



Hidden Degrees of Freedom in Implicit Vortex Filaments

SADASHIGE ISHIDA, Institute of Science and Technology Austria, Austria

CHRIS WOJTAN, Institute of Science and Technology Austria, Austria

ALBERT CHERN, University of California, San Diego, USA




Fig. 1. The time-evolution of a highly knotted vortex filament according to the rules of inviscid fluid dynamics. Our implicit description represents filaments as the zero levelsets of space-time complex-valued functions, which automatically handles topological changes of curves without yielding singularities.

This paper presents a new representation of curve dynamics, with applications to vortex filaments in fluid dynamics. Instead of representing these filaments with explicit curve geometry and Lagrangian equations of motion, we represent curves *implicitly* with a new co-dimensional 2 level set description. Our implicit representation admits several redundant mathematical degrees of freedom in both the configuration and the dynamics of the curves, which can be tailored specifically to improve numerical robustness, in contrast to naive approaches for implicit curve dynamics that suffer from overwhelming numerical stability problems. Furthermore, we note how these hidden degrees of freedom perfectly map to a *Clebsch representation* in fluid dynamics. Motivated by these observations, we introduce *untwisted* level set functions and *non-swirling* dynamics which successfully regularize sources of numerical instability, particularly in the twisting modes around curve filaments. A consequence is a novel simulation method which produces stable dynamics for large numbers of interacting vortex filaments and effortlessly handles topological changes and re-connection events.

CCS Concepts: • **Computing methodologies** → *Geometric topology; Physical simulation*; • **Mathematics of computing** → *Partial differential equations; Differential calculus; Discretization*.

Additional Key Words and Phrases: curve dynamics, implicit representation, Clebsch variables, fluid dynamics, vortex filaments

ACM Reference Format:

Sadashige Ishida, Chris Wojtan, and Albert Chern. 2022. Hidden Degrees of Freedom in Implicit Vortex Filaments. *ACM Trans. Graph.* 41, 6, Article 241 (December 2022), 14 pages. <https://doi.org/10.1145/3550454.3555459>

Authors' addresses: Sadashige Ishida, Institute of Science and Technology Austria, Austria, sadashige.ishida@ista.ac.at; Chris Wojtan, Institute of Science and Technology Austria, Austria, wojtan@ista.ac.at; Albert Chern, University of California, San Diego, USA, alchern@eng.ucsd.edu.



This work is licensed under a Creative Commons Attribution International 4.0 License.
© 2022 Copyright held by the owner/author(s).
0730-0301/2022/12-ART241
<https://doi.org/10.1145/3550454.3555459>

1 INTRODUCTION

The deformation of space curves is an interesting topic in many subjects such as differential geometry, low-dimensional topology, classical and quantum fluid mechanics, and electromagnetism. One example from fluid mechanics is the dynamics of vortex filaments. In a nearly inviscid fluid, vorticity originates from *codimension-1* interfaces or obstacle surfaces. The vortex sheets subsequently roll up into *codimension-2* vortex filaments, due to the Kelvin-Helmholtz instability. Hence, most physical inviscid fluids have their vorticity concentrated into a sparse set of space curves, rather than distributed evenly throughout space. Based on this observation, certain physical equations model fluids only with dynamically deforming space curves. Many fluid simulation methods take advantage of this sparsity structure.

One major challenge for an explicit (Lagrangian) filament-based fluid solver is to handle reconnection events when filaments collide. Without any reconnection, the total length of filaments can grow exponentially, exploding the computational cost and halting the solver. Hence, existing explicit filament simulators include a tedious process of collision detection followed by non-differentiable heuristic geometry surgeries.

To that end, implicit (Eulerian) curve representations are more appealing. The recently emerging *Clebsch representation* expresses vortex lines as level sets of a 2-dimensional-valued function called *Clebsch variables* [Clebsch 1859]. Like any level set method, topological changes of level set geometries occur gracefully. The difficulty, however, in a Clebsch-based fluid solver is in the dynamics of the Clebsch variables. The Clebsch variables satisfy the transport equation advected with the fluid velocity, which unfortunately behaves in a swirling motion with a high-spatial frequency and singularities near the vortex filaments. Such a rough transporting vector field is hard to resolve accurately in a computational grid. Even if the transport equation is computed accurately, the level set function will quickly evolve into a twisted and distorted function that is difficult to deal with.

This paper develops a new approach for describing the geometry and dynamics of filaments with implicit curve functions. Our main insight is that the problem has a huge number of redundant degrees of freedom: both the velocity field and the level set function (i.e., the Clebsch representation) can be varied in ways that do not change the solution. We exploit these additional degrees of freedom to ensure stable numerical simulation and automatic handling of topological changes, without sacrificing accuracy. In particular, we choose an *untwisted* Clebsch representation for the level set geometry, and *non-swirling* dynamics for advecting vortex filaments. We regularize these functions by identifying and constraining hidden degrees of freedom in their representations, allowing us to greatly improve numerical robustness compared to naive implementations.

Our algorithm is the first method for animating implicit vortex filament geometry with automatic topological changes. Our mathematical formulation offers new tools for future research on fluid simulation and curve geometry processing, and our results show greatly improved stability compared to a standard level set implementation, with fewer user parameters than for explicit Lagrangian filament techniques.

2 RELATED WORK

A majority of numerical schemes for the evolution of high-codimension geometries are developed in vortex methods in fluid animation and computational fluid dynamics. These vortex-capturing schemes seek a representation for the vorticity and solve their governing equations of motion.

Explicit vortex methods. In previous vortex methods, vortices are represented either as particles [Gamito et al. 1995; Park and Kim 2005; Selle et al. 2005; Zhang and Bridson 2014; Angelidis 2017], filaments [Cottet et al. 2000; Angelidis and Neyret 2005; Weißmann and Pinkall 2009; Weißmann and Pinkall 2010; Padilla et al. 2019], segments [Chorin 1990; Xiong et al. 2021], sheets [Brochu et al. 2012; Pfaff et al. 2012; Da et al. 2015] or volumes [Elcott et al. 2007; Zhang et al. 2015]. Vortex particle methods represent vortices as a disconnected point cloud. However, the strength of vortex per particle or per unit volume [Zhang et al. 2015] undergoes a numerically unstable *vortex stretching*, requiring an artificial clamping or diffusion that sacrifices accuracy. The stretching problem is avoided by representing vorticity per filament, segment, or sheet, or per unit area using differential 2-forms [Elcott et al. 2007]. However, describing vortex explicitly (Lagrangian method) with filaments, segments and sheets comes with a cost of sophisticated and heuristic treatment for changes of vortex topology. Volumetric (Eulerian) methods [Elcott et al. 2007] do not require managing topological changes, but they do not have a handle on codimensional structures. Our codimension-2 level set method is an Eulerian method that can represent sharp filament structures without any additional difficulty from vortex reconnection.

Clebsch representations. Another Eulerian representation of vorticity is to describe vortex lines as the level sets of a 2D-valued function known as the Clebsch variable [Clebsch 1859; Lamb 1895]. The representation was not widely adopted since all $\mathbb{R}^2$ -valued Clebsch variables can only describe fluids with zero *helicity* [Chern et al.

2017]. The helicity problem is solved by using a sphere-valued Clebsch variable [Kuznetsov and Mikhailov 1980] which can represent nonzero (though quantized) values of helicity. Since the recognition of these “spherical” Clebsch maps, they have become an established method for vortex representation in computer graphics [Chern et al. 2016, 2017; Yang et al. 2021]. However, Clebsch variables represent a *smooth vorticity field* by the continuum of level sets of all values. It remains a challenge to represent a *sharp codimension-2 filament*, especially with a limited grid resolution. By contrast, our method represents a sharp filament just as the *zero set* of a complex-valued function. Moreover, we show in Section 4.1 that our representation can actually resolve continuous values of helicity (as opposed to the quantized values mentioned earlier).

Dynamics of Clebsch variables. In addition to the implicit representation of vorticity, Clebsch variables also play significant roles in a variational and Hamiltonian formulation for the incompressible Euler equation [Clebsch 1859; Lamb 1895; Morrison 1998; Chern 2017]. In short, one (of many possible) governing equation(s) for the Clebsch variable is exceedingly simple: the Clebsch variable is advected by the fluid velocity. This equation of motion is recently adopted by Yang et al. [2021]. However, referred to as the *Lagrangian chaos*, a direct transportation by the fluid velocity quickly stirs and twists any variable to a distorted one unresolved by the finite computational grid [Qu et al. 2019]. The Clebsch variable is no exception under such dynamics. The method of *Schrödinger’s Smoke* [Chern et al. 2016] bypassed the Lagrangian chaos: its total energy (Hamiltonian) includes the Dirichlet energy of the Clebsch variable, which is therefore bounded for all time. However, while the dynamics of *Schrödinger’s Smoke* appear to be similar to that of Euler’s equation, it is only an approximation to the Euler fluid. There is still a large degree of freedom in the Clebsch representation and its dynamical system. Finding an equation of motion for the Clebsch variable that both describes the correct Euler fluid without Lagrangian chaos is an unexplored research topic. Our paper describes an instance of a Lagrangian-chaos-free dynamical system for an implicit representation of vortex filaments.

Implicit filament representations. We represent filaments as the zero set of a complex-valued function. These zero sets of a complex phase field are widely studied in condensed matter physics as *topological defects* appearing in superfluids and superconductors [Bethuel et al. 1994; Pismen et al. 1999]. These topological defect models also facilitate singularity placements in flow analysis and geometry processing [Weißmann et al. 2014; Solomon et al. 2017; Palmer et al. 2020]. Complex phase field models are taken more generally as high-codimensional level set representations by [Ambrosio and Sonner 1996; Ruuth et al. 2001; Burchard et al. 2001; Min 2004]. However, in the physics and geometry processing literature, the level set functions have specific physical meanings such as the phase of a wave, leaving little room for a smoother representative. In most cases, the phase field has norm 1 except for a sudden dip to 0 near the filaments, creating a configuration that is difficult to resolve efficiently on a computational grid. In the level set method literature, the norm-1 condition is often adopted for (re)initializing the level set functions despite the discontinuity [Ruuth et al. 2001]. The complex phase is constructed locally with little discussion about

global topological obstruction. Burchard et al. [2001] addressed the challenges in reinitializing the multi-component level set functions; by mimicking the codimension-1 signed distance functions, they propose a sophisticated reinitialization by solving a “manifold eikonal equation” along the isosurface of each function component. Unfortunately the process will not resolve the twists of the framed curve. To our knowledge, there has not been a thorough discussion about the degrees of freedom in the implicit filament representations or in their dynamics until now.

In our work, we explore the degrees of freedom of both the codimension-2 level set functions and their equations of motion. We further provide a simple reinitialization method comparable to the codimension-1 signed distance function. Comparisons show that exploiting these degrees of freedom are essential to a robust simulation.

3 REPRESENTATIONS FOR EVOLVING CURVES

We now begin our description of implicit filament dynamics. The main mathematical object is a union of closed space curves. This section describes an implicit representation for these curves and their dynamics, as well as the degrees of freedom in the representation.

3.1 Representations for curve configuration

Let the physical domain be an open region $M \subset \mathbb{R}^3$. We use γ to represent a collection of m closed curves

$$\gamma: \bigsqcup_{i=1}^m \mathbb{S}^1 \rightarrow M, \quad (1)$$

where $\bigsqcup$ denotes a disjoint union, $\mathbb{S}^1$ is the topological circle, m is the number of filaments, and γ is the mapping from the 1D curves into 3D. The configuration space $\mathcal{F}$ of these filaments is the space of all possible placements of these curves:

$$\mathcal{F} = \bigsqcup_{m=1}^{\infty} \left\{ \gamma: \bigsqcup_{i=1}^m \mathbb{S}^1 \rightarrow M \right\} / \text{reparametrizations}. \quad (2)$$

Although it is straightforward to represent curves *explicitly* via objects γ in (2) as parameterized curves or their discrete counterparts, topological changes of curves such as splitting or merging are difficult to describe mathematically and algorithmically. For instance, the number m of components can change when curves reconnect or split apart.

Instead of relying on an explicit curve representation, our work adopts an *implicit* representation for the elements in $\mathcal{F}$. We model every collection of closed curves in M as the zero set of a complex-valued level set function $\psi: M \rightarrow \mathbb{C}$:

$$\gamma = \{p \in M \mid \psi(p) = 0\} = \{\text{Re } \psi = 0\} \cap \{\text{Im } \psi = 0\}, \quad (3)$$

In other words, an alternative definition for γ is the set of all points p where *both* the real and imaginary components of a level set function ψ evaluate to zero as in Figure 2. We find it useful to draw an analogy to the scalar-valued level sets commonly used in computer graphics applications: the zero level-set of a scalar-valued function represents shapes of codimension 1 (*a.k.a.* surfaces), while our complex-valued level set has twice as many variables and thus represents shapes of codimension 2 (curves).

We next note that different functions ψ can represent the same collection of curves if they share the same zeros. To formulate this




Fig. 2. An example of level surfaces of two linked curves. The blue and red surfaces are cross sections of $\{\text{Im } \psi = 0\}$ and $\{\text{Re } \psi = 0\}$ respectively.

redundancy precisely, we define the following equivalence relation $\sim$ on the function space $(M \rightarrow \mathbb{C})$. We say $\psi_1 \sim \psi_2$ if and only if $\psi_1 = \varphi \psi_2$ for some nowhere-vanishing function $\varphi: M \rightarrow \mathbb{C} \setminus \{0\}$. This means $\psi_1 \sim \psi_2$ if and only if they share the same zero set.

For the implicit representation to work properly, we make a regularity assumption that the level set functions ψ are sufficiently “nice” around the zero set. Precisely, we assume ψ is smooth, and the differential $d\psi|_p: T_p M \rightarrow \mathbb{C} \cong \mathbb{R}^2$ is surjective, i.e. its matrix form has rank 2 at any point p on the zero-set.¹

To summarize, our configuration space of filaments $\mathcal{F}$ from (2) is replaced by the space of complex-valued functions ψ modulo a multiple of a non-vanishing function:

$$\mathcal{G} := \{\psi: M \rightarrow \mathbb{C}\} / \sim. \quad (4)$$

While $\mathcal{F}$ and $\mathcal{G}$ describe the same configuration space of filaments ($\mathcal{F} \cong \mathcal{G}$), the objects in $\mathcal{G}$ are much more continuous compared to the disjoint spaces of $\mathcal{F}$.

Relation to Clebsch representations. The representation of codimension-2 curve geometries in 3D is known in fluid dynamics as *Clebsch representations* [Clebsch 1859; Lamb 1895; Chern et al. 2017; Yang et al. 2021]. For a fluid flow with a smooth vorticity field, the vortices are geometrically depicted as fibrous vortex lines diffusely distributed over the fluid domain. A Clebsch representation aims at an implicit representation for such fibrous structure. The representation uses a map $s: M \rightarrow \Sigma$ from the 3D fluid domain M to a 2-dimensional manifold Σ with a measure σ to describe the *vortex lines* as preimages $s^{-1}\{p\}$ of points $p \in \Sigma$, and the density of the vortex lines as the pullback $s^*(\sigma)$ of the measure σ . A smooth Clebsch map s and a smooth measure σ yields a smooth distribution of vortex lines. To achieve a more singular and concentrated vorticity field such as vortex filaments, one would consider s with larger derivatives [Marsden and Weinstein 1983] (s sweeps out more measure σ over a small area in M).

In our setup, we want to represent singular curves with a Dirac- δ density, instead of diffused distribution of vortex lines. Previous

¹The full-rank requirement does not modify the degrees of freedom in the rescaling φ since $d(\varphi\psi)|_p = d\varphi|_p \underbrace{\psi_p}_{=0} + \varphi_p d\psi|_p = \varphi_p d\psi|_p$ and a non-zero complex multiplication by φ_p always preserves the rank of $d\psi|_p$.

considerations in Clebsch representations would set the Clebsch map s with enormous derivative. By contrast, we obtain such concentrated filaments by setting σ singular while keeping the Clebsch map smooth. Our complex level function $\psi: M \rightarrow \mathbb{C}$ is a Clebsch map with the target space $\Sigma = \mathbb{C}$ equipped with a δ -measure $\sigma = \delta_0$ at the origin.

An important discussion about Clebsch representations [Chern et al. 2017] is whether or not a fluid configuration can be represented with the choice of Σ and σ . Previous Clebsch representations [Chern et al. 2017; Yang et al. 2021] adopt $\Sigma = \mathbb{S}^2$, since the more straightforward choice of $\Sigma = \mathbb{R}^2 \cong \mathbb{C}$ with the standard area measure σ can only represent fluid flows with zero *helicity*. The helicity obstruction is reduced for $\Sigma = \mathbb{S}^2$ as it can admit a discrete set of nonzero helicity [Chern et al. 2017]. Our Clebsch representation can represent any space curve without any obstruction. In particular, the helicity of a vortex filament is proportional to its *writhe* [Arnold and Khesin 1998] which can take any real value.

3.2 Representations for curve dynamics

In the explicit representation, a first-order time evolution of curves $\gamma: (\sqcup^m \mathbb{S}^1) \times \mathbb{R} \rightarrow M$ can be described by an equation of the form:

$$\frac{\partial \gamma}{\partial t}(s, t) = \mathbf{V}_{\gamma_t}(s), \quad s \in \sqcup \mathbb{S}^1, t \in \mathbb{R}. \quad (5)$$

Here, s is the parameterization of the curve, t is time, and the velocity $\mathbf{V}_{(\cdot)}: \mathcal{F} \rightarrow (\sqcup \mathbb{S}^1 \rightarrow \mathbb{R}^3)$ is a dynamical model that tells the filament how to move based on the current filament shape and position.

Example 1. In the context of fluid dynamics, important examples for the velocity model $\mathbf{V}$ are the ones that govern the motion of vortex filaments. When $M = \mathbb{R}^3$, i.e. there are no obstacles or boundaries, the velocity models are the Biot–Savart model

$$\mathbf{V}_Y^{\text{BS}}(s) := \frac{\Gamma}{4\pi} \oint \gamma'(\tilde{s}) \times \frac{\gamma(s) - \gamma(\tilde{s})}{|\gamma(s) - \gamma(\tilde{s})|^3} d\tilde{s} \quad (6)$$

and the more regular Rosenhead–Moore model [Saffman 1992, pp. 213]

$$\mathbf{V}_Y^{\text{RM}}(s) := \frac{\Gamma}{4\pi} \oint \gamma'(\tilde{s}) \times \frac{\gamma(s) - \gamma(\tilde{s})}{\sqrt{e^{-3/2}a^2 + |\gamma(s) - \gamma(\tilde{s})|^2}^3} d\tilde{s} \quad (7)$$

where the constants Γ and a are the vortex strength and vortex thickness respectively, and the integrating measure $d\tilde{s}$ is the arclength element (set $a = 0$ for the Biot–Savart model). Note that (6) and (7) are the restrictions at the curve of the entire fluid velocity field over the 3D domain

$$\mathbf{U}_Y^{\text{RM}}(\mathbf{x}) := \frac{\Gamma}{4\pi} \oint \frac{\gamma'(\tilde{s}) \times (\mathbf{x} - \gamma(\tilde{s}))}{\sqrt{e^{-3/2}a^2 + |\mathbf{x} - \gamma(\tilde{s})|^2}^3} d\tilde{s}, \quad \mathbf{x} \in \mathbb{R}^3 \quad (8)$$

That is, $\mathbf{V}_Y^{\text{RM}}(s) = \mathbf{U}_Y^{\text{RM}}(\gamma(s))$.

Now, we translate the dynamical system (5) into an evolution equation for a time-dependent complex level function ψ . First, we note that the evolution of ψ around the zeros is given as the *transport equation* along some vector field, which is formally the following lemma.

Lemma 1. For any time-dependent complex level set function ψ with the regularity assumptions in Section 3.2, there exists a neighborhood $U \subset M$ of the zero set of ψ and a vector field $\mathbf{v}: U \rightarrow \mathbb{R}^3$ such that

$$\frac{\partial \psi}{\partial t} + \mathbf{v} \cdot \nabla \psi = 0 \quad \text{in } U. \quad (9)$$

PROOF. By the regularity assumptions in Section 3.2, there exists a neighborhood U of the zero set of ψ where the 3D-to-2D linear map $d\psi$ is full-rank and thus surjective. Hence at every point $\mathbf{x} \in U$ and for any value of $\partial\psi/\partial t(\mathbf{x})$, there exists a vector $\mathbf{v}_{\mathbf{x}} \in \mathbb{R}^3$ such that $d\psi|_{\mathbf{x}}(\mathbf{v}_{\mathbf{x}}) = -\partial\psi/\partial t(\mathbf{x})$. With this construction, we obtain a vector field $\mathbf{v}: U \rightarrow \mathbb{R}^3$ satisfying $\partial\psi/\partial t + \mathbf{v} \cdot \nabla \psi = 0$ in U . $\square$

Hence the representation of the dynamics for ψ can be encapsulated into a *vector field*. Observe that the zero level curve γ of ψ evolves by $\partial\gamma/\partial t = \mathbf{v}|_{\gamma}$ under the transport equation (9), and that $\ker(d\psi)$ at each point on γ is spanned by the tangent γ' as γ' lies on the tangent spaces of both $\{\text{Im}\psi = 0\}$ and $\{\text{Re}\psi = 0\}$. By matching these induced curve dynamics in Lemma 1 with (5) we conclude:

Theorem 1. The zero level curve γ of ψ evolves according to (5) if and only if ψ satisfies (9) for some vector field $\mathbf{v}$ that agrees with the curve velocity at the curve:

$$\mathbf{v}(\gamma(s)) = \mathbf{V}_{\gamma}(s) + f(s)\gamma'(s) \quad (10)$$

where $f\gamma'$ is the tangent vector γ' multiplied by an arbitrary scalar function f . The degrees of freedom of the dynamics for ψ are the degrees of freedom for choosing $\mathbf{v}$ with the condition (10).

Essentially, the only velocities that really matter for the evolution of the curves are the velocities located on the zero level set. Furthermore, the locations of the curves will not change if we slide them around their tangent direction (like spinning a circle around its axis of symmetry), so we only need to pin down their *normal and bi-normal components*. So we have a huge number of velocity variables (3 for each point in the 3D domain) with very few constraints (2 for each point on the 1D curves). This under-determined system gives us a *redundancy* in possible velocity fields for curve dynamics, which is largely unexplored by previous work.

Let us apply Theorem 1 to vortex filament dynamics (cf. Example 1): Plugging in (8) for $\mathbf{v}$ gives us the following dynamics for ψ :

$$\frac{\partial}{\partial t}\psi + \mathbf{v} \cdot \nabla \psi = 0, \quad \mathbf{v}(\mathbf{x}) = \mathbf{U}_Y^{\text{RM}}(\mathbf{x}). \quad (11)$$

This is the most straightforward way to do it: simply advect the level set ψ in the exact same way as the rest of the fluid. However, we know that $\mathbf{v} = \mathbf{U}_Y^{\text{RM}}$ is an extremely sensitive function to deal with numerically – it tends to infinity near γ , has unbounded derivatives, and rapidly changes direction in very tight swirls. Small errors in γ inevitably create huge errors in velocity, making simulations unstable, as demonstrated in our accompanying video. Fortunately, according to Theorem 1 we now know that there are infinitely many velocity fields that will all theoretically give us the same filament motions; our mission in the next section is to swap out this unstable Biot–Savart velocity field for one that is much more numerically robust.

4 UNTWISTED CLEBSCH VARIABLES AND NON-SWIRLING DYNAMICS

4.1 Untwisted Clebsch variables

Like the common codimension-1 real-valued level set methods, the implicit representation benefits from the regularity of the level set function. There, a level set function is well-conditioned if the magnitude of the gradient is close to one. For that reason, the level set function is typically initialized as the *signed distance function*, and this property is typically maintained as the level set is evolved (called *re-distancing*).

For our codimension-2 complex level set representation, we shall also characterize a set of desirable qualities of the complex level set function $\psi: M \rightarrow \mathbb{C}$. Due to the higher codimension, the discussion involves the notion of *twist* from the mathematical *ribbon theory*. Finally, we describe a concrete construction of ψ that will be used for initialization and re-distancing.

4.1.1 Conditioning of a complex level set function. We want ψ to be continuous everywhere and non-zero outside the curves $\gamma = \{\psi = 0\}$. Near the curves γ , we want the differential $d\psi$ to be well-conditioned: if we only consider the function restricted to the plane spanned by the curve normal and bi-normal, $d\psi|_{\gamma^\perp}$ is close to an isometry; that is $d\psi|_{\gamma^\perp}: \gamma^\perp \rightarrow \mathbb{C} \cong \mathbb{R}^2$ has singular values $\approx (1, 1)$. Notice, however, that even when $d\psi|_{\gamma^\perp}$ is well-behaved on each normal plane, the level set function ψ can still exhibit significant *shearing* if the complex phase varies significantly along the tangent direction.

This variation of complex phase along the curve's tangent can be seen more intuitively via the geometry of the surface S_ψ formed by the level sets of the positive real part of ψ ($S_\psi = \{\text{Re}(\psi) > 0, \text{Im}(\psi) = 0\} = \{\arg(\psi) = 0\}$). Note that the boundary of the surface S_ψ is γ , as illustrated in Figure 3. The tangential variation of the phase $\arg(\psi)$ is embodied by the twist of S_ψ at its boundary γ .

For the specific method presented in this paper, we take a simple construction of ψ that is known to have a smooth S_ψ with little twisting at γ . Readers who are only concerned about the method can skip ahead to Section 4.1.3; for researchers interested in extending this work, we will now lay out a few geometric and topological properties about the twist.

4.1.2 The twist of a complex level set function. To discuss the twist precisely, we consider the normal vector U_ψ of the curve γ that is tangent to the surface S_ψ .

Definition 1 (ψ -induced framing). *Each complex level set function $\psi: M \rightarrow \mathbb{C}$ for γ gives rise to a normal vector field $U_\psi: \gamma \rightarrow \mathbb{R}^3$, $|U_\psi| = 1$, $\langle \gamma', U_\psi \rangle = 0$ that points to the direction where ψ is real and positive. Explicitly, at each point on γ ,*

$$U_\psi := \frac{(d\psi|_{\gamma^\perp})^{-1}(1)}{|(d\psi|_{\gamma^\perp})^{-1}(1)|}. \quad (12)$$

This normal vector field U_ψ makes γ a *framed curve*.

Definition 2 (Twist). *The twist of ψ at each point on γ is given by*

$$\omega := U'_\psi \cdot (\gamma' \times U_\psi). \quad (13)$$

Here, $(\cdot)'$ is the derivative along curve's tangent relative to arclength.

When $d\psi|_{\gamma^\perp}: \gamma^\perp \rightarrow \mathbb{C}$ is conformal (which is true in our case introduced in Section 4.1.3 since it is isometric), the twist directly relates to the derivative of ψ along the curve:

$$\omega = -(d\theta)(\gamma') = -\text{Re} \left(\frac{(d\psi)(\gamma')}{i\psi} \right), \quad \theta = \arg \psi, \quad (14)$$

in a small neighborhood of γ .



Fig. 3. The surface S_ψ (translucent blue) and the ribbon U_ψ (opaque cyan).

Now, a natural consideration for designing ψ is to minimize the twist. For example, one may try to construct *Bishop's parallel frame* [Bergou et al. 2008], which has no twist. However, one would find the construction *impossible* for a closed curve and a union of closed curves in general. This is due to the following theorem that the *total twist*

$$\text{Tw}(\psi) := \frac{1}{2\pi} \oint_\gamma \omega \quad (15)$$

is fixed by the curve geometry $\gamma \in \mathcal{F}$ or the equivalent class of level functions $[\psi] \in \mathcal{G}$ and is generally non-zero.

Theorem 2 (Invariance of total twist). *Suppose M is simply connected. Let $[\psi] \in \mathcal{G}$. Then any representative $\psi \in [\psi]$ has the same total twist $\text{Tw}(\psi)$.*

PROOF. In the special case that γ is a single connected curve in $\mathbb{R}^3$, the statement is the result of the Călugăreanu Theorem: The total twist of the frame $\text{Tw}(\psi)$ and the total writhe of the curve $\text{Wr}(\gamma)$ must add up to the linking number between the two boundaries of the ribbon swept out by the frame U_ψ . Here, we note that any closed curve admits a Seifert surface i.e. a compact, connected, and oriented surface spanning the curve [Seifert 1935; Murasugi 1996]. Without loss of generality, we can assume a constant phase-shift to ψ so that S_ψ is a Seifert surface of γ . Since the ribbon U_ψ lies in a Seifert surface, the linking must be zero. Therefore, $\text{Tw}(\psi) = -\text{Wr}(\gamma)$, which only depends on γ . See Figure 3 for a visualization of S_ψ and U_ψ . For the definitions of quantities for curves such as writhe, twist, or linking numbers, we refer to [Arnold and Khesin 1998].

In the general case, observe that under any change of representative $\tilde{\psi} = \varphi\psi$ in $[\psi]$ with a nowhere-vanishing $\varphi: M \rightarrow \mathbb{C} \setminus \{0\}$, the twist transforms according to $\tilde{\omega} = \omega - d_{\gamma'}(\arg(\varphi))$. Since φ has no zero in a simply connected M , the angle $\arg(\varphi)$ can be defined as a real-valued function globally over M rather than a 2π -ambiguous angle-valued function. In particular, Stokes' theorem applies and the integral along the closed curve vanishes: $\oint_\gamma d(\arg(\varphi)) = 0$. So $\oint_\gamma \tilde{\omega} = \oint_\gamma \omega$. $\square$




Fig. 4. The solid-angle distance function ψ for a space curve γ is constructed by the distance and the angle subtended by the curve.

Note that even though the total twist is fixed, strong *local* twist can still be present.

4.1.3 Solid-angle distance function. As explained above, we desire a complex level set function ψ that is numerically well-conditioned: it should be nearly isometric near the curve γ and non-zero outside γ . In order to isolate these properties and enforce them explicitly, we model ψ with a complex wave function:

$$\psi(\mathbf{x}) = r(\mathbf{x})e^{i\theta(\mathbf{x})}. \quad (16)$$

Similar to *signed distance functions* for the codimension-1 level set functions, we set

$$|\psi(\mathbf{x})| = r(\mathbf{x}) := \text{dist}(\gamma, \mathbf{x}), \quad \mathbf{x} \in M \subset \mathbb{R}^3. \quad (17)$$

This setup ensures that the value $\psi(\mathbf{x})$ is non-zero for $\mathbf{x} \notin \gamma$. What remains is a choice for the complex phase $\theta: M \setminus \gamma \rightarrow \mathbb{S}^1 = \mathbb{R}_{\text{mod } 2\pi}$, which we set to the *half solid-angle subtended* by γ :

$$\theta(\mathbf{x}) := \frac{1}{2} \text{SolidAngle}(\gamma; \mathbf{x}) \pmod{2\pi} \quad (18)$$

where $\text{SolidAngle}(\gamma; \mathbf{x}) \in \mathbb{R}_{\text{mod } 4\pi}$ is a dimensionless quantity given by the signed spherical area² enclosed by the projection of γ on the unit sphere centered at $\mathbf{x}$ i.e. $\text{Proj}_{\gamma\mathbf{x}}(s) := (\gamma(s) - \mathbf{x})/|\gamma(s) - \mathbf{x}|$. We call this construct of ψ the *solid-angle distance function* (Figure 4).

The solid-angle distance function meets our desired conditions for the codimension-2 level set representation. On each normal plane $\theta|_{\gamma^\perp}$ is asymptotically the 2D angle function about the zero γ , so that $d\psi|_{\gamma^\perp}: \gamma^\perp \rightarrow \mathbb{C}$ is close to an isometry. Moreover, as studied by [Binysh and Alexander 2018], the surface $S_\psi = \{\theta = 0\}$ features little twist at γ . Figure 5 illustrates this concept.

We note that our choice of the Clebsch variables ψ is an instance of many reasonable ones that exploit degrees of freedom in curve representations rather than the optimal one for a specific dynamics such as vortex filaments. Nevertheless, our ψ has a number of desirable properties, which we list in Appendix A.

4.2 Non-swirling dynamics

In Section 4.1, we leveraged the degrees of freedom in the complex level set function ψ to design a sufficiently regular implicit representation. Here, we exploit similar degrees of freedom to construct

²Our definition of the solid angle is modulated by 4π as the spherical polygon of projected curves may cover the sphere multiple times.




Fig. 5. Plotting θ on a 2D plane which intersects a circular vortex ring at two points (left). The color indicates the value of the θ , and the black lines are its level curves. The curves meet where the filament intersects the plane. Zooming into the white box (right) shows evenly-spaced curves closer to the filament, where $e^{i\theta}$ resembles the complex plane $\mathbb{C}$.

evolution equations that produce theoretically equivalent dynamics but are more numerically robust.

Many dynamical systems for curves already come with a known physical evolution equation. For example, the vortex filament dynamics can be simulated with (11), i.e. by advecting ψ using the Biot–Savart or the Rosenhead–Moore flow (Example 1). Hence, redesigning the equation may seem unnecessary. However, when it comes to numerically advancing the variables, the highly oscillatory or discontinuous nature of the Biot–Savart and the Rosenhead–Moore flows near the vortex core (as illustrated in Figure 6) can cause significant interpolation error. We point out that these errors are avoidable by redesigning the flow of the advection.

We consider dynamical systems as discussed in Section 3.2. Suppose the evolution of the curve is given by $\frac{\partial \gamma}{\partial t} = \mathbf{V}_\gamma$ where $\mathbf{V}_\gamma$ is the velocity field defined on the curve. The evolution for ψ must be an advection by an extension $\mathbf{v}$ of the velocity field $\mathbf{V}_\gamma$ in a neighborhood of the curve (cf. (9)). A straightforward construction of $\mathbf{v}$ is a constant extrapolation. That is, $\mathbf{v}(\mathbf{x})$ is set to $\mathbf{V}_\gamma$ at the closest point on γ from $\mathbf{x}$. This extrapolation is, however, singular where closest points are not unique as in the middle row of Figure 6.

To gain continuity without changing the velocity on the filaments, we smooth away these singularities by taking the weighted average of the filament velocity as

$$\mathbf{v}(\mathbf{x}) := \frac{1}{\mathcal{N}(\mathbf{x})} \oint_\gamma \mathbf{V}_\gamma(\gamma(s)) w(\mathbf{x}, \gamma(s)) ds \quad (19)$$

Here, $\mathcal{N}(\mathbf{x})$ is the normalization factor

$$\mathcal{N}(\mathbf{x}) = \oint_\gamma w(\mathbf{x}, \gamma(s)) ds, \quad (20)$$

and w is some weight function that applies less smoothing as it gets closer to the filament, i.e. $w(\mathbf{x}, \gamma(s))/\mathcal{N}(\mathbf{x}) \rightarrow \delta(\gamma^{-1}(\mathbf{x}) - s)$ as $\text{dist}(\mathbf{x}, \gamma) \rightarrow 0$. For example, we observed that a Gaussian function with distance-dependent variance works stably:

$$w(\mathbf{x}, s) = \exp\left(-\frac{|\mathbf{x} - \gamma(s)|^2}{\sigma^2 \text{dist}(\mathbf{x}, \gamma)}\right), \quad \sigma \text{ some constant.} \quad (21)$$

To accelerate computation for $\mathbf{v}$, we can further multiply the integrand of Equation 19 by a smooth cutoff function which equals to 1 near γ and 0 far away from γ . Then $\mathbf{v}$ is non-vanishing only




Fig. 6. Different vector fields for two linked rings. The original Rosenhead–Moore model (top), the nearest point velocity field (middle), and a smooth weighted average field (bottom). While these three velocity fields coincide on the filaments, they differ significantly outside the filaments.

near γ . By applying this smooth cutoff, we only need to evaluate $\mathbf{v}$ in a narrow band close to the filament.

Note that $\mathbf{v}$ is in general not divergence-free. The velocity near the curve is determined according to the curve velocities so that the motion of the zeros of ψ emulates the motion of the curves. Imposing an additional constraint like incompressibility to the velocity field may trade off the fidelity to the original curve motion or the numerical smoothness of the surrounding vector field.

5 ALGORITHMS

In this section, we describe an algorithm for simulating filament dynamics.

Throughout the simulation, we maintain a complex level set function ψ . The main algorithm computes the transport equation of ψ along a velocity field $\mathbf{v}$ in a neighborhood of the zeros of ψ . This main algorithm is accompanied by a few subroutines for evaluating the velocity and redistancing: one subroutine extracts the zero set γ of ψ ; another subroutine constructs the solid-angle distance function ψ from γ (Section 4.1.3); a third subroutine evaluates the filament motion $\mathbf{V}_\gamma$ using γ ; and the last subroutine extends $\mathbf{V}_\gamma$ to a velocity field $\mathbf{v}$ in a neighborhood of γ (Section 4.2).

We store the level set function ψ and the velocity $\mathbf{v}$ on a 3D lattice and discretize the curves γ as oriented collections of line segments.

Algorithm 1 The main time integration

Input: Initial filament $\gamma \in \mathcal{F}$;
 1: $\psi \leftarrow$ construct ψ from γ ; ▷ Section 5.1.3
 2: **while** simulating **do**
 3: $\mathbf{v} \leftarrow$ evaluate $\mathbf{V}_\gamma$ and extend it on grids near γ ;
 4: $\psi \leftarrow$ advect ψ along $\mathbf{v}$; ▷ Section 5.1.1
 5: $\gamma \leftarrow$ extract the zero set of ψ ; ▷ Algorithm 2
 6: $\psi \leftarrow$ construct ψ from γ ; ▷ redistance; Section 5.1.3
 7: **end while**

Algorithm 2 Extract the zero curve γ from ψ

1: **for each** cell c **do**
 2: **for each** face f in c **do**
 3: Compute incidence $n_f \in \{-1, 0, 1\}$; ▷ Eq. (22)
 4: **if** $n_f = \pm 1$ **then**
 5: Find $\mathbf{p}_f^\pm = \psi^{-1}(0) \in \mathbb{R}^3$ in f ; ▷ bilinear interp.
 6: **end if**
 7: **end for**
 8: Connect $\mathbf{p}_f^-$ and $\mathbf{p}_g^+$ of some faces f, g in c ;
 9: **end for**

5.1 Details of the main algorithm

5.1.1 Advection. To advect ψ with a given flow $\mathbf{v}$, one can adopt any Eulerian advection scheme. In our implementation, we use the modified MacCormack method [Selle et al. 2008] with 4th order Runge–Kutta back-tracing.

5.1.2 Construction of γ from ψ . After updating ψ , we need to update the filaments γ by extracting $\{\psi = 0\}$. We summarize this subroutine in Algorithm 2, which is adopted from [Weißmann et al. 2014]. In our setting, each vertex of γ lives on a face f of the volumetric grid. We first evaluate for each face f the $\{-1, 0, +1\}$ -valued *signed intersection* n_f with the zero curve of ψ using the *argument principle*: If the vertices of a face f are i, j, k, ℓ in an oriented order, then

$$n_f = \frac{1}{2\pi} \left(\arg\left(\frac{\psi_j}{\psi_i}\right) + \arg\left(\frac{\psi_k}{\psi_j}\right) + \arg\left(\frac{\psi_\ell}{\psi_k}\right) + \arg\left(\frac{\psi_i}{\psi_\ell}\right) \right) \quad (22)$$

using the principal branch $-\pi < \arg(\cdot) \leq \pi$. Geometrically, (22) describes how many times the quadrilateral $\psi_i, \psi_j, \psi_k, \psi_\ell \in \mathbb{C}$ winds around the origin. For each face where $n_f \neq 0$ we evaluate the more precise location of the zero using a bilinear interpolation. That is, we regard $f = [0, 1]^2$ by scaling and ψ is bilinearly interpolated as

$$\begin{aligned} \psi_f(x, y) := & (1-x)(1-y)\psi(0, 0) + x(1-y)\psi(1, 0) \\ & + (1-x)y\psi(0, 1) + xy\psi(1, 1), \end{aligned} \quad (23)$$

and $\psi_f : f \rightarrow \mathbb{C}$ has the inverse when $n_f = \pm 1$. The location $\psi_f^{-1}(0)$ is a vertex $\mathbf{p}_f$ of the curve γ . Finally, we build the edges of γ by running over the grid cells where we connect the pairs of zeros on the face with a consistent orientation. Each cube c may have up to

two pairs of zeros with positive and negative n_f . A cube with two pairs of vertices has an ambiguity similarly to the marching cube algorithm [Lorensen and Cline 1987]. We resolve this ambiguity by connecting vertices arbitrarily in a way that preserves curve orientation. We have not investigated higher order algorithms to connect curves more accurately at the sub-grid scale.

5.1.3 Construction of ψ from γ . Given γ , we construct ψ as described in Section 4.1.3. For each grid point $\mathbf{x}$ near γ , we evaluate $|\psi(\mathbf{x})| = \text{dist}(\gamma, \mathbf{x})$ as the distance to the closest polygon edge. We evaluate $\arg(\psi(\mathbf{x})) = \theta(\mathbf{x}) = \frac{1}{2} \text{SolidAngle}(\gamma; \mathbf{x})$ by computing the signed area of a spherical polygon with vertices $\{p_i\}_{i=1}^N$ that are the discrete points of γ projected onto the unit sphere centered at $\mathbf{x}$. To ensure differentiability and the local isometry properties of $d\psi|_{\gamma^\perp}$, the sign of θ must depend on curve orientation: we compute the *signed* area by introducing a pole $Z := (1, 0, 0)$:

$$\text{SignedArea}(\{p_i\}_{i=1}^N) = \sum_{i=1}^N \text{sign}((p_i \times p_{i+1}) \cdot Z) \text{Area}(p_i, p_{i+1}, Z) \quad (24)$$

where the *unsigned* area of each spherical triangle $\{q_1, q_2, q_3\}$ is computed using a standard area formula [Bevis and Cambareri 1987; Arvo 1995],

$$\text{Area}(q_1, q_2, q_3) = -\pi + \sum_{i=1}^3 \arccos\left(\frac{(q_{i-1} \times q_i) \cdot (q_i \times q_{i+1})}{\|q_{i-1} \times q_i\| \|q_i \times q_{i+1}\|}\right). \quad (25)$$

These equations assume cyclic vertex indexing, so $p_{N+1} = p_1$, $q_{3+1} = q_1$, and $q_{1-1} = q_3$.

5.1.4 Extending velocity to grid. In order to advect ψ , we need to extend the velocity field defined on γ to the grid points near the curves. To produce non-swirling dynamics, we used a smooth average field (Equation 19, 20, and 21) for the examples in this paper, unless otherwise explained. In Equation 21, we observed that σ ranging from 0.1 to 10 times the grid size works stably without smoothing out detailed dynamics.

6 APPLICATIONS

This section discusses applications of our approach, specifically applied to vortex filament dynamics. We implemented our algorithms on Houdini 18.5.759 and ran all simulations on a MacBook Pro (13-inch, 2020) with a 2.3 GHz Quad-Core Intel Core i7 processor. For an example implementation, see <https://github.com/sdsgisd/ImplicitVortexFilaments>

We use our algorithm to animate two “leapfrogging” vortex rings in Figure 7 and our accompanying video. We note that the system remains stable and highly symmetric even at the end of a long simulation with several high-speed ring interactions. Figure 8 visualizes two vortex rings colliding with one another at right angles, reconnecting, and detaching into two new rings. We note that the final rings retain plenty of energy after the collision event, in contrast to Eulerian simulations of this phenomenon which tend to damp out over time. The visual detail in our simulations is also practically independent of grid resolution, as the motion of marker particles are




Fig. 7. Two filaments leapfrog through one another, dragging marker particles into the shape of a mushroom cloud. The initial filament geometry consists of two co-planar vortex rings, one with half the radius of the other.




Fig. 8. Two smoke rings (left) colliding at orthogonal angles (middle) and re-connecting (right). Rendered as filaments (top) and marker particles (bottom). The colliding rings leave swirling trails of smoke particles after their collision and reconnection.

described analytically by Biot-Savart-style velocities (Equation 8), instead of a vector field stored on a coarse grid.

Figure 9 illustrates a jet of smoke created by generating a new smoke ring at the left side of the domain every three time steps. For transporting smoke as a scalar field stored on grid points, we again used Equation 8. This simulation shows the robustness of our topology changes: each re-connection event is the result of a curve extraction from a level set function, so there is no possibility of any unexpected edge cases, and no need for any geometric intersection code. The simple set-up creates a large variety of chaotic motions resulting from fast leap-frogging rings squeezing in between others and reconnecting filaments causing sudden changes in direction.

When rings shrink smaller than the grid resolution, our algorithm deletes them (similar to codimension-1 level set methods).

Lastly, Figure 1 illustrates how our method can evolve intricate filament geometry, specifically the (5,8)-torus knot defined by

$$\gamma(s) = ((\cos(qs) + 2) \cos(ps), (\cos(qs) + 2) \sin(ps), -\sin(qs))$$

with $(p, q) = (5, 8)$ and $s \in [0, 2\pi)$.




Fig. 9. A jet of smoke, rendered as raw filament geometry (top) and an advected smoke density function (bottom).

6.1 Influence of Numerical Parameters

Figure 10 demonstrates the importance of each step in our approach by selectively removing or modifying different algorithmic components and illustrating the consequences. First we illustrate what happens when we vary the free degrees of freedom in the velocity field $\mathbf{v}$ used to advect the filaments. In agreement with the discussion in Section 3.2, we see that setting $\mathbf{v}$ to the fluid velocity (based on Biot-Savart kernels) field causes the level set function ψ to rapidly twist up and become unstable. Setting $\mathbf{v}$ to equal the velocity at the nearest point on the filament creates similar noise, presumably due to spatial discontinuities in the field. Compare these results to the smooth geometry generated by our velocity field at the bottom of Figure 10.

Lower down in the same figure, we illustrate the effect of varying the free degrees of freedom in our level set function ψ . Starting with an initially smooth ψ and advecting it without any re-distancing or regularization works well at the beginning, but it eventually accumulates topological noise. To illustrate the impact of ψ 's free parameters on numerical stability and accuracy, the fourth row in Figure 10 replaces our smooth choice of ψ with one that is intentionally twisted by a phase shift of $\Delta\theta(x) := 0.05 \text{ dist}(x, \gamma)$; the twisted ψ causes high-frequency geometric noise and artificially shrinks




Fig. 10. Comparisons with different settings. Initial state is the 1st row of Figure 13. The left and right columns show states at frames 20 and 45.

the filaments. Again, we can compare these results to the smooth geometry generated by our un-twisted ψ at the bottom of Figure 10.

Another important numerical parameter is the spatial resolution of the filament. Lagrangian methods constrain the curve resolution by subdividing and collapsing edges when they become too long or short. In contrast, our method controls the curve details via the resolution of the grid used for the level set ψ . Figure 11 illustrates a simulation of an evolving trefoil knot on a $50 \times 50 \times 60$ grid and one twice as detailed at $100 \times 100 \times 120$. As expected, higher resolution grids create filaments with sharper details. Note, however, that the




Fig. 11. Simulation of a trefoil knot with different grid resolutions, as viewed from the side (left column) and front (right column).

common way to visualize fluids is with marker particles or smoke densities, not by visualizing the filaments themselves. Thus, even very low resolution vortex filaments can still produce high resolution visual details. More detailed filaments make themselves evident via there more detailed velocity fields and complex smoke dynamics.




Fig. 12. Comparing a buoyant trefoil knot simulation by Weißmann and Pinkall [2010] to ours. The simulations evolve from left to right.

6.2 Comparisons with Lagrangian filaments

Next, we qualitatively compare implicit and explicit representations of curve dynamics using our algorithm and the Lagrangian vortex filament technique of Weißmann and Pinkall [2010], as implemented in Houdini software by SideFX. Figure 12 shows the evolution of a knotted vortex filament with both methods. The filament is initialized as

$$\gamma(s) = (\sin(s) + 2\sin(2s), \cos(t) - 2\cos(2s), -\sin(3s))$$

with $s \in [0, 2\pi)$. Aside from some small differences arising from the particulars of how filaments break apart and reconnect, the two methods produce roughly the same dynamics.




Fig. 13. Simulating two linked vortex rings (top) with our method (2nd panel), and with the explicit filament approach of Weißmann and Pinkall [2010] (bottom two panels). Lagrangian methods can be sensitive to numerical parameters for topological changes.

On the other hand, the two methods have significantly different mechanisms for handling topological changes, which can produce divergent results. The Lagrangian method depends on user parameters like the thresholds for distances and angles between curves; the only relevant user parameter for changing topology in our method is the grid spacing, which prescribes the resolution of the level set ψ . Figure 13 shows a simulation of two linked rings: our approach both preserves long thin tendrils and filters out topological noise. The Lagrangian simulation is sensitive to re-connection parameters: setting these parameters too aggressively leads to smooth geometry but loses thin features, while setting parameters too conservatively preserves thin features but creates noisy, persistent, high-speed “ringlets” that dominate the fluid velocity field. We stress here that the purpose of this comparison is to show how explicit and implicit

descriptions handle topological changes of curves result in qualitatively different ways; the accuracy of these methods are not easily comparable as they have different mechanisms, and the accuracy depends on the type of curve dynamics.

6.3 Obstacles

Like many methods for vortex dynamics, our method can also make filaments circumvent obstacles. A typical approach is to find a smooth harmonic potential ϕ such that for a given obstacle $B \subset M$, it solves

$$\langle \mathbf{v}_\gamma - \mathbf{v}_{\partial B} - \nabla \phi, \mathbf{n} \rangle_{\mathbb{R}^3} = 0 \text{ on } \partial B \quad (26)$$

$$\phi \rightarrow 0 \text{ at infinity} \quad (27)$$

where $\mathbf{v}_{\partial B}$ is the boundary velocity at each point of ∂B . Weißmann and Pinkall [2012] construct such a potential by regarding points inside B as sources of localized potentials so their weighted sum solves Equation 26, and Brochu et al. [2012] and Zhang et al. [2014] solve a similar system using boundary element techniques. Nabizadeh et al. [2021] address an equivalent problem by solving linear PDEs on infinite domains using the Kelvin transform.

Another approach by Weißmann and Pinkall [2010] is more specialized to vortex filaments and phenomena like vortex shedding; it regards the obstacle as a collection of artificial filaments γ_M such that the normal component of the velocity is zero. Other approaches, like that of Park and Kim [2005] and Da et al. [2015] add point constraints to the boundary which zero out both the normal and tangential velocity components.

Our vortex filament algorithm is compatible with any of these obstacle-handling methods; our particular implementation uses [Nabizadeh et al. 2021], as seen in Figure 14. We can observe that filaments near the obstacle are accelerated due to the induced mirrored image of themselves in the obstacle (or vortex sheet on the surface).

Note that for the evolution of curves, only the evaluation of the velocity field needs the treatment of the boundary. Other components of the algorithm including the construction of levelset function ignore the boundary and do not require any additional treatments.

6.4 Other types of filament dynamics

Our idea of implicit representation of filament dynamics is not limited to vortex filaments. In theory, this should be applicable to any first-order time evolution of curves. Figure 15 shows an example with the curve-shortening flow,

$$\frac{\partial \gamma}{\partial t}(s, t) = -\gamma''(s, t). \quad (28)$$

As expected, this velocity causes the filament to shrink over time, similar to mean-curvature flow for surfaces [Osher and Sethian 1988].

6.5 Discussion

Our method represents a fundamental new way to animate fluids and vortex filament dynamics. Its unique features give it some strengths and weaknesses relative to existing approaches.

First of all, we represent our filaments with implicit functions ψ . While there are many ways to encode an implicit function, our




Fig. 14. Turbulence caused by a spherical obstacle. We generate a new vortex filament ring every 10 frames.



Fig. 15. Curves evolving under the curve shortening flow, from left to right.

implementation uses a regular grid, which implies a finite bounding box. While common for Eulerian fluid simulations, bounded domains are a constraint not shared by Lagrangian methods. This constraint could be mitigated if we use sparse grids or trees [Museth et al. 2019].

The main parameter in our method is the grid resolution. As discussed earlier, this parameter affects the geometric detail and topology of our filaments, influencing the velocity field directly, but influencing the final visual results only indirectly. The grid resolution also directly influences topological changes; the only mechanism for topological changes in our algorithm is to merge curves when they intersect the same grid cell. We believe this automatic and robust method for handling topological changes is a strength of our method; it minimizes the need to fine-tune parameters (especially the relationship between minimum/maximum edge length and topological change interaction lengths) and seems immune to the types of numerical blow-ups that we have seen in Lagrangian methods when filaments get close or exhibit near-degenerate geometry.

On the other hand, this grid-based method for re-connecting curves will also delete small features when they shrink below the grid cell size (similar to level set methods for surfaces). This behavior

Table 1. Parameters and timing breakdown per frame for all simulations in this paper and the accompanying video. Symbols Γ and a are the intensity and the thickness of filaments. All simulations are 24fps and the time step size 1.0/24s except for Figure 15, which we used 120fps. Average timings are taken over the entire simulation. The “Other” column includes the remaining operations, including obstacle handling and generation of new curves at the sources in Figure 9 and Figure 14.

Scene	Resolution	Γ	a	Total time	Evaluate $\mathbf{v}$	Advect ψ	Construct γ	Construct ψ	Other
Figure 13	$100 \times 100 \times 100$	1.0	0.05	0.159s	0.041s	0.017s	0.041s	0.056s	0.002s
Figure 12	$100 \times 100 \times 120$	1.0	0.05	0.274s	0.056s	0.026s	0.049s	0.102s	0.001s
Figure 7	$60 \times 60 \times 200$	2.0	0.08	0.090s	0.013s	0.021s	0.023s	0.031s	0.001s
Figure 8	$80 \times 120 \times 140$	1.0	0.05	0.152s	0.031s	0.030s	0.046s	0.044s	0.001s
Figure 1	$100 \times 100 \times 140$	1.0	0.05	0.655s	0.144s	0.085s	0.091s	0.332s	0.002s
Figure 9	$100 \times 100 \times 180$	2.0	0.08	1.423s	0.177s	0.099s	0.812s	0.324s	0.011s
Figure 14	$100 \times 100 \times 160$	2.0	0.08	0.969s	0.047s	0.087s	0.473s	0.226s	0.135s
Figure 15	$100 \times 100 \times 100$	N/A	N/A	0.152s	0.028s	0.012s	0.068s	0.043s	0.001s

Table 2. Computational timings compared with Houdini’s built-in implementation of Weissmann & Pinkall [2010]. “Same DOF” refers to simulations with approximately the same number of computational degrees of freedom as our method: we set the relevant parameters (re-connection distance, minimum and maximum edge lengths) so that the number of explicit curve vertices are similar to ours.

Scene description	W&P	W&P (Same DOF)	Ours
Linked rings (Figure 13)	0.052s	0.065s	0.159s
Trefoil knot (Figure 12)	0.045s	0.101s	0.274s

is most evident in our jet example (Figure 9). We believe that the deletion of small features can be reduced in the future in a number of ways. More accurate advection schemes will probably preserve higher frequency features of ψ better without deleting them. Also, although we took care to introduce a $\mathbf{v}$ and ψ which alleviates egregious numerical stability problems, we have not quite optimized for accuracy or geometric durability for the specific dynamics of vortex filaments. Finding a pair of $\mathbf{v}$ and ψ that is more suitable for each curve dynamics in the degrees of freedom discussed in Section 3 would lead to even better numerical performance.

Although our method relies on Eulerian advection to evolve the geometry, it *does not* suffer from the artificial viscosity typically associated with Eulerian fluid simulations. Our fluid velocity is reconstructed from filament dynamics, so the velocity field is not recursively re-sampled and does not accumulate damping errors. Consequently, our method produces swirly and energetic fluid flows even at low grid resolutions.

Table 1 lists the simulation parameters and timing breakdown for each of the simulations in our paper and accompanying video. We stress that our prototype implementation is meant as a proof of concept, and it has plenty of room for optimization. Our current implementation employs regular grids for ease of implementation; future implementations can make great use of sparse grids, since our filaments only use a 1-dimensional path through the 3D grid. Our implementation of “Evaluate $\mathbf{v}$ ” and “Construct ψ ” iterates over the entire curve geometry for each point in the narrow band where γ and ψ are required; future implementations could use a sparse data structure and approximated fast summation via tree-codes like the fast multipole method to reduce total evaluation time. Additionally,

our current implementation of “Construct γ ” redundantly doubles the work per curve vertex, so that step can be sped up by at least a factor of 2.

With these inefficiencies in mind, our implementation appears to run modestly slower than Houdini’s optimized implementation of Weissmann & Pinkall [2010] for the scenarios we tested. The two methods have completely different numerical degrees of freedom, so we find it difficult to compare them directly. Table 2 compares the computational cost of the two methods for figures in this paper. We aimed to keep the number of curve vertices roughly the same as our method’s in the “Same DOF” simulations, so we believe these are the most relevant for comparing timings.

In our examples of vortex filaments, we assumed inviscid filaments, but we can handle viscous motions too once the time-derivative of viscous curves is specified. We also assumed a uniform vortex strength. Just as in the explicit description, our implicit description requires a well-defined equation of motion for curves, and it outputs only new curve configurations without processing additional quantities like vortex strengths. Simulating vortex filaments with different strengths is an independent challenge. Merging filaments with different strengths would create a graph of filaments rather than just disjoint closed curves, which would require a new implicit representation for filament graphs.

7 CONCLUSION AND FUTURE WORK

We have shown that implicit representations of geometric curves exhibit large degrees of freedom in both their mathematical representation, as well as their dynamics. We then took advantage of these redundant degrees of freedom to improve the stability of vortex filament simulations.

We see a number of avenues that can be explored in future work. Our current strategy exploits redundancy in the 3D velocity field and level set outside the filaments, but we can also incorporate free degrees of freedom in the tangential components of the velocity on the filament itself. More generally, one can attempt to formalize the regularization of $\mathbf{v}$ and ψ for various purposes. A possible instance is an optimization problem for certain energies aiming provable guarantees on numerical accuracy and stability. Besides, the use of

redundancy in dynamics should also be possible for other codimensional cases such as level surfaces in 3D or level sets of an arbitrary codimension in a higher dimensional space.

Our numerical scheme can be made more sophisticated as well: higher order advection schemes and geometric curve representations, as well as sparse and adaptive grids can make our method both more efficient and more numerically accurate.

Finally, this paper explores vortex filaments and curve shortening flows, but our ideas are not limited to these specific dynamical systems. We expect that the idea of exploiting hidden degrees of freedom in implicitly represented curve dynamics will generalize to many more types of dynamics appearing in both scientific fields and engineering applications.

ACKNOWLEDGMENTS

We thank the visual computing group at IST Austria for their valuable discussions and feedback. Houdini Education licenses were provided by SideFX software. This project was funded in part by the European Research Council (ERC Consolidator Grant 101045083 *CoDiNa*).

REFERENCES

- Luigi Ambrosio and Halil Mete Soner. 1996. Level set approach to mean curvature flow in arbitrary codimension. *Journal of differential geometry* 43, 4 (1996), 693–737.
- Alexis Angelidis. 2017. Multi-scale vorticle fluids. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–12.
- 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*. 87–96.
- Vladimir I. Arnold and Boris A. Khesin. 1998. *Topological Methods in Hydrodynamics*. Springer.
- James Arvo. 1995. Applications of irradiance tensors to the simulation of non-lambertian phenomena. *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, 335–342.
- Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grinspun. 2008. Discrete Elastic Rods. , Article 63 (2008), 12 pages. <https://doi.org/10.1145/1399504.1360662>
- Fabrice Bethuel, Haïm Brezis, Frédéric Hélein, et al. 1994. *Ginzburg-landau vortices*. Vol. 13. Springer.
- Michael G. Bevis and Greg Cambareri. 1987. Computing the area of a spherical polygon of arbitrary shape. *Mathematical Geology* 19 (1987), 335–346.
- Jack Binysh and Gareth P Alexander. 2018. Maxwell’s theory of solid angle and the construction of knotted fields. *Journal of Physics A: Mathematical and Theoretical* 51, 38 (aug 2018), 385202. <https://doi.org/10.1088/1751-8121/aad8c6>
- 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*. Citeseer, 87–95.
- Paul Burchard, Li-Tien Cheng, Barry Merriman, and Stanley Osher. 2001. Motion of curves in three spatial dimensions using a level set approach. *J. Comput. Phys.* 170, 2 (2001), 720–741.
- Albert Chern. 2017. *Fluid dynamics with incompressible Schrödinger flow*. Ph.D. Dissertation. California Institute of Technology.
- Albert Chern, Felix Knöppel, Ulrich Pinkall, and Peter Schröder. 2017. Inside fluids: Clebsch maps for visualization and processing. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–11.
- Albert Chern, Felix Knöppel, Ulrich Pinkall, Peter Schröder, and Steffen Weißmann. 2016. Schrödinger’s smoke. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–13.
- Alexandre Joel Chorin. 1990. Hairpin removal in vortex interactions. *J. Comput. Phys.* 91, 1 (1990), 1–21.
- A. Clebsch. 1859. Ueber die Integration der hydrodynamischen Gleichungen. *Journal für die reine und angewandte Mathematik* 56 (1859), 1–10. English translation by D. H. Delphenich, http://www.neo-classical-physics.info/uploads/3/4/3/6/34363841/clebsch_-_clebsch_variables.pdf.
- Georges-Henri Cottet, Petros D Koumoutsakos, et al. 2000. *Vortex methods: theory and practice*. Vol. 8. Cambridge university press Cambridge.
- Fang Da, Christopher Batty, Chris Wojtan, and Eitan Grinspun. 2015. Double bubbles sans toil and trouble: Discrete circulation-preserving vortex sheets for soap films and foams. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–9.
- Sharif Elcott, Yiyi Tong, Eva Kanso, Peter Schröder, and Mathieu Desbrun. 2007. Stable, circulation-preserving, simplicial fluids. *ACM Transactions on Graphics (TOG)* 26, 1 (2007), 4–es.
- Manuel Noronha Gamito, Pedro Faria Lopes, and Mário Rui Gomes. 1995. Two-dimensional simulation of gaseous phenomena using vortex particles. In *Computer Animation and Simulation ’95*. Springer, 3–15.
- Evgenii A Kuznetsov and Aleksandr V Mikhailov. 1980. On the topological meaning of canonical Clebsch variables. *Physics Letters A* 77, 1 (1980), 37–38.
- Horace Lamb. 1895. *Hydrodynamics*. Cambridge University Press.
- William E. Lorensen and Harvey E. Cline. 1987. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. *SIGGRAPH Comput. Graph.* 21, 4 (aug 1987), 163–169. <https://doi.org/10.1145/37402.37422>
- Jerrold Marsden and Alan Weinstein. 1983. Coadjoint Orbits, Vortices, and Clebsch Variables for Incompressible Fluids. *Physica D: Nonlinear Phenomena* 7, 1 (1983), 305–323.
- Chohong Min. 2004. Local level set method in high dimension and codimension. *Journal of computational physics* 200, 1 (2004), 368–382.
- Philip J Morrison. 1998. Hamiltonian description of the ideal fluid. *Reviews of modern physics* 70, 2 (1998), 467.
- Kunio Murasugi. 2008 - 1996. *Knot theory & its applications*. Birkhäuser, Boston.
- Ken Museth, Nick Avramoussis, and Dan Bailey. 2019. OpenVDB. In *ACM SIGGRAPH 2019 Courses*. 1–56.
- Mohammad Sina Nabizadeh, Albert Chern, and Ravi Ramamoorthi. 2021. Kelvin Transformations for Simulations on Infinite Domains. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 97:1–97:15.
- Stanley Osher and James A Sethian. 1988. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of computational physics* 79, 1 (1988), 12–49.
- Marcel Padilla, Albert Chern, Felix Knöppel, Ulrich Pinkall, and Peter Schröder. 2019. On bubble rings and ink chandeliers. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–14.
- David Palmer, David Bommes, and Justin Solomon. 2020. Algebraic representations for volumetric frame fields. *ACM Transactions on Graphics (TOG)* 39, 2 (2020), 1–17.
- 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*. 261–270.
- Tobias Pfaff, Nils Thuerey, and Markus Gross. 2012. Lagrangian vortex sheets for animating fluids. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–8.
- Len M Pismen, Len M Pismen, et al. 1999. *Vortices in nonlinear fields: from liquid crystals to superfluids, from non-equilibrium patterns to cosmic strings*. Vol. 100. Oxford University Press.
- Ziyin Qu, Xinxin Zhang, Ming Gao, Chenfanfu Jiang, and Baoquan Chen. 2019. Efficient and conservative fluids using bidirectional mapping. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–12.
- Steven J Ruuth, Barry Merriman, Jack Xin, and Stanley Osher. 2001. Diffusion-generated motion by mean curvature for filaments. *Journal of Nonlinear Science* 11, 6 (2001), 473–493.
- Philip Geoffrey Saffman. 1992. *Vortex Dynamics*.
- H. Seifert. 1935. Über das Geschlecht von Knoten. *Math. Ann.* 110, 1 (1935), 571–592.
- Andrew Selle, Ronald Fedkiw, Byungmoon Kim, Yingjie Liu, and Jarek Rossignac. 2008. An unconditionally stable MacCormack method. *Journal of Scientific Computing* 35, 2 (2008), 350–371.
- Andrew Selle, Nick Rasmussen, and Ronald Fedkiw. 2005. A vortex particle method for smoke, water and explosions. In *ACM SIGGRAPH 2005 Papers*. 910–914.
- Justin Solomon, Amir Vaxman, and David Bommes. 2017. Boundary element octahedral fields in volumes. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1.
- Steffen Weißmann and Ulrich Pinkall. 2009. Real-time Interactive Simulation of Smoke Using Discrete Integrable Vortex Filaments. In *Workshop in Virtual Reality Interactions and Physical Simulation “VRIPHYS” (2009)*, Hartmut Prautzsch, Alfred Schmitt, Jan Bender, and Matthias Teschner (Eds.). The Eurographics Association. <https://doi.org/10.2312/PE/vrphys/vrphys09/001-010>
- Steffen Weißmann and Ulrich Pinkall. 2010. Filament-Based Smoke with Vortex Shedding and Variational Reconnection. *ACM Trans. Graph.* 29, 4, Article 115 (jul 2010), 12 pages. <https://doi.org/10.1145/1778765.1778852>
- Steffen Weißmann and Ulrich Pinkall. 2012. Underwater Rigid Body Dynamics. *ACM Trans. Graph.* 31, 4, Article 104 (jul 2012), 7 pages. <https://doi.org/10.1145/2185520.2185600>
- Steffen Weißmann, Ulrich Pinkall, and Peter Schröder. 2014. Smoke Rings from Smoke. *ACM Trans. Graph.* 33, 4, Article 140 (jul 2014), 8 pages. <https://doi.org/10.1145/2601097.2601171>
- Shiying Xiong, Rui Tao, Yaorui Zhang, Fan Feng, and Bo Zhu. 2021. Incompressible Flow Simulation on Vortex Segment Clouds. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 98:1–98:11.
- Shuqi Yang, Shiying Xiong, Yaorui Zhang, Fan Feng, Jinyuan Liu, and Bo Zhu. 2021. Clebsch gauge fluid. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–11.

- Xinxin Zhang and Robert Bridson. 2014. A PPPM fast summation method for fluids and beyond. *ACM Transactions on Graphics (TOG)* 33, 6 (2014), 1–11.
- Xinxin Zhang, Robert Bridson, and Chen Greif. 2015. Restoring the missing vorticity in advection-projection fluid solvers. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–8.

A PROPERTIES OF THE UNTWISTED CLEBSCH VARIABLES

As explained in Section 4.1, our untwisted Clebsch variable is composed of the distance function and the solid angle of projected curves. It has following desirable properties for curve dynamics:

- (1) it has an explicit formula and is uniquely computed from given curves;
- (2) it is locally little twisted and $d\psi|_{\gamma^\perp}$ is locally nearly-isometry;
- (3) it is harmonic;
- (4) the zeros of the real and the imaginary parts intersect orthogonally, which ensures the zeros to be always codimension-2 and makes numerically extracting zeros robust.

Open curves. We also note that our construction of ψ is not limited to closed curves but is also valid for open curves with end points located on the boundary or obstacles subject to the following integrability condition. The orientation of each curve assigns a positive or negative signature for its two end points. The collection of curves is said to be integrable if all ends of the curves lie on the boundary and each connected component of the boundary contains an equal number of positive and negative ends. Note that vortex filaments must be integrable since this integrability condition is precisely the integrability condition for curl: The vorticity 2-form is the exterior derivative of a velocity 1-form (i.e. exact) if and only if it is closed (divergence-free) and its restriction to every boundary component has zero total flux. In this case, the integrability of curl ensures that the endpoints of open filaments landing on boundaries must give equal positive and negative ends per boundary component. Therefore, it is possible to pair the endpoints along the boundary and complete the filaments as closed curves, from which we know how to construct ψ .