

Curl-Flow: Boundary-Respecting Pointwise Incompressible Velocity Interpolation for Grid-Based Fluids

JUMYUNG CHANG, NVIDIA
RUBEN PARTONO, University of Waterloo
VINICIUS C. AZEVEDO, ETH Zurich
CHRISTOPHER BATTY, University of Waterloo




Fig. 1. **Effect of Interpolants on Particle Distribution in 3D:** The results of identical, initially uniform dense particle distributions advected through the same $10 \times 10 \times 10$ discretely incompressible field for 160 frames using different velocity interpolants. Left is standard trilinear velocity interpolation, center is monotonic cubic velocity interpolation [Fritsch and Carlson 1980], and right is our Curl-Flow method. The coloring indicates the per-particle local density estimate, with yellow/white being highest density. Basic polynomial interpolants lead to undesirable severe clustering and spreading due to grid-scale divergence, while our pointwise incompressible Curl-Flow interpolation preserves close-to-uniform particle distributions over long time periods.

We propose to augment standard grid-based fluid solvers with *pointwise* divergence-free velocity interpolation, thereby ensuring exact incompressibility down to the sub-cell level. Our method takes as input a discretely divergence-free velocity field generated by a staggered grid pressure projection, and first recovers a corresponding discrete vector potential. Instead of solving a costly *vector* Poisson problem for the potential, we develop a fast parallel sweeping strategy to find a candidate potential and apply a gauge transformation to enforce the Coulomb gauge condition and thereby make it numerically smooth. Interpolating this discrete potential generates a pointwise vector potential whose analytical curl is a pointwise incompressible velocity field. Our method further supports irregular solid geometry through the use of level set-based cut-cells and a novel Curl-Noise-inspired potential ramping procedure that simultaneously offers strictly non-penetrating velocities and incompressibility. Experimental comparisons demonstrate that the vector potential reconstruction procedure at the heart of our approach is consistently faster than prior such reconstruction schemes, especially those that solve vector Poisson problems. Moreover, in exchange for its modest extra cost, our overall *Curl-Flow* framework produces significantly improved particle trajectories that closely respect irregular obstacles, do not suffer

from spurious sources or sinks, and yield superior particle distributions over time.

CCS Concepts: • **Computing methodologies** → **Physical simulation**.

Additional Key Words and Phrases: divergence-free, stream function, vector potential, velocity interpolation, advection

ACM Reference Format:

Jumyung Chang, Ruben Partono, Vinicius C. Azevedo, and Christopher Batty. 2022. Curl-Flow: Boundary-Respecting Pointwise Incompressible Velocity Interpolation for Grid-Based Fluids. *ACM Trans. Graph.* 41, 6, Article 243 (December 2022), 21 pages. <https://doi.org/10.1145/3550454.3555498>

1 INTRODUCTION

The assumption of incompressibility is pervasive in computer animation of fluids. Since compressive effects are imperceptible in many (but of course not all) visually relevant liquid and gas scenarios, neglecting fast-moving compression waves is often justifiable in practice and yields significant efficiency gains. Mathematically, incompressibility implies that the fluid velocity field \mathbf{u} should be divergence-free: $\nabla \cdot \mathbf{u} = 0$. Popular staggered grid-based schemes rely on this assumption, using finite difference or finite volume ideas, combined with Lagrangian or semi-Lagrangian advection methods, to transform the continuous incompressible flow equations into discrete, computable algorithms [Bridson 2015]. This fertile mathematical soil has sprouted diverse numerical tools for visual simulation of drifting cigarette smoke, coiling honey, crashing ocean waves, and more [Enright et al. 2002; Fedkiw et al. 2001; Larionov et al. 2017], which are widely integrated into industrial software like *Houdini*, *Maya*, and *Blender*. Yet despite the widespread use of

Authors' addresses: Jumyung Chang, NVIDIA, Toronto, jumyungc@nvidia.com; Ruben Partono, University of Waterloo, Waterloo, rparton@uwaterloo.ca; Vinicius C. Azevedo, ETH Zurich, Zurich, vinicius.azevedo@inf.ethz.ch; Christopher Batty, University of Waterloo, Waterloo, christopher.batty@uwaterloo.ca.

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 the author(s) 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.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2022/12-ART243 \$15.00 <https://doi.org/10.1145/3550454.3555498>

grid-based incompressible flow animation techniques, the pointwise velocity vector fields they produce are in fact *not incompressible*.

To understand this potentially surprising statement, one must distinguish *discrete* incompressibility from its *continuous* counterpart. Under a finite volume approach, the discrete velocity components, stored at cell face midpoints, will indeed satisfy discrete incompressibility: the sum of discrete fluxes across each cell's boundary is zero. However, interpolation is often required to provide *pointwise* velocity values everywhere in the simulation domain, to support popular Lagrangian or semi-Lagrangian discretizations of advection for passive tracer particles, density, velocity, temperature, and so on [Jiang et al. 2015; Stam 1999; Zhu and Bridson 2005]. Applying basic polynomial interpolants on discretely incompressible grid data affords no guarantee that the interpolated velocity fields will be *pointwise analytically incompressible*, and in practice they are not.

Spurious compressibility has non-negligible implications. Advecting particles through vector fields with artificial sources and sinks damages volume conservation and causes uniformly distributed particles to clump and spread. Alternatively, reducing discretization error by significantly increasing grid resolution fails to address the root cause and is too costly regardless: simulation time typically scales at least quartically with grid resolution. Irregular solid boundaries further exacerbate the issue because naive grid-based interpolants are essentially oblivious to obstacles; obstacle-aware cut-cell interpolants [Azevedo et al. 2016] improve the no-normal-flow enforcement at the cost of worsening compression artifacts.

To guarantee an analytically divergence-free interpolated velocity field by construction, we instead form the velocity field \mathbf{u} from the curl ($\nabla \times$) of a second vector field, $\boldsymbol{\psi}$, known as a *vector potential*:

$$\mathbf{u} = \nabla \times \boldsymbol{\psi}. \quad (1)$$

In 2D, $\boldsymbol{\psi}$ reduces to a scalar *stream function*, ψ . Incompressibility is thus enforced by a basic vector calculus identity,

$$\nabla \cdot \mathbf{u} = \nabla \cdot \nabla \times \boldsymbol{\psi} = 0. \quad (2)$$

We exploit this relationship in the discrete and continuous settings to ensure pointwise incompressibility through three steps:

- (1) Given discretely incompressible fluid velocity values, construct corresponding discrete vector potential values.
- (2) Interpolate the discrete vector potential values to yield a pointwise vector potential field.
- (3) Take the analytical curl of the interpolated vector potential field to yield a *pointwise incompressible* velocity field.

Bao et al. [2017] first proposed this three-step recipe for deriving discrete delta functions in an immersed boundary method for weakly coupled fluid-structure interaction in computational fluid dynamics. However, their method is limited to perfectly uniform rectangular domains (i.e., no irregular solids) with periodic boundaries. Even so, their potential reconstruction approach requires the solution of a costly vector Poisson problem, making it several times more expensive than pressure projection in a standard fluid solver.

We extend this conceptual framework to be practical for animation applications in two and three dimensions by overcoming several challenges, including enabling much more efficient recovery of discrete vector potentials, enforcing exact no-flow or prescribed

flux boundary conditions for exterior boundaries, and supporting irregular cut-cell obstacles. Our proposed *Curl-Flow* framework offers the following technical contributions:

- An efficient parallel sweeping algorithm to recover velocity-consistent discrete vector potentials on uniform and level set-based cut-cell grids;
- A gauge correction strategy to make the candidate vector potential well-suited to interpolation (i.e., numerically smooth) using a single fast, scalar Poisson solve;
- The design of an efficient, low-order vector potential interpolant that efficiently provides continuous and pointwise incompressible velocities in free space;
- A new additive ramping strategy for improved free-slip boundaries in the "Curl-Noise" method of Bridson et al. [2007];
- The application of this ramping method during velocity interpolation to enforce divergence-free, *exact* no-normal-flow conditions on axis-aligned exterior domain boundaries and interior irregular static obstacles in both 2D and 3D.

2 RELATED WORK

2.1 Within Computer Graphics

2.1.1 Recovering Discrete Vector Potentials. A few vector potential-based Eulerian/hybrid solvers have been proposed in the fluid animation literature, but most operate solely in the discrete realm and require solving a costly *vector* Poisson equation. Elcott et al. [2007] solved a vector Poisson problem to recover the vector potential from vorticity within a simplicial discretization of the vorticity equation. Ando et al. [2015] developed a vector Poisson-based alternative to the standard grid-based pressure projection. Notably, Ando et al. [2015] were the first to suggest employing vector potential-based velocity interpolation, in their discussion of future work. For an ideal Cartesian grid with only axis-aligned boundaries and certain choices of boundary conditions, the vector Poisson system can be decoupled into three scalar Poisson problems for efficiency; Sato et al. (likely) used this approach to recover the vector potential from a velocity field in fluid control applications [Sato et al. 2021, 2015]. Our proposed discrete vector potential reconstruction strategy offers a significantly more efficient solution, finding the desired potential at essentially the cost of a single scalar Poisson solve if the input velocity field is divergence-free, or two otherwise.

For 2D flow visualization, Biswas et al. [2016] recovered the discrete (scalar) stream function from velocity data and used marching squares to construct streamlines. They proposed an axis-based sweeping approach for 2D stream function recovery; we generalize this idea to irregular cut-cell boundaries and 3D vector potentials.

Relatedly, discrete vector field (Hodge) decomposition techniques have long been of interest in graphics [Tong et al. 2003]. Recently, more elaborate five-component discrete decompositions were considered by Poelke and Polthier [2016] and Zhao et al. [2019] for vector fields over triangulated surfaces and tetrahedral volumes, respectively, including various gauges and boundary conditions. These approaches determine the vector potential directly from the vector field via a vector Poisson solve, and in the five-component case, further require solving eigenproblems.

2.1.2 Pointwise Divergence-Free Fields. Various procedural techniques have used pointwise divergence-free velocity fields. Stam and Fiume [1993] first explored this idea for gaseous phenomena, treating incompressibility by restricting cross-spectral density functions in the Fourier domain. Curl-Noise [Bridson et al. 2007] uses a continuous vector potential to design animated divergence-free vector fields and the related scheme of DeWolf [2006] uses the cross-product of two gradient fields. Divergence-free sub-grid turbulence models subsequently built on these ideas [Kim et al. 2008; Schechter and Bridson 2008]. Incidentally, Schechter and Bridson [2008] also used (but did not describe) a divergence-free hybrid constant-linear velocity interpolant (see their supplemental video at time 2:18-2:27, left), which we show can be derived from bilinear stream function interpolation. This is the *only* prior instance of divergence-free grid-based interpolation in computer animation we know of. Pointwise divergence-free vector fields have also been used in geometry processing for volume-preserving modeling [Von Funck et al. 2006] and shape interpolation [Eisenberger et al. 2018].

2.1.3 Divergence-Free Simulators. Lagrangian vorticity-based simulation methods also employ a secondary vector variable to construct analytically divergence-free velocity fields; specifically, velocity is computed from vorticity via the Biot-Savart law. These methods come in many forms, including particles [Park and Kim 2005], filaments [Angelidis and Neyret 2005], and sheets [Brochu et al. 2012; Pfaff et al. 2012]; closely related boundary integral/element (surface-only) methods also offer divergence-free fields [Da et al. 2016]. Model-reduction methods based on Laplacian eigenfunctions offer pointwise divergence-free fields via analytical basis functions [Cui et al. 2021; De Witt et al. 2012], albeit in simple domains.

2.1.4 Subdivision Schemes. Subdivision schemes based on discrete exterior calculus consider how to preserve properties of discrete differential operators on triangulated surface meshes undergoing refinement [De Goes et al. 2016; Wang et al. 2006]. These papers focus on surfaces and do not consider volumetric uniform grids or polyhedral cut-cells. More fundamentally, they achieve pointwise incompressibility only in the limit of infinite refinement. It would be interesting to adapt these ideas to polyhedra.

2.1.5 Enhanced Advection. A wide array of fundamental improvements to advection have been proposed to go beyond basic (trilinear) interpolation and (semi-Lagrangian) advection schemes [Stam 1999], ranging from monotonic cubic interpolation [Fedkiw et al. 2001] to McCormack advection [Selle et al. 2008] to variants of particle-in-cell schemes [Jiang et al. 2015; Zhu and Bridson 2005]. Such methods tend to focus on vorticity-preservation, higher order accuracy, or diffusion-reduction. Because they typically use the underlying interpolant as a black box, they are orthogonal, and often complementary, to our scheme, as we demonstrate for FLIP [Zhu and Bridson 2005].

2.2 Outside Computer Graphics

2.2.1 Divergence-Free Finite Element Methods. In applied math, a wide range of discontinuous Galerkin (DG) and other non-conforming finite element methods (FEM) have been developed to offer pointwise divergence-free fields, often by adopting carefully designed incompressible per-element basis functions (e.g., [Cockburn et al.

2004; Lehrenfeld and Schöberl 2016; Rhebergen and Wells 2018]). In particular, Maljaars et al. [2018] combined an exactly divergence-free Hybridized Discontinuous Galerkin scheme with a particle-in-cell method (in 2D, without obstacles) and demonstrated better particle distributions. However, this family of FEM methods tends to have a more complex implementation and do not naturally integrate with standard approaches in animation. Moreover, they possess inter-element field discontinuities by their very nature, which can be problematic for visual applications; Maljaars et al. [2018] observed that reducing the magnitude of these inter-element discontinuities helps to improve particle distributions. Guzmán and Neilan [2014] tackled the more challenging task of designing a *fully* conforming/continuous Stokes finite element method yielding divergence-free solutions on 3D tetrahedra. Doing so required a specialized finite element basis consisting of a combination of cubic polynomials and divergence-free rational functions, and due to the construction's complexity the authors note that "practical significance of the proposed elements may be questionable". In isogeometric analysis, Evans and Hughes [2013] developed an exactly divergence-free Navier-Stokes simulation framework based on geometrically mapped rectangular B-spline grids; approaches in this vein require complex mesh construction for non-trivial domains, in contrast to the simpler and more efficient cut-cell techniques often preferred in animation. In fact, all of the approaches above would necessitate replacing the entire Navier-Stokes simulator, while ours provides a convenient plug-in upgrade to the advection phase of industry-standard visual effects methods. Thus, somewhat closer to our approach are methods that post-process flow solver solutions to exactly recover a divergence-free velocity field. For example, Linke [2012] converts the solution of a staggered triangulated discretization into a strictly divergence-free field in terms of Raviart-Thomas elements. Similarly, Lederer et al. [2017] proposed a velocity reconstruction operator that maps discretely divergence-free fields of Taylor-Hood and mini elements to exactly divergence-free ones, and suggested using their method as a postprocessing of the discrete solution in the conclusion. However, like the DG methods discussed above, the resulting finite element spaces (i.e., velocity fields) are not continuous between elements.

2.2.2 Direct Interpolation of Finite Volume Solutions. In a computational fluid dynamics context, Jenny and colleagues [Jenny et al. 2001; Meyer and Jenny 2004] derived 2D divergence-free node-based velocity interpolants for uniform grids and demonstrated improved particle distributions in particle-in-cell schemes. In geodynamics, Wang et al. [2015] developed a 3D extension and Pusok et al. [2017] adapted it for face-based data by first averaging to nodes. These approaches augment multilinear interpolants with corrective terms to satisfy pointwise incompressibility per cell. In astrophysics (magnetohydrodynamics), Balsara [2001, 2004] proposed very similar divergence-free vector field reconstruction strategies based on local piecewise fitting on each cell or tetrahedron; one presupposes a polynomial basis and uses the divergence-free condition to determine constraints. Balsara [2009] presented an extension of this approach to Cartesian grid WENO schemes. These approaches do not consider cut-cell geometries, and exhibit kinks between cells.

Conveniently, the approaches above do not require recovering a vector potential at all. However, the recovered fields exhibit pervasive discontinuities between cells, and to achieve a given order of accuracy or support a particular geometry, they must be developed each time on a case-by-case basis. Our choice to instead recover and exploit an explicit discrete vector potential makes it comparatively straightforward to derive generalizations of our framework to higher order accuracy, global continuity/smoothness, or different element shapes, using relatively *standard* interpolation techniques, including spline interpolants or mesh-free methods. It could also be fruitfully combined with fluid control tools that already rely on vector potentials [Sato et al. 2021, 2015].

Concurrent work by Schroeder et al. [2022] proposes a direct spline-based interpolation method for staggered regular grid velocity fields that yields pointwise divergence-free fields of different orders. Some variants of their method can also exactly interpolate the cell normal components. This method avoids recovering the vector potential, but is limited to uniform grids without obstacles.

An alternative meshless divergence-free interpolation strategy [Lowitzsch 2005; McNally 2011] exploits matrix-valued radial basis functions and properties of the vector Laplacian. While this method offers divergence-free fields, as a meshless point-based method it struggles to precisely enforce boundaries and, more critically, preliminary tests showed it to be massively expensive, since it uses a fully dense global solve to simultaneously determine coefficients for all of the radial basis functions. Even attempting to substitute local kernels for the method's global kernels remains expensive: a large support radius (i.e., particle neighborhood) is needed to avoid local oscillations and thus the system matrix is still unattractively dense.

2.2.3 Vector Potentials and Gauge Conditions. The methods most similar to ours are those that use explicit vector potentials. Since the vector potential for a given velocity field is not unique, various gauge conditions are used to select a particular potential from this null space. The gauge choice provides one dimension along which to categorize the schemes discussed below. Another is whether they apply a local approach to find the vector potential, based on a cell-by-cell (possibly branching) traversal, or find a simultaneous global solution by solving a linear system. We straddle these categories: our parallel sweeping approach is a fast cell-by-cell scheme that yields an initial potential, which we correct with a relatively inexpensive scalar Poisson solve (via Discrete Sine Transform or DST) to ensure smoothness by enforcing the Coulomb gauge, $\nabla \cdot \psi = 0$. Furthermore, beyond our innovations in the uniform grid setting, our work is unique in considering solid geometry via cut-cells.

The Coulomb gauge condition picks out a unique smooth vector potential by enforcing that it be divergence-free; it is the most common choice and has been used in graphics (e.g., [Ando et al. 2015] among others.) The initial implicit gauge condition used in our 3D parallel sweeping method is conceptually similar to that of Ravu et al. [2016], who sought to directly construct spline-based vector potentials. Like us, they set the ψ_z component to zero. However, their overall approach differs in that (a) their velocity and vector potential degrees of freedom are colocated at nodes, (b) they directly seek the coefficients of continuous cubic spline functions rather than edge-based discrete vector potential values (hence four times as many

unknowns), and (c) compared to our fast parallel sweeping, they solve an expensive global linear system to find the many coefficients.

Another implicit gauge condition falls out from the tree-cotree approach, explored for edge-based finite element methods by Albanese and Rubinacci [1990] on grids and Manges and Cendes [1995] on meshes. This global approach determines a spanning tree of an edge-based grid or mesh that one can safely set to zero, in order to explicitly eliminate the null space of the matrix representing the discrete curl operator. This idea has natural connections to the traversal patterns of cell-by-cell approaches.

Bao et al. [2017] proposed a divergence-free interpolation strategy on uniform staggered grids that reconstructs a discrete vector potential field from velocities, to improve coupling and reduce volume conservation errors in the immersed boundary (IB) method [Peskin 2002]. (IB uses smeared delta functions for approximate coupling with solids, in contrast to our "sharp" cut-cells.) Similar to methods in graphics [Ando et al. 2015; Tong et al. 2003], they recover the vector potential under the Coulomb gauge using a vector Poisson system. They also assume periodic boundary conditions, which are undesirable for most graphics applications, and only suggest non-periodic domains as future work. Their periodic, obstacle-free setting decouples the vector problem into three scalar Poisson solves, which they treat via the fast Fourier transform (FFT). Our method recovers the same potential with an efficient sweep process followed by just one scalar Poisson solve (which can likewise be done via FFT on simple domains). We further support interior obstacles and closed domains with a fast DST-based solve. For interpolation, they use tensor-product B-splines, similar to ours. Casquero et al. [2018] also considered the immersed boundary method on uniform grids, but like Evans and Hughes [2013] and in contrast to Bao et al. [2017], they perform the entire simulation using divergence-free spline basis functions rather, than post-processing staggered data.

Silberman et al. [2019] adopt a staggered uniform grid configuration and seek a discrete edge-based vector potential, similar to our setting, but for applications in electromagnetics. They propose a sequential, cell-by-cell flooding strategy which is more costly than our approach. Specifically, our axis-based sweeping safely assumes zero values in the z -component, which reduces the computation per cell, the number of unknowns (by a third), and the number of special cases, while enabling a fast parallel solution via 1D sweeps. Like us, they use a Poisson-based correction to apply the Coulomb gauge, but like Bao et al. [2017] they support only uniform grids, and they use infinite boundary conditions to allow solution by FFT. Silberman et al. [2019] additionally suggest an alternative "global linear algebra" approach that is loosely similar to the vector Poisson method of Ando et al. [2015] in that it both removes divergence and recovers a potential satisfying the Coulomb gauge; they note that it scales worse than their flood-then-correct approach.

In summary, our vector potential recovery approach is most closely related to the uniform staggered grid methods of Bao et al. [2017] and Silberman et al. [2019], but is more efficient and generalizes these ideas to the more flexible domain boundary conditions and irregular cut-cell solids favored in graphics applications.

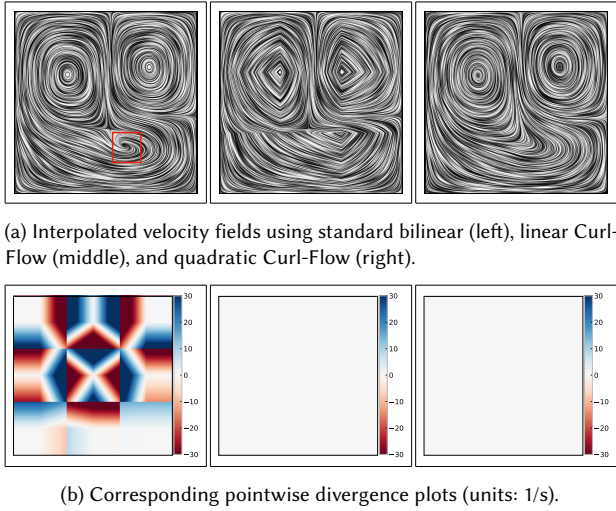


Fig. 2. **Pointwise Divergence Comparison:** Left: Bilinear interpolation of discretely divergence free velocity data on a 3×3 grid. The red box in (a) contains a spurious sink. Middle: Linear Curl-Flow interpolation is divergence-free, but has piecewise constant components that induce kinks. Right: Quadratic Curl-Flow interpolation is divergence-free and smooth.

3 PROBLEM SETTING

We build on a standard staggered grid-based fluid solver for the incompressible Euler equations [Bridson 2015; Fedkiw et al. 2001]. We represent solids using node-based level sets, leading to simple marching cubes cut-cells, and assume that all solids are static. For cut-cell pressure projection, we use the finite volume Poisson stencil of Ng et al. [2009]. For velocity advection, we use semi-Lagrangian backtracing with multilinear interpolation and third-order Runge-Kutta, except where noted.

We develop a new velocity interpolant for use during advection,

$$\frac{\partial q}{\partial t} + \mathbf{u} \cdot \nabla q = 0, \quad (3)$$

where \mathbf{u} is an incompressible velocity field and q is any scalar (e.g., density) or vector (e.g., velocity) quantity to be advected. Common advection discretizations use semi-Lagrangian [Stam 1999] or Lagrangian [Zhu and Bridson 2005] methods that trace trajectories through the flow and require interpolation to query the velocity.

The input discrete velocity field \mathbf{u} to be used for advection is assumed to be discretely divergence-free due to the underlying simulator's incompressibility enforcement (i.e., pressure projection). Discrete velocity components are placed at grid face centers, so the discrete divergence-free condition is

$$\sum_{\text{faces } f} A_f u_f = 0, \quad (4)$$

where A_f is face area and u_f is the face's outward oriented normal velocity. Given such discrete data \mathbf{u} , our task is to generate a pointwise vector field that is analytically divergence-free, while ideally being continuous, smooth, and boundary-respecting.

Bilinear interpolation applied to each staggered velocity component independently is the simplest common interpolant, but the

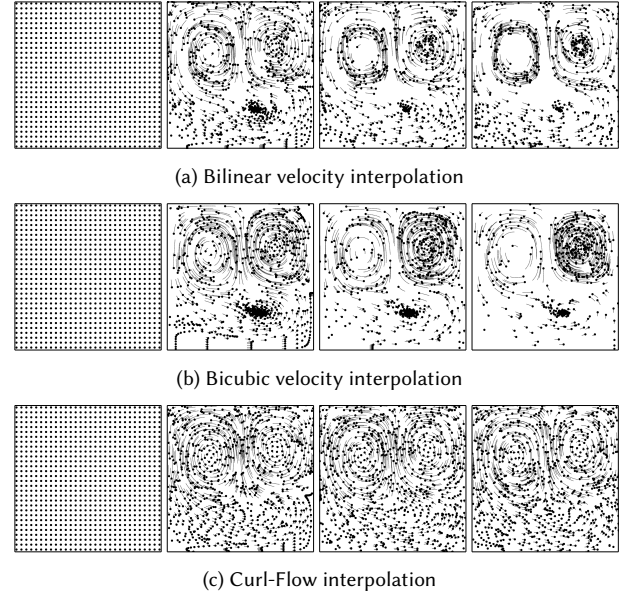


Fig. 3. **Effect on Particle Distribution:** Initially uniform particles advected through a static 2D vector field on a 3×3 grid. Frames 1, 100, 200, and 300 are shown left to right. (a): With direct bilinear velocity interpolation, particles become clustered, leaving large empty voids. (b): With higher order direct velocity interpolation (monotonic cubic [Fedkiw et al. 2001; Fritsch and Carlson 1980]), clustering and spreading remain significant. (c): With quadratic Curl-Flow interpolation, particles remain better distributed.

resulting vector field's analytical divergence is nonzero in general (see Section 1 in supplemental material). Figure 2, left column, visualizes such a flow and its divergence; if one advects uniformly sampled particles through this field (Figure 3, top row), a sink region absorbs a large number of particles while many more cluster into large and small rings. Unfortunately, neither higher order interpolation (Figure 3, middle row) nor using node-based / collocated velocity data will resolve these issues, since those interpolants remain divergence-oblivious. We use staggered multilinear interpolation as our baseline *direct velocity interpolant* throughout the paper, unless stated otherwise.

The middle and right columns in Figure 2 visualize the linear and quadratic variants of our proposed approach, respectively, developed in the next section. Since the linear case exhibits velocity kinks we ultimately prefer quadratic, but neither flow contains sinks or sources and their pointwise divergence fields are exactly zero.

Section 2.2.2 mentioned prior approaches that aim for divergence-free interpolation using velocity data directly, rather than needing intermediate potentials. In particular, Jenny et al. [2001] and Balsara [2001] used the same staggered grid layout as us in 2D (and in practice yield identical fields to one another). These *direct incompressible velocity interpolants* produce pointwise divergence-free velocities within each cell, but kinks are visible *between* most cells (see Figure 4) due to their piecewise construction and use of the minmod limiter [Jenny et al. 2001], unlike our (quadratic) Curl-Flow method.




Fig. 4. **Direct Incompressible Velocity Interpolation vs. Curl-Flow.** (a) Discrete and discretely divergence-free velocity field to interpolate. (b) Direct incompressible velocity interpolation [Balsara 2001; Jenny et al. 2001]. The velocity field inside each cell is incompressible, but discontinuities between cells are pronounced. (c) Curl-Flow interpolation is pointwise incompressible and smooth both inside and across cells.




Fig. 5. **Discretization in 2D:** Red squares represent discrete velocity samples and blue circles represent discrete stream function samples. (a) A uniform grid cell. (b) A cut-cell induced by clipping with a solid object. (c) The relationship between edge normal velocity and stream function samples. For this case, $v_n = -(\psi_j - \psi_i)/l$.

4 CURL-FLOW INTERPOLATION IN 2D

We first describe our approach in 2D without obstacles. Velocity \mathbf{u} has two components, u and v , and the vector potential has one scalar component, ψ_z , called the *stream function* and denoted by non-bolded ψ . The relationship (1) simplifies to

$$u = \frac{\partial \psi}{\partial y}, \quad v = -\frac{\partial \psi}{\partial x}. \quad (5)$$

That is, \mathbf{u} is a 90° rotation of $\nabla \psi$, denoted $\nabla \psi^\perp$. To discretize, we place ψ samples at cell vertices (Figure 5) and assume ψ varies linearly on edges. Figure 5(c) illustrates two nodal ψ values and the normal component of velocity, v_n , on a shared edge e_{ij} having unit tangent \mathbf{e}_{ij} and length l . The gradient theorem,

$$\int_{e_{ij}} \nabla \psi(\mathbf{r}) \cdot d\mathbf{r} = \psi(\mathbf{x}_j) - \psi(\mathbf{x}_i), \quad (6)$$

discretized on each edge, gives the relationship

$$\frac{\psi_j - \psi_i}{l} = \nabla \psi \cdot \mathbf{e}_{ij} = \nabla \psi^\perp \cdot \mathbf{e}_{ij}^\perp = v_n(\mathbf{n} \cdot \mathbf{e}_{ij}^\perp). \quad (7)$$

Since \mathbf{e}_{ij}^\perp is oriented, unit length, and matches \mathbf{n} up to a sign flip, the dot product simply determines the sign.

4.1 Uniform Grids in 2D





Fig. 6. **Parallel Sweeping** efficiently recovers a uniform grid discrete stream function field. Left: First, sequentially compute all ψ values at $x = x_0$, starting from $\psi_0 = 0$ and sweeping in the y direction. Right: Next, sweep in the x direction, processing rows in parallel.

4.1.1 *Recovering discrete ψ .* For a single uniform grid cell as in Figure 5(a), equation 7 leads to

$$\begin{aligned} u_l &= \frac{\psi_2 - \psi_0}{h}, & u_r &= \frac{\psi_3 - \psi_1}{h}, \\ v_b &= -\frac{\psi_1 - \psi_0}{h}, & v_t &= -\frac{\psi_3 - \psi_2}{h}, \end{aligned} \quad (8)$$

consistent with finite differences on (5). We seek ψ_i satisfying the given velocities. These equations have a 1D null space: constant offsets of all ψ_i do not change the velocities. (Physically, discrete incompressibility ensures one flux is the negated sum of the others, implying one equation is linearly dependent.) We select a unique solution by arbitrarily setting $\psi_0 = 0$. We determine the other ψ_i by traversing the cell's edge graph, computing each subsequent ψ value from its predecessor on the edge and the edge's velocity component. Discrete incompressibility ensures discrete integrability, i.e., looping back to ψ_0 gives a consistent result.

This graph traversal scheme applies likewise to a whole uniform grid of cells, thereby avoiding a costly global Poisson solve for ψ based on (2). Any ordering suffices, so we aim to maximize parallelism. As proposed by Biswas et al. [2016] and illustrated in Figure 6, we first set $\psi_0 = 0$, and sweep vertically to obtain all the ψ values at $x = x_0$ (i.e., $\psi_1 = \psi_0 + hu_0$ and so on). We then sweep horizontally, in parallel, to obtain all remaining ψ values. This approach is compatible with any standard exterior domain boundary conditions (inflow/outflow/open/closed), since it only needs the velocity data.



4.1.2 *Interpolating ψ .* Interpolation of the grid ψ values provides an analytical stream function at every point. Applying the analytical curl, we obtain a velocity field which is *pointwise* divergence-free by construction. As a simple example, for bilinear interpolation (referring to the inset figure),

$$\begin{aligned} \psi(x, y) &= \text{lerp}(\text{lerp}(\psi_0, \psi_2, \beta(y)), \text{lerp}(\psi_1, \psi_3, \beta(y)), \alpha(x)), \\ u(x, y) &= \frac{\text{lerp}(\psi_2 - \psi_0, \psi_3 - \psi_1, \alpha(x))}{h} = \text{lerp}(u_l, u_r, \alpha(x)), \\ v(x, y) &= -\frac{\text{lerp}(\psi_1 - \psi_0, \psi_3 - \psi_2, \beta(y))}{h} = \text{lerp}(v_b, v_t, \beta(y)), \end{aligned} \quad (9)$$

where $\text{lerp}(a, b, t) = (1-t)a + tb$ and the rightmost equalities follow from (8). In this specific case, the analytical divergence is clearly

zero because it exactly equals the finite difference divergence:

$$\nabla \cdot \mathbf{u} = \frac{u_r - u_l + v_t - v_b}{h} = 0. \quad (10)$$

The derivatives in the curl operator inherently induce a different polynomial degree per axis: here, velocity components are piecewise constant in one axis and piecewise linear in the other, causing undesired tangential discontinuities between cells (Figure 2, middle).

Fortunately, because we recovered the discrete ψ , we can achieve velocity continuity simply by upgrading to a higher order interpolant (Figure 2, right). Although many choices are possible, we adopt quadratic dyadic B-spline kernels [Jiang et al. 2016; Steffen et al. 2008]; they are sufficiently smooth, their analytical curl is straightforward to derive, and they possess small stencils for efficiency. Bao et al. [2017] similarly used tensor-product B-splines when constructing divergence-aware Dirac delta functions.

For the special case of bilinear ψ interpolation on 2D uniform grids, (9) shows that the velocity interpolant reduces to a simple function of the input discrete u, v velocities, eliminating the need for explicit ψ values. This effect occurs because, in 2D, each velocity component is defined by a single partial derivative of ψ and the null space is a constant shift. However, deriving simple direct incompressible velocity interpolants in this way becomes unwieldy or impossible in the combined presence of cut-cell geometry, higher order interpolation, and especially 3D, where each velocity component depends on two vector potential components and the null space is nontrivial. We therefore prefer to explicitly recover ψ .


4.2 Cut-Cell Obstacles in 2D

We next extend our scheme to support irregular solid obstacles.

4.2.1 Recovering discrete ψ . Considering cut-cell solids during stream function recovery only requires accounting for truncated grid edges and non-axis-aligned "cut edges". Applying (7) we get an update rule for sweeping through such edges:

$$\psi_{i+1} = \psi_i + I_{i,i+1} v_{i,i+1} (\mathbf{e}_{i,i+1}^\perp \cdot \mathbf{n}_{i,i+1}). \quad (11)$$

For the static obstacles we consider, the recovered ψ will be constant on the surface, i.e., the surface is an isocontour that particles slide along in a free-slip manner [Bridson et al. 2007]. Static obstacles also preserve the remarkable simplicity of our parallel sweeping method: nodal stream function values at the entry and exit point of a grid line passing through an obstacle are identical, so we can sweep *through* obstacles by assuming zero velocities inside.



4.2.2 Interpolating ψ . Cut-cell solids cause the nodal ψ values (blue in the inset) to no longer lie in a convenient uniform grid, thereby complicating interpolation. One could use moving least squares (MLS) interpolation based on ψ data in some neighborhood, but MLS requires solving small dense linear systems and computing the necessary analytical

derivatives is difficult [Huerta et al. 2004]. Instead we extrapolate ψ values to uniform grid nodes inside the solid (i.e., ψ_s in the inset) and use quadratic B-spline kernels as in the uniform grid case. This extrapolation can be done in various ways, such as using MLS or

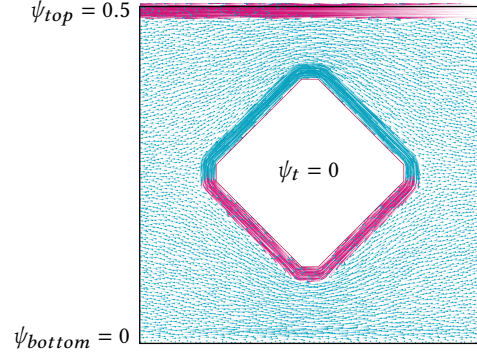


Fig. 7. **Boundary Ramping Issues:** Multiplicative Curl-Noise ramping with a target value of zero ($\psi_g = 0$) applied to a steady rightward flow, where the constant inflow/outflow rate is described by a vertical linear stream function value from $\psi_{top} = 0.5$ to $\psi_{bottom} = 0$. Blue trails indicate the right direction (+x) and red trails indicate (erroneous) trajectories in the reverse (-x) direction. Because the expected discrete ψ_g value should differ at each surface, always ramping to $\psi_g = 0$ causes spurious tangential flow near the obstacle, undesired no-slip at the bottom where ψ is already zero, and large reverse flow at the top.

simply copying (e.g., from ψ_g to ψ_s), assuming discrete flux is zero on interior solid edges (dashed lines); we adopt the latter for simplicity. At the outer axis-aligned domain boundaries, the quadratic stencil is also missing data samples, so we extrapolate an extra uniform layer of ψ samples outside the domain. However, regardless of these various extrapolation and interpolation choices near boundaries, the desired pointwise no-normal-flow boundary condition is (so far) enforced only approximately; we propose a final correction below.

4.2.3 A Curl-Noise enhancement for exact 2D boundary enforcement. To ensure that the interpolated ψ will be *strictly* constant along the polygonal boundary for static obstacles, we adapt ideas from the Curl-Noise method [Bridson et al. 2007]. That method modulates the existing ψ value near obstacles, forcing it to a constant zero value on the boundary by multiplying against a smooth ramp function based on boundary proximity. Defining d_0 as an influence radius (we use the grid width h), $d(\cdot)$ as the distance to the surface, and $ramp(\cdot)$ as a smooth ramp function (we use a classic smoothstep function), Bridson's multiplicative correction to the ambient ψ is the product $\psi'(\mathbf{x}) = \alpha\psi(\mathbf{x})$, where $\alpha = ramp(d(\mathbf{x})/d_0)$.

However, depending on the background ψ , the appropriate target boundary value ψ_g is often *not* zero; arbitrarily choosing zero leads to dramatic spurious flow deviations (Figure 7). One can modify Bridson's method to "ramp in" the nonzero target boundary value ψ_g :

$$\psi'(\mathbf{x}) = \alpha\psi(\mathbf{x}) + (1 - \alpha)\psi_g. \quad (12)$$

Unfortunately, the initial multiplication still damages the *normal derivatives* of ψ . Since $\mathbf{u} = \nabla\psi^\perp$ implies $(\mathbf{u} \cdot \mathbf{t}) = \frac{\partial\psi}{\partial n}$, where \mathbf{t} is the boundary tangent vector, (12) induces undesired tangential damping, acceleration, or no-slip behavior (Figure 8, bottom half).

We propose a novel *additive* ramping procedure that instead enforces ψ_g by adding a compensating offset. The required offset is determined by evaluating the existing ψ at the closest boundary

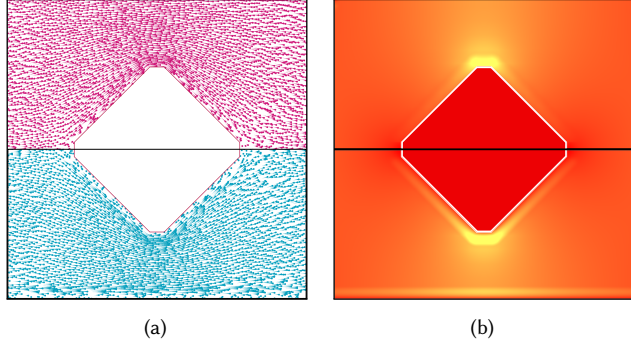


Fig. 8. **Additive vs. Multiplicative Ramping:** A steady rightward flow with boundary ψ_g value adjustment. Left: Particles with speed-dependent trail lengths. Right: Velocity magnitudes. Top half: Our additive ramping exactly satisfies the boundary with no apparent damping of free-slip velocities. Bottom half: Despite modifying Bridson’s multiplicative ramp to target the correct ψ_g value, tangential damping occurs along the bottom wall and near the solid, since $\partial\psi/\partial\mathbf{n}$ is damaged. Note the more pronounced banding in the bottom magnitude plot.

point and subtracting it from ψ_g . Letting $cp(\cdot)$ be a function returning the closest boundary point, we modify ψ as:

$$\psi'(\mathbf{x}) = \psi(\mathbf{x}) + (\psi_g - \psi(cp(\mathbf{x}))) (1 - \alpha). \quad (13)$$

This expression ramps the full correction precisely on at the boundary ($d(\mathbf{x})/d_0 = 0$) and blends it smoothly off at the edge of the influence region ($d(\mathbf{x})/d_0 = 1$). The inset shows an example input ψ curve (black), and the result after using multiplicative (red) and our additive (green) ramping to $\psi_g = 0$.

Both methods correct the value, but additive ramping better preserves the derivative near the boundary and thus yields more faithful free-slip flow (Figure 8, top half). We again use the analytical curl to determine this added velocity contribution.




Fig. 9. **Flow Near Solids:** Particle trajectories flowing from bottom to top under a fixed (steady state) flow past a solid disk in 2D. Grid resolution: 30×15 . (a) Direct velocity interpolation induces spurious trailing gaps, when trajectories colliding with the solid are terminated. (b) With the same velocity interpolant, projecting penetrating particles back to the surface does not significantly help divergent trajectories to “close up” and unwanted particle clumping occurs. (c) Our Curl-Flow method better respects solid boundaries and produces less spurious empty space behind the solid. Notice that the trajectories are simply isocontours of the 2D ψ field.

By its nature, applying ramping on piecewise flat (polygonal or polyhedral) geometry can introduce some discontinuity in velocity near sharp edges or vertices. This is because the recovered velocity $\mathbf{u}' = \nabla\psi'^{\perp}$ involves a $\nabla\alpha$ term from (13) (i.e., gradient of the normalized distance field), and $\nabla\alpha$ is not continuous in general. Nevertheless, our Curl-Flow method produces precisely boundary-respecting and divergence-free flows, and the discontinuities are often acceptable in practice, since the applied corrections only occur within one cell-width of obstacles.

In the context of our overall interpolation scheme, the ambient ψ is the interpolated grid-based ψ field and the target boundary ψ_g values are known at solid surface nodes, since they are recovered during sweeping. We apply our ramping strategy to closed exterior domain boundaries and static solid surfaces. In those cases, ψ_g is a constant value, since the boundary’s normal flux is zero. For exterior

domain boundaries that have nonzero inflow/outflow, the discrete ψ edge endpoint values of a segment differ (e.g., ψ_{c_0}, ψ_{c_1}). One can set the ramping target value ψ_g to the linearly interpolated value of ψ_{c_0} and ψ_{c_1} at the closest point to achieve the desired constant flow in the normal direction across that edge. However, for simplicity, we did not apply our ramping-based enhancement to the exterior domain boundaries with inflow/outflow conditions, since exact enforcement there is not visually critical.

Figure 9 compares direct (bilinear) velocity interpolation vs. our full Curl-Flow method (including additive ramping boundary correction) on a simulated steady-state flow past a disk. To approximate free-slip flow for the comparison direct method, we extrapolate fluid velocities (with solid-normal components projected out) to samples inside the solid [Houston et al. 2003; Rasmussen et al. 2004].

5 CURL-FLOW INTERPOLATION IN 3D

In three dimensions, the scalar stream function ψ is replaced by the vector potential $\boldsymbol{\psi} = (\psi_x, \psi_y, \psi_z)$. Each velocity component is now dictated by the *interactions* of two ψ components’ derivatives (e.g., $u = \partial\psi_z/\partial y - \partial\psi_y/\partial z$), which complicates boundary enforcement, especially for irregular geometry. The curl operator also possesses a multi-dimensional null space, necessitating enforcement of a *gauge condition* to find an appropriate unique $\boldsymbol{\psi}$.

5.1 Uniform Grids in 3D

We place vector potential components on cell edges and velocity normal components on cell faces [Ando et al. 2015; Bao et al. 2017; Elcott et al. 2007] (Figure 10). Given discretely divergence-free face velocities, we seek corresponding edge vector potential values.

For simplicity, we assume for now that the input axis-aligned exterior domain boundaries have zero normal velocity. We can safely




Fig. 10. **Discretization in 3D:** Discrete vector potential samples are located on cell edges (circles) and velocity samples are on cell faces (squares). Red, green, and blue represent x , y , and z components, respectively.

achieve this with a constant value (i.e., 0) for the tangential components of the boundary ψ (denoted ψ_{tan}). (Our method extends straightforwardly to prescribed inflow/outflow velocities, with appropriate variations in ψ_{tan} .)

With this choice, the continuous problem we are solving is

$$\begin{aligned} \nabla \times \psi &= \mathbf{u} & \text{in } \Omega, \\ \psi_{tan} &= 0 & \text{on } \partial\Omega. \end{aligned} \quad (14)$$

The flux across a given surface, in terms of the vector potential, is

$$\iint_S \mathbf{u} \cdot d\mathbf{S} = \iint_S (\nabla \times \psi) \cdot d\mathbf{S} = \int_C \psi \cdot d\mathbf{r}, \quad (15)$$

where S is an oriented smooth surface, C is the oriented boundary curve of the surface, and \mathbf{r} is the boundary tangent direction. (The second equality holds by Stokes' theorem.) Discretizing (15) with midpoint quadrature for the single face in Figure 10, right, we have

$$u_f h^2 = \sum_{e \in E} \psi_e \cdot (h\mathbf{e}_e) = h(\psi_{y_0} + \psi_{z_1} - \psi_{y_1} - \psi_{z_0}), \quad (16)$$

where h is a cell width and \mathbf{e}_e is the oriented unit vector along an edge. (Finite differences on (14) yields the same.) Equation 16 relates four unknown edge ψ_e values to one face u_f velocity. Stacking the equations for all grid faces yields a global sparse linear system.

Equation 14 is an inverse curl problem with infinitely many solutions, e.g., a single cell has a seven-dimensional null space: 12 edge ψ_e and five linearly independent face u_f (the sixth is redundant by incompressibility). We select a unique solution using the popular *Coulomb* gauge condition (with our boundary condition), which enforces $\nabla \cdot \psi = 0$. This gauge offers maximal smoothness of the ψ field, which will be attractive for interpolation.

Directly taking the curl of (14) and assuming $\nabla \cdot \psi = 0$ yields

$$\nabla \times (\nabla \times \psi) = -\nabla^2 \psi = \nabla \times \mathbf{u}, \quad (17)$$

i.e., a *vector* Poisson problem for ψ , around three times larger than the pressure projection (e.g., [Ando et al. 2015]). Fortunately, unlike Ando et al., our input discrete velocities are *already* incompressible, enabling us to solve (14) more efficiently. Observe that

$$\mathbf{u} = \nabla \times \psi = \nabla \times (\psi + \nabla\phi) = \nabla \times \psi', \quad (18)$$

where ϕ is an arbitrary scalar field. Since $\nabla \times \nabla\phi$ is always zero (a vector calculus identity), defining $\psi' = \psi + \nabla\phi$ gives another

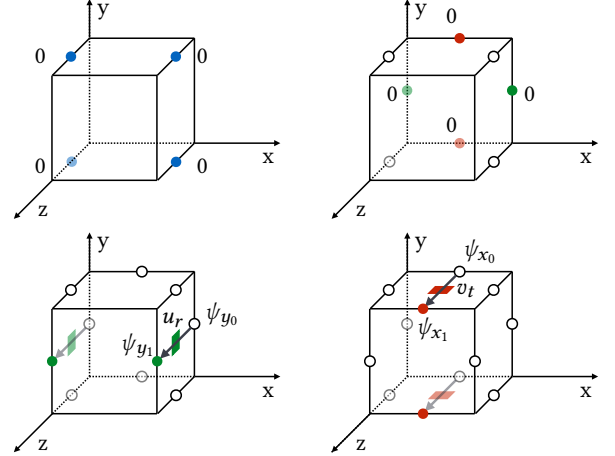


Fig. 11. **Parallel Sweeping in 3D:** Top-left: Set all ψ_z values to be zero. Top-right: Compute satisfying ψ_x and ψ_y values at $z = z_{min}$. In this example they are all set to zero. Bottom-left: Starting from ψ_y values at $z = z_{min}$, we can compute all ψ_y values in the entire domain in one sweep using the relation $\psi_{y_1} = \psi_{y_0} - hu_r$. Bottom-right: Starting from ψ_x values at $z = z_{min}$, we can compute all ψ_x values in the entire domain in one sweep using the relation $\psi_{x_1} = \psi_{x_0} + hv_t$.

valid vector potential field for the same velocity. Leveraging this characterization of the null space, we propose a two step approach. First, use an efficient 3D extension of our 2D parallel sweeping scheme (Section 5.1.1) to find a ψ field that is *velocity-consistent but has arbitrary gauge*. Second, modify this ψ field to satisfy the Coulomb gauge $\nabla \cdot \psi = 0$ and our boundary condition $\psi_{tan} = 0$ using a carefully constructed scalar potential ϕ (Section 5.1.2).

5.1.1 Recovering a Vector Potential by Parallel Sweeping. Our first task is to efficiently find a velocity-consistent discrete vector potential field on a box-shaped domain, irrespective of boundary conditions or gauge choice. Our proposed fast 3D parallel sweeping strategy is illustrated in Figure 11.

- (1) Set $\psi_z = 0$ (or a constant) *everywhere* in the domain (Figure 11, top-left). This is safe because the remaining two components (ψ_x, ψ_y) still suffice to represent any set of three velocity components (u, v, w) [Ravu et al. 2016].
- (2) Compute velocity-satisfying vector potential values for the $z = z_{min}$ boundary plane. For zero boundary normal fluxes, we can simply set all ψ_x and ψ_y to zero (Figure 11, top-right).
- (3) Compute the remaining vector potential values by parallel sweeping. Equation 16 and $\psi_z = 0$ give $\psi_{y_1} = \psi_{y_0} - hu$ (Figure 11, bottom-left), and likewise $\psi_{x_1} = \psi_{x_0} + hv$ (Figure 11, bottom-right). In like manner, values of ψ_{x_n} and ψ_{y_n} , along with the discrete velocities, dictate $\psi_{x_{n+1}}$ and $\psi_{y_{n+1}}$ as we sweep across the entire domain.

When this process ends, the velocity condition (16) is met on all grid faces, $\psi_z = 0$ everywhere, and we have zero ψ_x and ψ_y values on all the outer boundaries except for $z = z_{max}$. (For nonzero boundary fluxes, more ψ_x and ψ_y values will be nonzero.) We must next modify this ψ to enforce the desired gauge and boundary conditions.




Fig. 12. **Gauge Correction Boundary Conditions:** We construct $\nabla\phi_{BC}$ to update boundary ψ components to satisfy no-normal-flow. Black/white circles, black/white diamonds, and purple diamonds represent known vector potential values, pinned ϕ values, and unknown ϕ values, respectively. Unknown ϕ values are computed from the pinned or previously computed ϕ values ($\phi_1 = \phi_0 - h\psi_{x_0}$, $\phi_3 = \phi_2 - h\psi_{x_1}$).




Fig. 13. **Gauge Correction:** A uniform grid *scalar* Poisson solve is used to find a nodal scalar field ϕ (purple diamonds). Adding $\nabla\phi$ to the vector potential values, ψ_x (red disks), ψ_y (green disks), and ψ_z (blue disks), satisfies the Coulomb gauge condition, $\nabla \cdot \psi' = \nabla \cdot (\psi + \nabla\phi) = 0$.

5.1.2 Boundary Conditions and Gauge Correction. To satisfy the boundary condition $\psi_{tan} = \mathbf{0}$, we first construct a discrete scalar field ϕ_{BC} , which is defined at nodes of cells so components of the discrete $\nabla\phi_{BC}$ collocate with edge-based vector potential components. The desired boundary condition $\psi'_{tan} = \mathbf{0}$ gives

$$\psi'_{tan} = \psi_{tan} + \nabla\phi_{BC} = \mathbf{0} \text{ on } \partial\Omega. \quad (19)$$

Given ψ_{tan} , already known from parallel sweeping, we must find ϕ_{BC} on boundary nodes. Furthermore, since our sweeping process ensured that the only nonzero boundary ψ_{tan} values left to be eliminated are ψ_x and ψ_y on the $z = z_{max}$ plane, finding nonzero ϕ_{BC} values only on $z = z_{max}$ suffices. We set the outer boundary loop of nodal ϕ_{BC} values on this plane (black/white diamonds in Figure 12) to zero, and compute the interior ϕ_{BC} values in the plane by traversing the interior edges using $\phi_{next} = \phi_{prev} - h\psi_e$.

Adding $\nabla\phi_{BC}$ directly to ψ (including any interior edges touching the nonzero ϕ values) would satisfy our exterior boundary condition

without breaking consistency between the discrete \mathbf{u} and ψ . However, the (arbitrary) gauge of the resulting interior field still offers no guarantees on the presence or absence of large, discontinuous jumps [Silberman et al. 2019]. Thus, while applying interpolation and the analytical curl will yield *some* pointwise divergence-free velocity field, its observed pointwise behavior can be quite irregular depending on interactions between the discrete data jumps and componentwise interpolants, because velocity is determined by differences of vector potential component derivatives. Fortunately, we can enforce the Coulomb gauge on the interior to achieve an optimally smooth vector potential field for interpolation.

Defining a new global nodal ϕ field, we apply a gauge correction similar to Silberman et al. [2019] (although they consider a Fourier-based solution with exterior boundary conditions at infinity and do not handle interior solids). The Coulomb condition says that $\nabla \cdot \psi' = \nabla \cdot (\psi + \nabla\phi) = 0$, giving a node-based *scalar* Poisson problem for ϕ with Dirichlet boundary conditions:

$$\begin{aligned} \nabla \cdot \nabla\phi &= -\nabla \cdot \psi & \text{in } \Omega, \\ \phi &= \phi_{BC} & \text{on } \partial\Omega. \end{aligned} \quad (20)$$

For the box-shaped domains we consider, this Dirichlet problem is amenable to solution with a classic *fast Poisson solver* [Van Loan 1992] based on the Discrete Sine Transform (DST), even after we incorporate irregular obstacles (Section 5.2). After solving for ϕ we update the vector potential as $\psi' = \psi + \nabla\phi$.


5.1.3 Inflow/outflow exterior boundaries. The extension to exterior domain boundaries with nonzero fluxes (inflow and outflow) can be carried out similarly to the 2D case. For example, to match prescribed u velocities on the x boundary, we can use linear ψ_z values in the y direction on the boundary, while keeping ψ_y values at zero. The sweeping and gauge correction steps are easily modified to preserve these boundary values throughout.

5.1.4 Interpolation. Similar to 2D, we handle 3D uniform grid interpolation using low-order dyadic spline kernels, separately on each staggered component of ψ , but choose their polynomial degrees to ensure continuous velocities as follows. The curl operator applied to one component of ψ (e.g., ψ_x) involves its partial derivatives in the other two axes (e.g., $\partial\psi_x/\partial y$ and $\partial\psi_x/\partial z$). To ensure the resulting velocity is at least (piecewise) linear (and continuous) in all directions, we use a mix of linear and quadratic kernels (e.g., for ψ_x , linear x , quadratic y and z). Uniformly quadratic or higher order interpolants would also suffice, in exchange for higher cost.


5.2 Cut-Cells in 3D

The interference of cut-cell solids requires adaptations to our parallel sweeping and gauge correction steps, and a modified 3D ramping strategy to exactly enforce the desired boundary behavior.

5.2.1 Parallel Sweeping with 3D Cut-Cells. By carefully removing redundant DOFs and choosing traversal orders in Sec. 5.1.1, we efficiently obtained velocity-consistent ψ values on uniform grids. With cut-cells present, the number and orientation of vector potential components on some faces change, and thus our sweeping approach must be adapted. Conceptually, we achieve this by converting the



(a) Purple triangle (left) is collapsed to the purple point (right) and the new ψ' values in regular cells are length fractions times the original ψ values ($\psi'_{x_1} = \frac{l_0}{h}\psi_{x_1}$, $\psi'_{z_1} = \frac{l_1}{h}\psi_{z_1}$).



(b) Two purple edges (left) are collapsed (right) and the ψ values are rescaled accordingly ($\psi'_{z_0} = \frac{l_0}{h}\psi_{z_0}$, $\psi'_{z_1} = \frac{l_1}{h}\psi_{z_1}$, $\psi'_{x_1} = \frac{l_2}{h}\psi_{x_1}$).

Fig. 14. **Adapting Parallel Sweeping to 3D Cut-Cells:** Fluid cut-edges are converted to uniform edges by assigning zero ψ to solid purple edges and (conceptually) collapsing them, enabling uniform grid sweeping to proceed.

irregular cut-cell geometry to a nearby voxelized state (Figure 14), applying uniform grid parallel sweeping, and transforming back.

To achieve this mathematically, we rely on two assumptions. First, our earlier assumption of static obstacles implies cut-faces interior to the solid have zero normal flux (purple triangle in Figure 14(a)). Second, we assume that solid portions of partially cut grid-edges (e.g., the two axis-aligned purple segments in Figure 14(a)) have zero vector potential (the fluid parts of the edge are unchanged). Then, applying (15) to the solid triangle implies zero vector potential on the sloped purple edge too. Since the solid edges contribute nothing, we collapse them and rescale the fluid edges.

For example, we can now discretize the (pentagonal) top fluid face in Figure 14(a), left, using (15) as

$$v_t A W_t = h\psi_{z_0} + l_0\psi_{x_1} - l_1\psi_{z_1} - h\psi_{x_0} + 0 \quad (21)$$

where v_t is the velocity normal component on the top face, $A = h^2$ is the area of a regular (non-cut) cell face, the l_i are the lengths of partial fluid edges, W_t is the fluid area fraction of the top face, and the 0 comes from the diagonal purple edge. Through a change of variables, we can finally rewrite this in a form consistent with the uniform case of (16). We hold the weighted contribution of each edge and face fixed, but rescale potentials by length and velocity by area. Letting $\psi'_{x_1} = \frac{l_0}{h}\psi_{x_1}$, $\psi'_{z_1} = \frac{l_1}{h}\psi_{z_1}$ and $v'_t = v_t W_t$, (21) becomes:

$$v'_t A = h(\psi_{z_0} + \psi'_{x_1} - \psi'_{z_1} - \psi_{x_0}). \quad (22)$$

With these rescalings applied to our input, we can perform uniform grid sweeping across such faces, and then undo the scalings to recover the desired discrete vector potential field.

This face-centric approach extends naturally to cover all geometries of grid faces cut by level set solids (i.e., the marching squares cases). As another example, in Figure 14(b), the original top face consists of a fluid part (nonzero flux) and a solid part (zero flux). These two (sub-)faces imply two equations,

$$\begin{aligned} v_t A W_t &= l_0\psi_{z_0} + l_2\psi_{x_1} - l_1\psi_{z_1} - h\psi_{x_0}, \\ 0 &= l_2\psi_{x_1} - h\psi'_{x_1}. \end{aligned} \quad (23)$$

where in the second equation we again assume zero potential on the two axis-aligned purple solid edges. Rescaling and sweeping suffices to recover the potentials, as before.

5.2.2 Approximate Gauge Correction with 3D Cut-Cells. Applying gauge correction is more difficult with polyhedral cut-cell solids, since discrete vector potential edge components can correspond to arbitrary, rather than Cartesian, directions. A polyhedral PDE solver could be used, e.g., mimetic finite differences [Lipnikov et al. 2006] or the diamond Laplacian [Bunge et al. 2021]. However, such a solver is more costly and interpolating the resulting non-Cartesian edge components would also be unwieldy. Moreover, as in 2D, we still require a ramping correction to precisely respect solid boundaries. Since our motivation for gauge correction is simply to ensure smoothness, minor local deviations from the Coulomb gauge are acceptable. We therefore enforce the gauge in a simple but approximate manner: we use the rescaled data from the voxelized view of our geometry determined above (Figure 14), and perform gauge correction on the *entire* uniform grid (including through the interior of solids). This voxelized approach also still lets us retain the use of the fast DST Poisson solver. Thus, our method efficiently provides smooth axis-aligned ψ values everywhere, which we use for B-spline kernel interpolation as in the uniform case.

5.3 Precise boundary enforcement in 3D

As in 2D, we propose an additive ramping-based approach for exact boundary enforcement in 3D, applied as a correction atop the approximate ψ values produced by grid interpolation. The 3D adaptation requires caution in treating the interacting ψ components.

5.3.1 Exterior domain boundaries. The domain exterior consists of axis-aligned planes with two tangential components of ψ lying on each boundary plane (Figure 15(a)). Recall that for a u -face, the flux is determined by ψ_y and ψ_z through $u = \frac{\partial\psi_z}{\partial y} - \frac{\partial\psi_y}{\partial z}$. To achieve exactly zero flux everywhere on the plane and avoid leakage, these terms must precisely cancel: we set the discrete ψ_y and ψ_z values to zero for simplicity, and consequently ramp the pointwise ψ_y and ψ_z values to zero as a particle approaches the border. (A careless choice of the ramp values can badly distort the velocity, as we saw in 2D in Figure 7.) Since we apply ramping to each component of ψ independently, this ramping strategy is identical to the 2D case (Section 4.2.3). As in 2D, we do not enforce exact ramping on inflow/outflow boundaries as its effects are not visually apparent.

5.3.2 Solid obstacles. For our explicit representation of the level set-based solid obstacles, we use triangle meshes obtained from

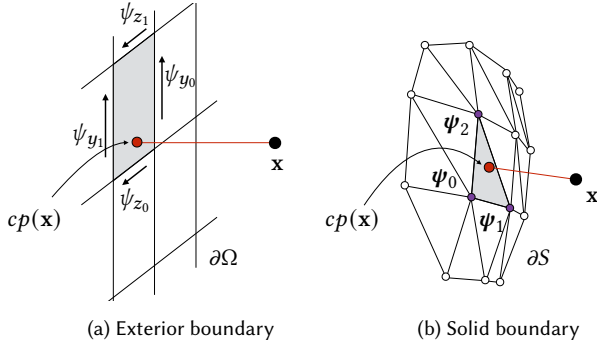


Fig. 15. **Boundary enforcement:** The black and red disks represent a particle near the boundary and its closest point on the boundary, respectively. (a) For planar exterior boundaries, the discrete ψ values (e.g., ψ_{y0} , ψ_{y1} , ψ_{z0} , ψ_{z1}) are zero by our boundary conditions, and the relevant axis components ($\psi_y(\mathbf{x})$ and $\psi_z(\mathbf{x})$) are ramped to zero for exact enforcement. (b) For triangulated solids, we first find discrete ψ values at the triangle vertices (e.g., ψ_0 , ψ_1 , ψ_2) which imply a perfectly tangential surface velocity, under barycentric interpolation of nodal ψ . When querying velocities, we ramp the full ψ vector towards this interpolated surface ψ at the closest point.

marching cubes. A key component of exact boundary enforcement with ramping is the proper choice of target ψ values (i.e., ψ_g) on the solid surface. Specifically, ψ_g should produce a zero normal velocity,

$$\mathbf{u} \cdot \mathbf{n} = (\nabla \times \psi_g) \cdot \mathbf{n} = 0, \quad (24)$$

where \mathbf{n} represents the solid normal vector.

For exterior boundaries, we were able to simply set constant discrete ψ_{tan} values in the solve, and by using those same constant values for the target ψ_g during ramping, (24) was trivially satisfied. However, one cannot in general find a single constant vector ψ_g that precisely compensates the existing (spatially varying) ambient interpolated ψ on irregular solids. Moreover, recall that an *arbitrary* choice of ψ_g will still severely distort the nearby tangential velocities, even if it satisfies (24).

We instead leverage the characteristics of triangular barycentric coordinates: the gradient of the barycentric interpolant is constant, thus $\nabla \times \psi_g$ yields a constant velocity (per triangle). (We also tried other interpolants satisfying (24) continuously on a triangulated surface (piecewise constant, Whitney), but found that barycentric produced the smoothest velocity fields.) Therefore, an appropriate choice of ψ at the triangle vertices (ψ_0 , ψ_1 , ψ_2 in Figure 15(b)), can satisfy (24), exactly and continuously. At the same time, we wish to minimize the perturbation of the velocity field induced by ramping, which we can do by encouraging the discrete ψ_i values at triangle vertices to be as close to the ambient ψ (i.e., the interpolated pointwise ψ before applying ramping) as possible. This yields an equality constrained quadratic minimization problem:

$$\begin{aligned} & \underset{\mathbf{y}}{\operatorname{argmin}} && \frac{1}{2}(\mathbf{y} - \boldsymbol{\psi}_{amb})^T(\mathbf{y} - \boldsymbol{\psi}_{amb}) \\ & \text{subject to} && \left(\nabla \times \sum_{i \in \{0,1,2\}} w_i(\mathbf{x}) \mathbf{y}_{t_i} \right) \cdot \mathbf{n}_t = 0 \quad \text{for } t \in T, \end{aligned}$$




Fig. 16. **Flow Past A Sphere:** A uniform layer of particles travels left to right under a steady state flow field past a sphere. Particles that exit the domain are frozen in place to emphasize their final distribution (right column). Grid resolution: $40 \times 20 \times 20$. (a) Direct velocity interpolation, with colliding particles halted on contact: Collided particles densely blanket the sphere. The hole in the final particle layer (right) highlights the erroneous trailing gaps in the flow behind the solid. (b) Direct velocity interpolation, with colliding particles projected back out: Repeated collision and projection leads to irregular particle patterns and the severe hole in the particle distribution persists. (c) Curl-Flow interpolation, with colliding particles halted on contact: Particles smoothly traverse the sphere with very few colliding. The flow also tightly follows behind the solid, so the final particle layer remains much more uniform (right).

where T is a set of (connected) solid triangles, $w_i(\mathbf{x})$ is the barycentric coordinate function for the i^{th} vertex for a spatial position \mathbf{x} , t_i is the global vertex index of the i^{th} vertex in the i^{th} triangle, \mathbf{n}_t is the triangle normal, and $\boldsymbol{\psi}_{amb}$ is a stack of the ambient ψ values at the triangle vertices (before ramping). If we write the constraints as $A\mathbf{y} = \mathbf{0}$, the optimality conditions yield the linear system

$$\begin{bmatrix} I & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \mathbf{y} \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\psi}_{amb} \\ \mathbf{0} \end{bmatrix} \quad (25)$$

or $AA^T \boldsymbol{\lambda} = A\boldsymbol{\psi}_{amb}$. We can find the desired discrete ψ_i at the solid vertices via $\mathbf{y} = \boldsymbol{\psi}_{amb} - A^T \boldsymbol{\lambda}$. Since the normal component of velocity only depends on the tangential components of ψ , one need only ramp the tangential components using the target values of $\boldsymbol{\psi}_{g,tan} = \boldsymbol{\psi}_g - (\boldsymbol{\psi}_g \cdot \mathbf{n})\mathbf{n}$. For ramping, $\boldsymbol{\psi}_g$ and $\boldsymbol{\psi}_{g,tan}$ can be queried by barycentric interpolation \mathbf{y} in the closest triangle.

Letting $\psi_{nor}(\mathbf{x})$ be the normal component of the ambient ψ , the final modified pointwise ψ with ramping is

$$\begin{aligned}\psi'(\mathbf{x}) &= \psi'_{tan}(\mathbf{x}) + \psi_{nor}(\mathbf{x}) \\ &= \psi_{tan}(\mathbf{x}) + (\alpha(\mathbf{x}, \mathbf{cp}(\mathbf{x})) - 1)(\psi_{tan}(\mathbf{cp}(\mathbf{x})) - \psi_{g,tan}(\mathbf{cp}(\mathbf{x}))) \\ &\quad + \psi_{nor}(\mathbf{x})\end{aligned}\tag{26}$$

Again, we use $cp(\mathbf{x})$ to represent the closest point on the solid triangles and $\alpha(\mathbf{x}) = ramp(d(\mathbf{x}/d_0))$ for a smooth transition toward $\psi_{g,tan}$. Here, d_0 is the influence radius (h in our case), and $d(\cdot)$ is the distance to the closest point $cp(\mathbf{x})$. We again employ the smoothstep function for $ramp(\cdot)$. This approach is much like (13) in 2D, except that the target $\psi_{g,tan}$ is not a constant since it is computed from barycentric interpolation.

The corresponding free-slip velocity is

$$\begin{aligned}\mathbf{u}'(\mathbf{x}) &= \nabla \times (\psi'_{tan}(\mathbf{x}) + \psi_{nor}(\mathbf{x})) \\ &= \nabla \times \psi(\mathbf{x}) + \nabla \alpha(\mathbf{x}, \mathbf{cp}(\mathbf{x})) \times (\psi_{tan}(\mathbf{cp}(\mathbf{x})) - \psi_{g,tan}(\mathbf{cp}(\mathbf{x}))) \\ &\quad + (\alpha(\mathbf{x}, \mathbf{cp}(\mathbf{x})) - 1) \{ \nabla \times (\psi_{tan}(\mathbf{cp}(\mathbf{x})) - \psi_{g,tan}(\mathbf{cp}(\mathbf{x}))) \}\end{aligned}\tag{27}$$

As the position \mathbf{x} approaches $\mathbf{cp}(\mathbf{x})$, $\nabla \times \psi_{tan}(\mathbf{x})$ and $\nabla \times \psi_{tan}(\mathbf{cp}(\mathbf{x}))$ cancel each other out, and due to barycentric interpolation, $\nabla \times \psi_{g,tan}(\mathbf{cp}(\mathbf{x}))$ is constant per triangle, which we constrained to have zero normal velocity in (25); thus exact zero flux on the solid surface. (Other components do not influence normal velocity.)

Due to the piecewise constant component from barycentric interpolation in (27), it is possible for additional velocity discontinuities to be introduced, beyond those coming from non-smooth changes in $\mathbf{cp}(\mathbf{x})$ (Section 4.2.3). However, we found these effects to be a reasonable tradeoff for incompressibility and precise boundary satisfaction. Moreover, the correction applied by our scheme is often small and any possible discontinuities can only occur within one cell width of solid obstacles or domain boundaries, whereas the discontinuities mentioned in related work (e.g., DG methods (Section 2.2.1) and most direct incompressible velocity interpolants (Section 2.2.2)) arise between adjacent elements or cells everywhere throughout the domain, even in the absence of obstacles.

To illustrate the improved behavior of 3D Curl-Flow with cut-cells in a basic scenario, we consider a flow past a solid sphere with a steady state velocity field (Figure 16), and release a single planar layer of passive particles. With trilinear advection, many of the particles collide with the obstacle, leaving a large gap in the particle field. With Curl-Flow, the particles instead flow more naturally around the obstacle.

6 RESULTS AND DISCUSSION

To consistently compare our Curl-Flow method against direct velocity interpolation while isolating interpolation-related effects, most of our comparisons change only the interpolation of velocity used for *passive* particle tracing, keeping the underlying simulated flow data the same. (The exception are tests that apply Curl-Flow to tracing the FLIP particle paths for FLIP advection.) Whenever a particle penetrates a solid object due to poor quality velocity fields, large timesteps, or insufficiently accurate path integration, a typical choice is to "resolve" the collision by projecting the particle back out of the solid, although this exacerbates clumping. To visually




Fig. 17. **Dynamic 2D Flow Comparison:** Uniformly sampled particles advected through a *dynamic* simulated 2D vector field with turbulent forcing on a 20×20 grid. Frames 1, 50, 100, and 150 are shown from left to right. Particle distribution remains more uniform with Curl-Flow.

highlight the pervasiveness of such errors, we freeze penetrating particles in place, unless otherwise noted (such as in Figure 16(b)). Simulation timings were gathered on a 2.8 GHz, 4-Core Intel Core i7 processor.

6.1 Particle Distribution Comparisons in 2D

The 2D results of Figures 3, 7, 8, and 9 used static velocity fields for illustration; below, we consider time-evolving (i.e., dynamic) flow simulations.

In Figure 17, we densely seed passive particles with blue noise sampling and advect them under a dynamic flow on a coarse 20×20 grid. We apply an upward force at the center of the domain and a smoothly varying random (noise) force over the entire domain in each frame for a turbulent effect. Visually, the resulting distribution again remains more uniform under our method. To better quantify this effect, we conceptually lay a three times finer grid of sub-cells over the domain and plot histograms of the sub-cells' particle counts in the first and last (150th) frame of the same example (Figure 18). Initially, most sub-cells contain about 5–10 particles, and the distribution is tightly clustered. The final Curl-Flow graph remains closely clustered, indicating still fairly uniform particle distribution, while the direct velocity interpolation graph spreads out severely. At the extreme ends, the direct case has 137 empty sub-cells and one sub-cell containing 149 particles.

In Figure 19 (top half), we exercise our 2D cut-cell treatment by observing passive particles under a dynamic horizontal flow past circular and jagged solids. Spurious gaps in the initially dense particle sampling quickly arise with direct velocity interpolation. Curl-Flow closely obeys the boundary while being strictly incompressible, sharply reducing the occurrence of gaps. Because interpolation errors depend on cell size, tripling the grid resolution (Figure 19, right pair of columns) proportionally reduces the direct method's gaps, but is costly and such errors persist for any finite resolution.

In FLIP-type simulations [Zhu and Bridson 2005], particles are not merely passive tracers but are instead used to transport velocity; as a result, particle density errors can be equally or more

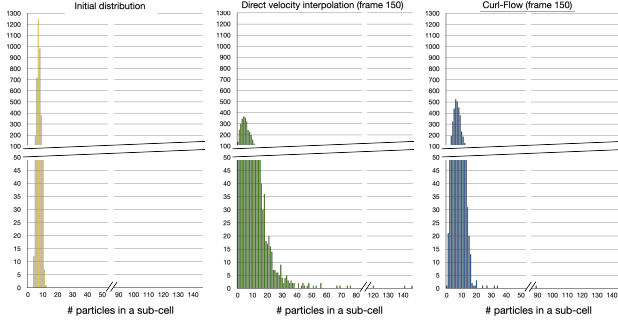


Fig. 18. **Particle Density Histograms:** Histograms of particles per sub-cell for the first and last (150th) frames of the example in Figure 17. The sub-cell width for measurement is $1/3$ the simulation cell width. The x-axis indicates the number of particles in a sub-cell, and the y-axis represents the number of sub-cells having that particle count. The graph for direct velocity interpolation has a longer and thicker tail of high density sub-cells, and dramatically more sub-cells with no particles at all.

problematic in that setting. Figure 19 (bottom half) applies FLIP using Curl-Flow for particle tracing (but *not* particle-to-grid (P2G) or grid-to-particle (G2P) transfers). Aside from FLIP particle seeding at the start and left inflow (8 particles per grid cell), no artificial density interventions were applied, e.g., reseeding, particle position adjustment, etc. Since P2G fails for empty cells, such cells fall back to semi-Lagrangian advection. We observe trailing gaps with direct velocity interpolation and, furthermore, its motion is more damped compared to FLIP with Curl-Flow.

6.2 Particle Distribution Comparisons in 3D

Static Flow In A Box. To explore incompressibility’s effects in 3D at near- and sub-grid scales, Figure 1 compares our full Curl-Flow method (including exterior boundary ramping) against direct velocity interpolants with (componentwise) trilinear and monotonic tricubic schemes [Fedkiw et al. 2001; Fritsch and Carlson 1980], as these are simple and common choices. The data is a static, extremely coarse $10 \times 10 \times 10$ discretely incompressible discrete vector field. The direct interpolants quickly exhibit particle clustering and spreading, forming visible low density regions and dense ring-like structures. By contrast, our result remains more uniformly distributed over time, despite no particle resampling, positional perturbations, or other remedies being applied. Clearly our method succeeds at enforcing incompressibility *at even the finest scales*.

Dynamic Flows Past Smooth and Jagged Obstacles. At medium scales, we next compare Curl-Flow to direct velocity interpolation under separate 3D dynamic horizontal flows past a sphere and a jagged obstacle (Figure 20). Here we consider grids ranging from $40 \times 20 \times 20$ to $80 \times 60 \times 60$. Only the particles within a narrow slice plane are shown to better visualize their density. The particles are purely passive, but for illustration we seed them throughout the domain at a density representative of typical particle-in-cell schemes, i.e., 8 per cell [Zhu and Bridson 2005]. As in the corresponding 2D examples, Curl-Flow retains more uniform particle sampling.

The average computational costs per timestep for our method applied to the dynamic flows above are given in Table 1. For reference, we also include timings for the underlying fluid simulator’s pressure projection, as it is often a bottleneck in fluid solvers. For these tests, we solve *both* pressure projection and gauge correction with the basic conjugate gradient method in the Eigen library [Guennebaud et al. 2010], with relative tolerance 1×10^{-8} . We used Eigen’s LDLT direct solver for the solid surface problem (25), since we found it can be ill-conditioned. The observed total cost of reconstructing ψ (for the grid and solid surfaces) is noticeably less than for pressure projection (around 50% – 75% as much). The cost of advecting particles with our interpolant is slightly more than for trilinear interpolation (around 10% – 30% more).

6.3 Smoke Simulations and Performance

In the following two examples, we consider rising smoke plumes in 150^3 grids with different obstacles: a stylized calf (Figure 21) and Houdini’s rubber toy (Figure 22). For a simple age-based dissipation effect, we removed particles after 4 seconds of existence, but to emphasize collision errors we did not delete any particles that collided with the obstacle. We supplied up to 4M purely passive smoke particles for each example.

Our ramping method enforces a pointwise boundary-respecting velocity field under Curl-Flow, but particles advected with discrete timesteps will still collide with solid obstacles, especially with fast dynamic flows and/or complex solids. For the smoke examples (our version and the baseline), we apply substepping to (only) particle advection to partially ameliorate this innate problem, constraining the timestep so that particles cannot move more than one grid cell per substep. Using more substeps would further reduce collisions, at the cost of increasing computational time. We additionally apply a simple “collision-aware” substepping upon first contact: if a particle hits an obstacle, we push it slightly back from the hit position (10% of the distance the particle moved) and then perform a new partial substep from that position for the remaining time. (If it is still colliding afterwards, it is then frozen.)

To demonstrate Curl-Flow’s practicality, these smoke plume tests were carried out by implementing our method into Houdini’s smoke solver (*Pyro*), with various modifications (e.g., we used our own cut-cell pressure projection, solved with Eigen’s CG). Unlike the prior examples, we used a MacCormack scheme (Houdini’s default) for velocity advection and we used our DST Poisson solver for fast gauge correction. In both examples, Curl-Flow shows superior boundary-respecting behavior; particles flow over and around the shape, instead of into it and halting. This effect can be observed by inspecting the density and brightness of smoke in the later images. Timings are provided in Table 2.

Curl-Flow incurs preprocessing costs at each step to reconstruct discrete ψ on the regular grid and on solid vertices, i.e., primarily a volumetric Poisson problem and a surface-based least squares problem. Fortunately, our DST solver makes gauge correction efficient and the fact that the surface solve is lower dimensional likewise keeps its scaling reasonable. We investigate these costs in the ψ reconstruction scaling tests of Table 3. At 256^3 , ψ construction is two orders of magnitude faster than cut-cell pressure projection

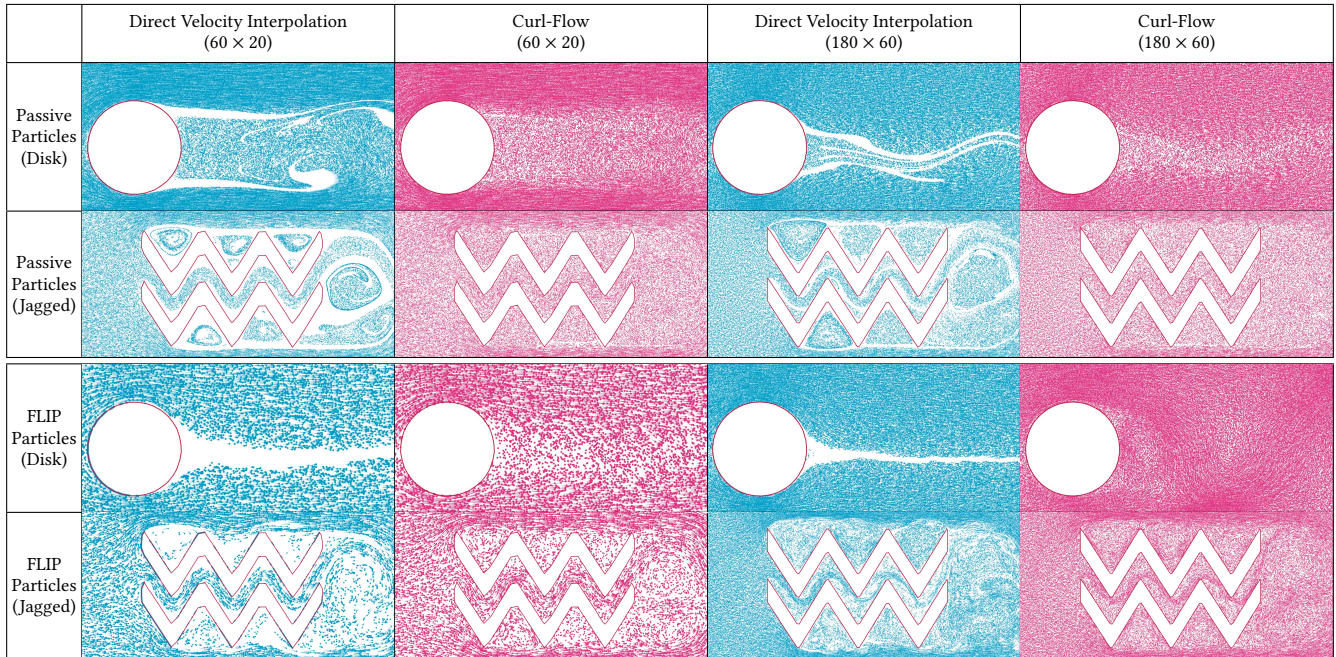


Fig. 19. **Dynamic 2D Flow in A Wind Tunnel Past A Solid:** Passive tracer particles (top two rows) and FLIP particles (bottom two rows) undergo a dynamic horizontal flow past static obstacles. Particles colliding into objects are halted in place. Direct velocity interpolation creates spurious gaps (blue), even at triple the grid resolution; Curl-Flow interpolation tightly and incompressibly follows the solid objects, significantly reducing the gaps in the flow (red).




Fig. 20. **Dynamic 3D Flow in A Wind Tunnel Past A Solid (middle frame):** A slice-plane view of particles undergoing a time-dependent horizontal flow past a solid object in 3D. The color of the particles represents the particle densities: the density increases from red to yellow. Particles colliding into objects are halted in place for illustration. Regardless of the grid resolution, direct velocity interpolation creates spurious gaps, while our Curl-Flow interpolation tightly follows the solid object, significantly reducing the gaps in the flow and maintaining better particle distribution.

with Eigen's CG (Fourier methods do not support cut-cells). With major implementation effort, advanced multigrid and GPU solvers can sharply reduce the cost of pressure projection [Raateland et al. 2022; Shao et al. 2022]; however, even in the worst case, since both

are Poisson problems, pressure projection and gauge correction would have around the same cost (e.g., see rows 3 and 4 of Table 4).

The remaining additional overhead of our method arises during particle advection: Curl-Flow has a larger stencil size for interpolation (quadratic vs. linear), and our ramping strategy requires

Table 1. Average computational time per timestep for didactic examples.

Examples	Interpolation Method	Pressure Projection (s)	Grid ψ Construction (s)	Solid ψ Construction (s)	Particle Advection (s)	# Particles	# Solid Vertices (Triangles)
Flow Past A Sphere ($40 \times 20 \times 20$)	Direct Velocity Interpolation	2.214×10^{-2}	–	–	3.962×10^{-1}	117,482	726
	Curl-Flow		9.040×10^{-3}	5.887×10^{-3}	4.443×10^{-1}	117,483	(1448)
Flow Past A Sphere ($80 \times 40 \times 40$)	Direct Velocity Interpolation	4.268×10^{-1}	–	–	3.010	949,272	2934
	Curl-Flow		1.812×10^{-1}	3.888×10^{-2}	3.708	949,277	(5864)
Flow Past A Jagged Obstacle ($40 \times 30 \times 30$)	Direct Velocity Interpolation	9.942×10^{-2}	–	–	8.158×10^{-1}	271,144	3140
	Curl-Flow		4.262×10^{-2}	2.907×10^{-2}	1.050	271,148	(6272)
Flow Past A Jagged Obstacle ($80 \times 60 \times 60$)	Direct Velocity Interpolation	2.070	–	–	6.885	2,035,118	12454
	Curl-Flow		9.001×10^{-1}	2.070×10^{-1}	7.723	2,035,124	(24900)



Fig. 21. **A Smoke Plume Simulation with A Solid Shaped like A Calf.** In the direct velocity interpolation case, a large number of particles erroneously collide with the solid and are then frozen, especially around the ears and underside of the head; with Curl-Flow, more particles flow smoothly over and around the object's surface and since more smoke remains active it is denser and brighter (middle column, above the face, and see video).

additional operations, e.g., finding the closest point on the solid, evaluating a second interpolation, etc., unlike simple grid-based velocity interpolation. In the calf and dragon examples, collided (frozen) particles also no longer undergo advection, which significantly benefits the direct approach. In Figure 21, 1,508,617 (top) and 885 (bottom) particles are frozen and in Figure 22, 2,136,520 (top) and 2,236 (bottom) particles are frozen at the last frame. Advection costs are heavily dependent on the number of active particles and can

become the bottleneck for massive particle counts. We anticipate that interpolation could be further optimized. Most importantly, since particle advection is done in a fully parallel manner, more compute cores immediately translate into speedups. Genuinely passive particles can also be traced in an entirely parallel post-process, after simulation completes.




Fig. 22. **A Smoke Plume Simulation with A Rubber Toy.** The lingering outline of dense smoke from collided particles (frozen for emphasis) is conspicuous on the belly in the direct velocity interpolation case. For Curl-Flow, the smoke instead flows around and past the object, so it remains visibly denser after going by (middle and right columns).

Table 2. Average computational time per timestep for smoke simulations of Figure 21 and Figure 22. 400K particles were supplied every second, and they fade away after 4 seconds.

	Interpolation Method	Pressure Projection (s)	Grid ψ Construction (s)	Solid ψ Construction (s)	Particle Advection (s)	# Solid Vertices (Triangles)
Stylized Calf (150^3)	Direct Velocity Interpolation	—	—	—	1.213	16,450 (32,896)
	Curl-Flow	3.898×10^1	3.807×10^{-1}	3.571×10^{-1}	2.749×10^1	16,450 (32,896)
Rubber Toy (150^3)	Direct Velocity Interpolation	—	—	—	4.049×10^{-1}	21,050 (42,096)
	Curl-Flow	3.847×10^1	3.687×10^{-1}	4.431×10^{-1}	1.626×10^1	21,050 (42,096)

Table 3. Average computational time per timestep for discrete field construction of Curl-Flow at various scales.

	Pressure Projection (s)	Grid ψ Construction (s)	Solid ψ Construction (s)	# Solid Vertices (Triangles)
Stylized Calf (32^3)	4.875×10^{-2}	6.540×10^{-3}	4.026×10^{-3}	726 (1,448)
Stylized Calf (64^3)	8.419×10^{-1}	2.920×10^{-2}	2.483×10^{-2}	2,974 (5,944)
Stylized Calf (128^3)	1.984×10^1	2.158×10^{-1}	1.987×10^{-1}	11,946 (23,888)
Stylized Calf (256^3)	3.685×10^2	1.713	1.570	47,930 (95,856)
Rubber Toy (32^3)	4.804×10^{-2}	1.229×10^{-2}	9.849×10^{-3}	1,836 (3,664)
Rubber Toy (64^3)	8.998×10^{-1}	2.954×10^{-2}	3.257×10^{-2}	7,508 (15,008)
Rubber Toy (128^3)	1.997×10^1	2.111×10^{-1}	2.509×10^{-1}	30,600 (61,192)
Rubber Toy (256^3)	3.844×10^2	1.705	2.166	123,284 (246,560)

6.4 Vector Potential Reconstruction Comparisons

Several alternative, but more costly, methods to reconstruct edge-based vector potential fields from face-based vector fields have previously been proposed [Ando et al. 2015; Bao et al. 2017; Sato et al. 2015; Silberman et al. 2019]. Since those methods mostly do not handle cut-cell obstacles, to fairly compare computational costs we consider a flow in an empty box domain. We adapted the methods of

Ando et al. and Silberman et al. into our framework, replacing our ψ reconstruction step. We consider the combined cost of pressure projection and ψ reconstruction for each method in Table 4, since Ando et al. achieve both simultaneously. For a fair comparison, the first four rows use Eigen’s CG for all solves. Since our DST-based gauge correction can always be used whether obstacles are present or not, we present this result in its own row. Despite different

Table 4. Average computational time per timestep for discrete ψ construction on the 150^3 grid simulation of Figure 23 for various methods.

	Pressure Projection (s)	Initial ψ Construction (s)	Gauge Correction (s)	Total ψ Construction (s)
Vector Poisson, Dimensionally-Coupled (Ando et al. [2015])	-	-	-	152.0
Vector Poisson, Componentwise (e.g., Bao et al. [2017])	-	-	-	81.87
Pressure Projection + Gauge Correction (CG) (Silberman et al. [2019])	37.13	5.723	25.51	68.37
Ours: Pressure Projection + Gauge Correction (CG) + Parallel Sweeping	37.13	0.05904	24.33	61.51
Ours: Pressure Projection + Gauge Correction (DST) + Parallel Sweeping	37.13	0.05781	0.3304	37.52




Fig. 23. **A Rising Plume Using Curl-Flow** : Five different ψ construction methods are plugged into Curl-Flow to produce this rising plume scene. They all yield the same results shown above, with different computational costs (Table 4).

computational costs, the solutions are numerically consistent and the results visually indistinguishable (Figure 23).

Comparison to Ando et al. [2015]: Ando et al. solve a single, dimensionally coupled vector Poisson problem for ψ from a *divergent* velocity field, avoiding a separate pressure projection. We modified this approach to satisfy our domain boundary conditions (Section 5.3.1). As shown in Table 4, its computational cost (topmost row) is expensive compared to our method (bottom two rows): about $2.5 - 4\times$ slower for the 150^3 grid of Figure 23. Unless boundary conditions couple them (as is often the case for free surfaces and some solid boundary treatments [Ando et al. 2015]), the dimensionally-coupled vector Poisson equation can be split into three independent scalar Poisson equations for efficiency. Applying this special case optimization nearly halves the cost (Table 4, second row), but it remains about $1.3\times - 2.2\times$ the cost of our approach.

Table 5. Average computational time for constructing ψ , within the code of Bao et al. [2017] for an immersed deforming membrane. For the "Curl-Flow" comparison, we replace only their ψ construction part.

	ψ construction time (s)
Bao et al. [2017]	2.324×10^{-3}
Curl-Flow	9.397×10^{-4}

Comparison to Silberman et al. [2019]: The reconstruction method of Silberman et al. from electromagnetics is most similar to ours: it recovers ψ from an initially divergence-free input vector field, using a cell-by-cell construction of an initial vector potential, followed by Poisson-based gauge correction. Silberman et al. assumed infinite boundary conditions to enable gauge correction by FFT. We adapted this idea by matching our boundary conditions and linear solver (CG) – see Table 4, third row. Compared to our sweeping approach, the initial cell-by-cell ψ construction of Silberman et al. is slower because it is inherently serial and it builds all three components of ψ , rather than the two in our sweeping scheme. It also requires complicated case-by-case code: it constructs 12 edge values on a cell, and adjusts them for consistency depending on which neighbors have already been processed. The cost of our parallel sweeping strategy is *two orders of magnitude faster* than Silberman et al. (Table 4, third column), and negligible compared to the Poisson solves. Overall, the (adapted) method of Silberman et al. [2019] is 11% slower than ours, or 82% slower when we use DST. (We emphasize that they also do not handle our boundary conditions nor interior obstacles, so their FFT approach never applies to our setting.)

Comparison to Bao et al. [2017]: Bao et al. use a vector potential in their immersed boundary method. They assume divergence-free input and exploit the lack of (cut-cell) obstacles to use the componentwise vector Poisson form (three scalar Poisson problems). They assume periodic boundaries to solve with the FFT. For a fair comparison, we used their (MATLAB) code and replaced only the discrete ψ construction step with our method, similarly using FFT for gauge correction. Our method for ψ construction is $\sim 2.5\times$ faster, as shown in Table 5.

Although our focus is fluid velocity interpolation, vector potential reconstruction has already been used in graphics (for fluid control [Sato et al. 2021, 2015] and visualization [Biswas et al. 2016]), as well as electromagnetics and beyond. Our new approach can thus offer immediate speedups in all of these domains, e.g., based on Table 4,




Fig. 24. **Deforming elastic membrane test from [Bao et al. 2017]**: A spherical membrane immersed in fluid is deformed with an initial velocity, before surface tension returns it to a spherical shape. We replace Bao’s ψ reconstruction method with ours. Our methods produces identical results but faster (Table 5).

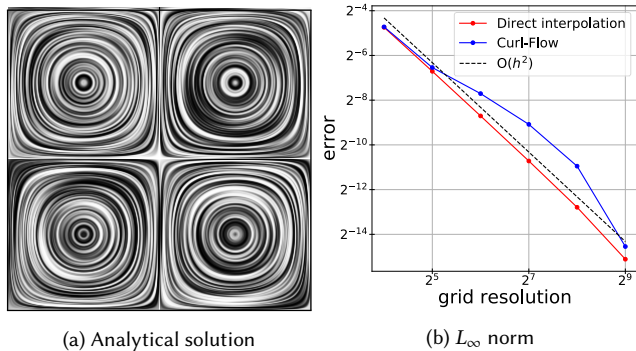


Fig. 25. **Convergence of Velocity Interpolation**: We sample discrete grid velocities from the analytically divergence-free field of (a), enforce discrete incompressibility, interpolate using direct (bilinear) velocity interpolation and Curl-Flow interpolation, and measure the error. Both methods show approximately second-order convergence. See error data in Table 1 in supplemental material.

perhaps doubling or quadrupling the speed of the most expensive step of the fluid guiding method of Sato et al. [2021].

6.5 Convergence of Curl-Flow Velocity Interpolation

Since Curl-Flow uses intermediate vector potentials rather than directly interpolating grid velocities, one may ask if it is a valid interpolant at all: does the resulting continuous field converge? We study this question by evaluating how well the interpolated field agrees with an input analytically divergence-free field under refinement. Since the method’s input should be a discretely divergence-free velocity field defined on the staggered grid, we first sample grid velocities from an input analytical field and apply pressure projection for discrete incompressibility, then apply the chosen interpolant (at a fine sampling of points) and evaluate their error with respect to the original analytical field. We use the field

$$\mathbf{u}(x, y) = (\sin(2\pi x)\cos(2\pi y), -\cos(2\pi x)\sin(2\pi y)),$$

shown in Figure 25(a). Since Curl-Flow’s quadratic kernel will query some sample positions outside the domain, we linearly extrapolate ψ values to such samples for the convergence test. Direct (bilinear) velocity interpolation yields second-order convergence and Curl-Flow shows similar behavior (Figure 25).

7 CONCLUSIONS AND FUTURE WORK

Large time steps in particle tracing, poorly enforced boundaries, and inadequate pressure solver tolerances are known to cause density/volume drift; many post-compensation strategies have been developed [Ando et al. 2012; Kugelstadt et al. 2019; Sato et al. 2018b; Takahashi and Lin 2019]. Our work is the first in computer graphics to identify divergent velocity interpolation as an additional factor. Curl-Flow, tailored to plug into popular grid-based cut-cell fluid simulators, addresses this issue up front by constructing pointwise divergence-free velocity fields that also respect obstacles. It offers better long-term particle distributions and natural motions around obstacles and, when used for FLIP particles, yields more dynamic motion and/or reduces the need for post-compensation. We believe our work opens a new direction in the design of advection schemes for fluid animation and raises compelling questions for future study.

We did not consider moving obstacles, which are a useful element of many practical scenarios. Allowing for non-zero boundary fluxes during vector potential reconstruction seems feasible; however, ensuring exactly collision-safe trajectories near Eulerian cut-cell solids (i.e., the aim of the present work) is significantly complicated by the Lagrangian nature of moving obstacles. We have also only considered grid domains that are fully open, fully closed, or “wind-tunnel”-like. More general boundary conditions are an obvious next step, including free surface boundaries for liquid animation.

For our ramping ψ towards solids, we found that seeking a minimal perturbation to ψ gave visually natural results, but other criteria are possible (e.g., controlling $\nabla \times \psi_t$). A drawback of boundary handling in Curl-Noise-type schemes is that, though always incompressible, they can yield discontinuities when the closest polygon facet changes. A challenging extension would be to develop a strictly continuous variant of Curl-Noise, perhaps by incorporating a smooth distance measure [Madan and Levin 2022].

While not a salient issue for simply tracing particle paths, interpolation of higher order than linear can yield overshooting. We sidestepped this possibility in our FLIP tests by not using the Curl-Flow field for P2G or G2P transfers: any new extrema will not feed back into the dynamic velocity field. Using Curl-Flow for G2P (or all the interpolations within a semi-Lagrangian update) may necessitate some form of limiter. Straightforward velocity clamping [Selle et al. 2008] would sacrifice incompressibility, so falling back to linear Curl-Flow (with kinks) may be preferable.

Lastly, the ability to edit fluid simulations in an *efficient* post-process remains highly desirable [Pan et al. 2013; Sato et al. 2018a] due to their cost and limited controllability. Our work, and that of Sato et al. [2021; 2015], hint that vector potentials may unlock fast, divergence-free editing capabilities by combining discrete *simulation* tools and continuous *procedural* tools.

REFERENCES

- R. Albanese and G. Rubinacci. 1990. Magnetostatic field computations in terms of two-component vector potentials. *Internat. J. Numer. Methods Engrg.* 29, 3 (mar 1990), 515–532. <https://doi.org/10.1002/nme.1620290305>
- Ryoichi Ando, Nils Thuerey, and Chris Wojtan. 2015. A stream function solver for liquid simulations. *ACM Transactions on Graphics* 34, 4 (jul 2015), 53:1–53:9.
- Ryoichi Ando, Nils Thuerey, and Reiji Tsuruno. 2012. Preserving fluid sheets with adaptively sampled anisotropic particles. *IEEE Transactions on Visualization and Computer Graphics* 18, 8 (2012), 1202–1214.

- Alexis Angelidis and Fabrice Neyret. 2005. Simulation of smoke based on vortex filament primitives. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*. ACM, 87–96.
- Vinicius C Azevedo, Christopher Batty, and Manuel M Oliveira. 2016. Preserving geometry and topology for fluid flows with thin obstacles and narrow gaps. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 97.
- Dinshaw S Balsara. 2001. Divergence-free adaptive mesh refinement for magnetohydrodynamics. *J. Comput. Phys.* 174, 2 (2001), 614–648.
- Dinshaw S Balsara. 2004. Second-order-accurate schemes for magnetohydrodynamics with divergence-free reconstruction. *The Astrophysical Journal Supplement Series* 151, 1 (2004), 149.
- Dinshaw S Balsara. 2009. Divergence-free reconstruction of magnetic fields and WENO schemes for magnetohydrodynamics. *J. Comput. Phys.* 228, 14 (2009), 5040–5056.
- Yuanxun Bao, Aleksandar Donev, Boyce E Griffith, David M McQueen, and Charles S Peskin. 2017. An Immersed Boundary method with divergence-free velocity interpolation and force spreading. *Journal of computational physics* 347 (2017), 183–206.
- Ayan Biswas, Richard Strelitz, Jonathan Woodring, Chun-Ming Chen, and Han-Wei Shen. 2016. A scalable streamline generation algorithm via flux-based isocontour extraction. In *Proceedings of the 16th Eurographics Symposium on Parallel Graphics and Visualization*. 69–78.
- R. Bridson. 2015. *Fluid Simulation for Computer Graphics, Second Edition*. Taylor & Francis.
- Robert Bridson, Jim Houriham, and Marcus Nordenstam. 2007. Curl-Noise for Procedural Fluid Flow. *ACM Trans. Graph.* 26, 3 (July 2007), 46–es. <https://doi.org/10.1145/1276377.1276435>
- Tyson Brochu, Todd Keeler, and Robert Bridson. 2012. Linear-time smoke animation with vortex sheet meshes. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. Eurographics Association, 87–95.
- Astrid Bunge, Mario Botsch, and Marc Alexa. 2021. The Diamond Laplace for Polygonal and Polyhedral Meshes. In *Computer Graphics Forum*, Vol. 40. Wiley Online Library, 217–230.
- Hugo Casquero, Yongjie Jessica Zhang, Carles Bona-Casas, Lisandro Dalcin, and Hector Gomez. 2018. Non-body-fitted fluid–structure interaction: Divergence-conforming B-splines, fully-implicit dynamics, and variational formulation. *J. Comput. Phys.* 374 (2018), 625–653.
- Bernardo Cockburn, Fengyan Li, and Chi-Wang Shu. 2004. Locally divergence-free discontinuous Galerkin methods for the Maxwell equations. *J. Comput. Phys.* 194, 2 (2004), 588–610.
- Qiaodong Cui, Timothy Langlois, Pradeep Sen, and Theodore Kim. 2021. Spiral-spectral fluid simulation. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–16.
- Fang Da, David Hahn, Christopher Batty, Chris Wojtan, and Eitan Grinspun. 2016. Surface-only liquids. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–12.
- Fernando De Goes, Mathieu Desbrun, Mark Meyer, and Tony DeRose. 2016. Subdivision exterior calculus for geometry processing. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–11.
- Tyler De Witt, Christian Lessig, and Eugene Fiume. 2012. Fluid simulation using Laplacian eigenfunctions. *ACM Transactions on Graphics (TOG)* 31, 1 (2012), 10.
- Ivan DeWolf. 2006. *Divergence-free noise*. Technical Report. Martian Labs., 2005.
- Marvin Eisenberger, Zorah Löhner, and Daniel Cremers. 2018. Divergence-Free Shape Interpolation and Correspondence. *arXiv preprint arXiv:1806.10417* (2018).
- Sharif Elcott, Yiyang Tong, Eva Kanso, Peter Schröder, and Mathieu Desbrun. 2007. Stable, circulation-preserving, simplicial fluids. *ACM Transactions on Graphics (TOG)* 26, 1 (2007), 4.
- Douglas Enright, Stephen Marschner, and Ronald Fedkiw. 2002. Animation and rendering of complex water surfaces. *ACM Transactions on Graphics (TOG)* 21, 3 (2002), 736–744.
- John A Evans and Thomas JR Hughes. 2013. Isogeometric divergence-conforming B-splines for the unsteady Navier–Stokes equations. *J. Comput. Phys.* 241 (2013), 141–167.
- Ronald Fedkiw, Jos Stam, and Henrik Wann Jensen. 2001. Visual simulation of smoke. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. 15–22.
- F. N. Fritsch and R. E. Carlson. 1980. Monotone Piecewise Cubic Interpolation. *SIAM J. Numer. Anal.* 17, 2 (1980), 238–246.
- Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. <http://eigen.tuxfamily.org>.
- Johnny Guzmán and Michael Neilan. 2014. Conforming and divergence-free Stokes elements in three dimensions. *IMA J. Numer. Anal.* 34, 4 (2014), 1489–1508.
- Ben Houston, Chris Bond, and Mark Wiebe. 2003. A unified approach for modeling complex occlusions in fluid simulations. In *ACM SIGGRAPH 2003 Sketches*. 1–1.
- Antonio Huerta, Yolanda Vidal, and Pierre Villon. 2004. Pseudo-divergence-free element free Galerkin method for incompressible fluid flow. *Computer Methods in Applied Mechanics and Engineering* 193, 12 (2004), 1119 – 1136.
- P Jenny, SB Pope, M Muradoglu, and DA Caughey. 2001. A hybrid algorithm for the joint PDF equation of turbulent reactive flows. *J. Comput. Phys.* 166, 2 (2001), 218–252.
- Chenfanfu Jiang, Craig Schroeder, Andrew Selle, Joseph Teran, and Alexey Stomakhin. 2015. The affine particle-in-cell method. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 51.
- Chenfanfu Jiang, Craig Schroeder, Joseph Teran, Alexey Stomakhin, and Andrew Selle. 2016. The material point method for simulating continuum materials. In *ACM SIGGRAPH 2016 Courses*. 1–52.
- Theodore Kim, Nils Thürey, Doug James, and Markus Gross. 2008. Wavelet Turbulence for Fluid Simulation. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 1–6. <https://doi.org/10.1145/1360612.1360649>
- Tassilo Kugelstadt, Andreas Longva, Nils Thürey, and Jan Bender. 2019. Implicit Density Projection for Volume Conserving Liquids. *IEEE Computer Architecture Letters* 01 (2019), 1–1.
- Egor Larionov, Christopher Batty, and Robert Bridson. 2017. Variational stokes: a unified pressure-viscosity solver for accurate viscous liquids. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 101.
- Philip L Lederer, Alexander Linke, Christian Merdon, and Joachim Schöberl. 2017. Divergence-free reconstruction operators for pressure-robust Stokes discretizations with continuous pressure finite elements. *SIAM J. Numer. Anal.* 55, 3 (2017), 1291–1314.
- Christoph Lehrenfeld and Joachim Schöberl. 2016. High order exactly divergence-free hybrid discontinuous Galerkin methods for unsteady incompressible flows. *Computer Methods in Applied Mechanics and Engineering* 307 (2016), 339–361.
- Alexander Linke. 2012. A divergence-free velocity reconstruction for incompressible flows. *Comptes Rendus Mathématique* 350, 17–18 (2012), 837–840.
- Konstantin Lipnikov, Mikhail Shashkov, and Danil Svyatskiy. 2006. The mimetic finite difference discretization of diffusion problem on unstructured polyhedral meshes. *J. Comput. Phys.* 211, 2 (2006), 473–491.
- Svenja Lowitzsch. 2005. Matrix-valued radial basis functions: stability estimates and applications. *Advances in Computational Mathematics* 23, 3 (2005), 299–315. <https://doi.org/10.1007/s10444-004-1786-8>
- Abhishek Madan and David IW Levin. 2022. Fast Evaluation of Smooth Distance Constraints on Co-Dimensional Geometry. *ACM Transactions on Graphics (TOG)* (to appear) (2022).
- Jakob M Maljaars, Robert Jan Labeur, and Matthias Möller. 2018. A hybridized discontinuous Galerkin framework for high-order particle–mesh operator splitting of the incompressible Navier–Stokes equations. *J. Comput. Phys.* 358 (2018), 150–172.
- J.B. Manges and Z.J. Cendes. 1995. A generalized tree-cotree gauge for magnetic field computation. *IEEE Transactions on Magnetics* 31, 3 (may 1995), 1342–1347.
- Colin P McNally. 2011. Divergence-free interpolation of vector fields from point values—exact $\text{div} \mathbf{B} = 0$ in numerical simulations. *Monthly Notices of the Royal Astronomical Society: Letters* 413, 1 (2011), L76–L80.
- DW Meyer and P Jenny. 2004. Conservative velocity interpolation for PDF methods. In *PAMM: Proceedings in Applied Mathematics and Mechanics*, Vol. 4. Wiley Online Library, 466–467.
- Yen Ting Ng, Chohong Min, and Frédéric Gibou. 2009. An Efficient Fluid-solid Coupling Algorithm for Single-phase Flows. *J. Comput. Phys.* 228, 23 (Dec. 2009), 8807–8829.
- Zherong Pan, Jin Huang, Yiyang Tong, Changxi Zheng, and Hujun Bao. 2013. Interactive localized liquid motion editing. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 1–10.
- Sang Il Park and Myoung Jun Kim. 2005. Vortex fluid for gaseous phenomena. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*. ACM, 261–270.
- Charles S Peskin. 2002. The immersed boundary method. *Acta numerica* 11 (2002), 479–517.
- Tobias Pfaff, Nils Thürey, and Markus Gross. 2012. Lagrangian vortex sheets for animating fluids. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 112.
- Konstantin Poelke and Konrad Polthier. 2016. Boundary-aware Hodge decompositions for piecewise constant vector fields. *Computer-Aided Design* 78 (2016), 126–136.
- Adina E Pusok, Boris JP Kaus, and Anton A Popov. 2017. On the quality of velocity interpolation schemes for marker-in-cell method and staggered grids. *Pure and Applied Geophysics* 174, 3 (2017), 1071–1089.
- Wouter Raateland, Torsten Hadrich, Jorge Alejandro Amador Herrera, Daniel T Banuti, Wojciech Paubicki, Sören Pirk, Klaus Hildebrandt, and Dominik L Michels. 2022. DCGrid: An Adaptive Grid Structure for Memory-Constrained Fluid Simulation on the GPU. (2022).
- Nick Rasmussen, Douglas Enright, Due Nguyen, Sebastian Marino, Nigel Sumner, Willi Geiger, Samir Hoon, and Ronald Fedkiw. 2004. Directable photorealistic liquids. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*. 193–202.
- Bharath Ravu, Murray Rudman, Guy Metcalfe, Daniel Lester, and Devang Khakhar. 2016. Creating analytically divergence-free velocity fields from grid-based data. *J. Comput. Phys.* 323 (07 2016). <https://doi.org/10.1016/j.jcp.2016.07.018>
- Sander Rhebergen and Garth N Wells. 2018. A hybridizable discontinuous Galerkin method for the Navier–Stokes equations with pointwise divergence-free velocity field. *Journal of Scientific Computing* 76, 3 (2018), 1484–1501.
- Syuhel Sato, Yoshinori Dobashi, and Theodore Kim. 2021. Stream-Guided Smoke Simulation. *ACM Transactions on Graphics (TOG)* 40, 4, Article 161 (2021).

- Syuhei Sato, Yoshinori Dobashi, and Tomoyuki Nishita. 2018a. Editing fluid animation using flow interpolation. *ACM Transactions on Graphics (TOG)* 37, 5 (2018), 1–12.
- Syuhei Sato, Yoshinori Dobashi, Yonghao Yue, Kei Iwasaki, and Tomoyuki Nishita. 2015. Incompressibility-preserving deformation for fluid flows using vector potentials. *The Visual Computer* 31, 6 (2015), 959–965.
- Takahiro Sato, Christopher Wojtan, Nils Thuerey, Takeo Igarashi, and Ryoichi Ando. 2018b. Extended narrow band FLIP for liquid simulations. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 169–177.
- Hagit Schechter and Robert Bridson. 2008. Evolving sub-grid turbulence for smoke animation. In *Proceedings of the 2008 ACM SIGGRAPH/Eurographics symposium on Computer animation*. Eurographics Association, 1–7.
- Craig Schroeder, Ritoban Roy Chowdhury, and Tamar Shinar. 2022. Local divergence-free polynomial interpolation on MAC grids. *J. Comput. Phys.* 468 (2022), 111500.
- Andrew Selle, Ronald Fedkiw, Byungmoon Kim, Yingjie Liu, and Jarek Rossignac. 2008. An unconditionally stable MacCormack method. *Journal of Scientific Computing* 35, 2-3 (2008), 350–371.
- Han Shao, Libo Huang, and Dominik L. Michels. 2022. A Fast Unsmoothed Aggregation Algebraic Multigrid Framework for the Large-Scale Simulation of Incompressible Flow. *ACM Transaction on Graphics* 41, 4, Article 49 (07 2022).
- Zachary J. Silberman, Thomas R. Adams, Joshua A. Faber, Zachariah B. Etienne, and Ian Ruchlin. 2019. Numerical generation of vector potentials from specified magnetic fields. *J. Comput. Phys.* 379 (2019), 421 – 437. <https://doi.org/10.1016/j.jcp.2018.12.006>
- Jos Stam. 1999. Stable Fluids. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 121–128. <https://doi.org/10.1145/311535.311548>
- Jos Stam and Eugene Fiume. 1993. Turbulent wind fields for gaseous phenomena. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*. 369–376.
- Michael Steffen, Robert M Kirby, and Martin Berzins. 2008. Analysis and reduction of quadrature errors in the material point method (MPM). *International journal for numerical methods in engineering* 76, 6 (2008), 922–948.
- Tetsuya Takahashi and Ming C Lin. 2019. A geometrically consistent viscous fluid solver with two-way fluid-solid coupling. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 49–58.
- Yiyong Tong, Santiago Lombeyda, Anil N Hirani, and Mathieu Desbrun. 2003. Discrete multiscale vector field decomposition. *ACM Transaction on Graphics (TOG)* 22, 3 (2003), 445–452.
- Charles Van Loan. 1992. *Computational frameworks for the fast Fourier transform*. SIAM.
- Wolfram Von Funck, Holger Theisel, and Hans-Peter Seidel. 2006. Vector field based shape deformations. In *ACM Transactions on Graphics (TOG)*, Vol. 25. ACM, 1118–1125.
- Hongliang Wang, Roberto Agrusta, and Jeroen van Hunen. 2015. Advantages of a conservative velocity interpolation (CVI) scheme for particle-in-cell methods with application in geodynamic modeling. *Geochemistry, Geophysics, Geosystems* 16, 6 (2015), 2015–2023.
- Ke Wang, Yiyong Tong, Mathieu Desbrun, and Peter Schröder. 2006. Edge subdivision schemes and the construction of smooth vector fields. *ACM Transactions on Graphics (TOG)* 25, 3 (2006), 1041–1048.
- Rundong Zhao, Mathieu Desbrun, Guo-Wei Wei, and Yiyong Tong. 2019. 3D Hodge decompositions of edge-and face-based vector fields. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–13.
- Yongning Zhu and Robert Bridson. 2005. Animating sand as a fluid. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 965–972.