. J.: Level set equations on surfaces via the Closest Point Method. *J. Sci. Comput.* 35, 2–3 (June 2008), 219–240. doi:10.1007/s10915-008-9196-6. 3, 4, 7
- [MR09] MACDONALD C. B., RUUTH S. J.: The implicit Closest Point Method for the numerical solution of partial differential equations on surfaces. *SIAM J. Sci. Comput.* 31, 6 (2009), 4330–4350. doi:10.1137/080740003. 3, 9
- [NMZ07] NEILL P., METOYER R., ZHANG E.: Fluid flow on interacting deformable surfaces. In *ACM SIGGRAPH 2007 posters* (2007), SIGGRAPH '07, ACM. 3
- [NNRW09] NEMITZ O., NIELSEN M. B., RUMPF M., WHITAKER R.: Finite element methods on very large,

- dynamic tubular grid encoded implicit surfaces. *SIAM J. on Scientific Computing* 31, 3 (2009), 2258–2281. 1
- [Pan11] PANTALEONI J.: Voxelpipe: a programmable pipeline for 3d voxelization. In *Proceedings of the ACM SIGGRAPH Symposium on High Performance Graphics* (New York, NY, USA, 2011), HPG '11, ACM, pp. 99–106. 3, 4
- [RM08] RUUTH S. J., MERRIMAN B.: A simple embedding method for solving partial differential equations on surfaces. *J. Comput. Phys.* 227 (January 2008), 1943–1961. 1, 2, 3, 7
- [RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-beltrami spectra as "shape-dna" of surfaces and solids. *Computer-Aided Design* 38, 4 (2006), 342–366. 3
- [SS10] SCHWARZ M., SEIDEL H.-P.: Fast parallel surface and solid voxelization on gpus. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 29 (2010), 179:1–179:9. 3, 4
- [Sta99] STAM J.: Stable fluids. *SIGGRAPH '99*, ACM, pp. 121–128. 2, 8, 9
- [Sta03] STAM J.: Flows on surfaces of arbitrary topology. *ACM Trans. Graph.* 22 (July 2003), 724–731. 1, 2
- [SY04] SHI L., YU Y.: Inviscid and incompressible fluid simulation on triangle meshes: Research articles. *Comput. Animat. Virtual Worlds* 15 (July 2004), 173–181. 3
- [Tau95] TAUBIN G.: A signal processing approach to fair surface design. *SIGGRAPH '95*, ACM, pp. 351–358. 11
- [TMR09] TIAN L., MACDONALD C. B., RUUTH S. J.: Segmentation on surfaces with the Closest Point Method. In *Proc. ICIP09, 16th IEEE International Conference on Image Processing* (Cairo, Egypt, 2009), pp. 3009–3012. 3
- [Tur91] TURK G.: Generating textures on arbitrary surfaces using reaction-diffusion. *SIGGRAPH '91*, ACM, pp. 289–298. 1, 3
- [Ver67] VERLET L.: Computer "experiments" on classical fluids. i. thermodynamical properties of Lennard-Jones molecules. *Phys. Rev.* 159, 1 (Jul 1967), 98. 8
- [WMT07] WANG H., MILLER G., TURK G.: Solving general shallow wave equations on surfaces. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation* (Aire-la-Ville, Switzerland, Switzerland, 2007), SCA '07, Eurographics Association, pp. 229–238. 3

RECENT REPORTS

33/11	Data assimilation using bayesian filters and B-spline geological models	Duan Farmer Hoteit Lu Moroz
34/11	Review of nonlinear Kalman, ensemble and particle filtering with application to the reservoir history matching problem	Luo Hoteit Duan Wang
35/11	Modelling a Tethered Mammalian Sperm Cell undergoing Hyper-activation	Curtis Kirkman-Brown Connolly Gaffney
36/11	A simple mathematical model for investigating the effect of cluster roots on plant nutrient uptake	Zygalakis Roose
37/11	Frequency jumps in the planar vibrations of an elastic beam	Neukirch Frelat Goriely Maurini
38/11	Ice-lens formation and con- nemenent-induced supercooling in soils and other colloidal materials	Style Cocks Peppin Wettlaufer
39/11	An asymptotic theory for the re-equilibration of a micellar surfac- tant solution	Griffiths Bain Breward Chapman Howell Water
40/11	Higher-order numerical methods for stochastic simulation of chemical reaction systems	Székely Burrage Erban Zygalakis
41/11	On the modelling and simulation of a high pressure shift freezing process	Smith Peppin Ángel M. Ramos
42/11	An efficient implementation of an implicit FEM scheme for fractional-in-space reaction-diffusion equations	Burrage Hale Kay
43/11	Coupling fluid and solute dynamics within the ocular surface tear film: a modelling study of black Line osmolarity	Zubkov Breward Gaffney
44/11	A prototypical model for tensional wrinkling in thin sheets	Davidovitch Schroll Vella Adda-Bedia Cerde

47/11	Effects of demographic noise on the synchronization of a metapopulation in a fluctuating environment	Lai Newby Bressloff
48/11	High order weak methods for stochastic differential equations based on modified equations	Abdulle Cohen Vilmart Zygalakis
49/11	The kinetics of ice-lens growth in porous media	Style Peppin
50/11	Wound healing angiogenesis: the clinical implications of a simple mathematical model	Flegg Byrne Flegg McElwain
51/11	Wound healing angiogenesis: the clinical implications of a simple mathematical model	Du Gunzburger Lehoucq Zhou
52/11	Image inpainting based on coherence transport with adapted distance functions	März
53/11	Surface growth kinematics via local curve evolution	Moulton Goriely
54/11	A multiple scales approach to evaporation induced Marangoni convection	Hennessey Münch
55/11	The dynamics of bistable liquid crystal wells	Luo Majumdar Erban

Copies of these, and any other OCCAM reports can be obtained from:

**Oxford Centre for Collaborative Applied Mathematics
Mathematical Institute
24 - 29 St Giles'
Oxford
OX1 3LB
England**

www.maths.ox.ac.uk/occam