Inside Fluids: Clebsch Maps for Visualization and Processing

ALBERT CHERN, Caltech

FELIX KNÖPPEL, TU Berlin

ULRICH PINKALL, TU Berlin

PETER SCHRÖDER, Caltech

Clebsch maps encode velocity fields through functions. These functions contain valuable information about the velocity field. For example, closed integral curves of the associated vorticity field are level lines of the vorticity Clebsch map. This makes Clebsch maps useful for visualization and fluid dynamics analysis. Additionally they can be used in the context of simulations to enhance flows through the introduction of subgrid vorticity. In this paper we study *spherical* Clebsch maps, which are particularly attractive. Elucidating their geometric structure, we show that such maps can be found as minimizers of a non-linear Dirichlet energy. To illustrate our approach we use a number of benchmark problems and apply it to numerically given flow fields. Code and a video can be found in the ACM Digital Library.

CCS Concepts: • **Mathematics of computing → Partial differential equations**; • **Computing methodologies → Physical simulation**; • **Applied computing → Physics**;

Additional Key Words and Phrases: Fluid dynamics, Clebsch maps, flow visualization, analysis, and enhancement

ACM Reference format:

Albert Chern, Felix Knöppel, Ulrich Pinkall, and Peter Schröder. 2017. Inside Fluids: Clebsch Maps for Visualization and Processing. *ACM Trans. Graph.* 36, 4, Article 142 (July 2017), 11 pages.

DOI: http://dx.doi.org/10.1145/3072959.3073591

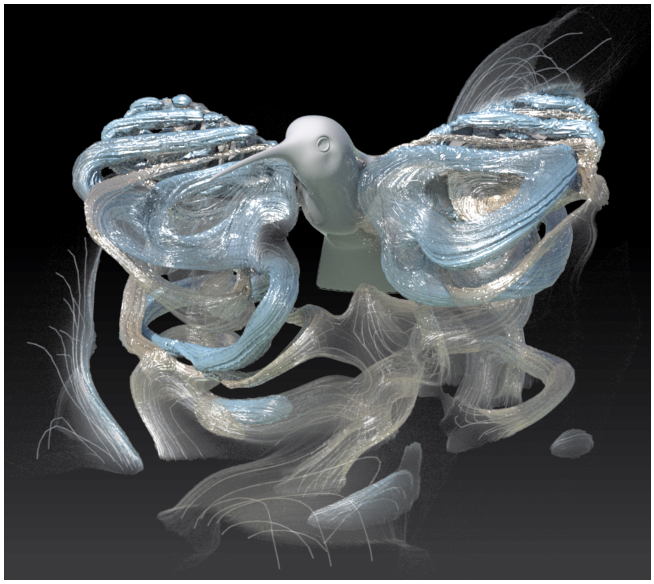

Fig. 1. Visualization of the vorticity field produced by the flapping wings of a Hummingbird. Flow data and photogrammetrically acquired bird geometry courtesy Haibo Dong, Flow Simulation Research Group, University of Virginia [Ren et al. 2016]. See the video at 02:00.

1 INTRODUCTION

In his study of the equations of hydrodynamics, Alfred Clebsch in 1859 proposed what is now known as the *Clebsch representation* [Clebsch 1859]. It encodes the fluid velocity field $\mathbf{u}\colon M \to \mathbb{R}^3$ ($M \subset \mathbb{R}^3$) with the aid of a function $\psi := (\lambda, \mu, \phi)\colon M \to \mathbb{R}^3$, writing

$$\mathbf{u} = \lambda \operatorname{grad} \mu - \operatorname{grad} \phi. \tag{1}$$

The vorticity vector field $\mathbf{w}$, in turn, is represented by the function $s := (\lambda, \mu)\colon M \to \mathbb{R}^2$

$$\mathbf{w} = \operatorname{curl} \mathbf{u} = \operatorname{grad} \lambda \times \operatorname{grad} \mu \tag{2}$$

due to curl(grad) = 0.

DOI: http://dx.doi.org/10.1145/3072959.3073591

Such "encodings" of velocity resp. vorticity fields are useful for flow visualization [Kotiuga 1991], analysis [Jeong and Hussain 1995], simulation [Brandenburg 2010; Cartes et al. 2007; He and Yang 2016; Yang and Pullin 2010], and enhancement among others. For example, Eq. (2) implies that pre-images of points, $s^{-1}(\{p\})$, correspond to closed *vortex lines* (integral curves of the vorticity field), and pre-images of regions, $s^{-1}(\Omega)$, to *vortex tubes* (Fig. 1). Similarly, given s on some computational grid one can add properly aligned turbulence at the subgrid level through simple manipulations involving s (Fig. 11).

So what's not to like? There are a number of problems with Clebsch' original proposal. For example, it can only represent fields with zero helicity (Thm. 4.1). Moreover, near points of vanishing vorticity, smooth functions ψ may not exist, even locally [Graham and Henyey 2000]. These problems can be addressed with *generalized* Clebsch maps [Cartes et al. 2007; Graham and Henyey 2000; Zakharov and Kuznetsov 1997]. Unfortunately, most of these no longer yield a level set representation for vortex lines. To maintain the latter we need to replace the range space of s with a surface more general than $\mathbb{R}^2$ (and similarly for ψ). Indeed, using the 2-sphere $\mathbb{S}^2$ for s and

the 3-sphere $\mathbb{S}^3$ for ψ yields *spherical* Clebsch maps. These are the only known example capable of representing fields with non-zero helicity *and* giving a level set representation for vortex lines [Chern et al. 2016; Kotiuga 1991; Kuznetsov and Mikhailov 1980].

What is missing, however, is an algorithm to find a Clebsch map for a numerically given vector field. In this paper we give an algorithm to produce spherical Clebsch maps ψ approximating a vector field $\mathbf{u}$ on some domain M. (While such vector fields often arise from incompressible fluid simulations, there is no requirement that the field be divergence free.) The desired ψ is found through minimizing a non-linear Dirichlet energy, which encapsulates fidelity and smoothness terms with a variable trade-off.

Overview. Before we delve into the details of our algorithm, we set the stage with a more abstract definition of Clebsch maps (Sec. 2). To make this manageable in the face of manifold-valued functions we employ the formalism of *exterior calculus*. (Readers wishing for accessible explanations of differential geometric concepts are directed to [Frankel 2012], while [Abraham et al. 2001, Ch. 7] provides a comprehensive review of exterior calculus.) With that machinery available we are equipped to describe spherical Clebsch maps (Sec. 3) and formulate a strategy for finding such maps for a given velocity field (Sec. 3.1). The remainder of the paper is devoted to validation (Sec. 4) of the basic method and suggestions for applications (Sec. 5). Proofs and implementation details are delegated to Apps. A to E.

Table 1. Notation.

Notation	Meaning
$\mathbf{r}$	a point on M
$\mathbf{u}$	velocity vector field on M
$\eta = \mathbf{u} \cdot d\mathbf{r}$	velocity 1-form
$\mathbf{w} = \operatorname{curl} \mathbf{u}$	vorticity vector field
$\omega = d\eta = \mathbf{w} \cdot (d\mathbf{r} \times d\mathbf{r})$	vorticity 2-form
$\mathbb{H} \equiv \mathbb{R}^4$	quaternions
$\langle \cdot, \cdot \rangle$	scalar product on $\mathbb{R}^4$
$\Omega^k(M; V)$	V valued k-forms on M
$\psi^* f = f \circ \psi$	pullback of a function f
$\psi^* \alpha(\cdot) = \alpha(d\psi(\cdot))$	pullback of a 1-form α
$\psi^* \sigma(\cdot, \cdot) = \sigma(d\psi(\cdot), d\psi(\cdot))$	pullback of a 2-form σ

2 GENERAL CLEBSCH MAPS

Instead of working with the vector field $\mathbf{u}$ we now work with the corresponding 1-form (co-vector field) η.

Let $\eta \coloneqq \mathbf{u} \cdot d\mathbf{r}$ be a velocity 1-form on M, then Eq. (1) reads

$$\eta = \lambda d\mu - d\phi.$$

This can be expressed succinctly as the *pullback* of a fixed 1-form $\alpha = x dy - dz$ by a function $\psi = (\lambda, \mu, \phi)$

$$\eta = \psi^* \alpha.$$

The formalism of pullback encapsulates the underlying change of variables. Since exterior derivative and pullback commute Eq. (2) becomes

$$\omega = d\eta = d(\psi^* \alpha) = \psi^* d\alpha = \psi^*(dx \wedge dy),$$

where we used $d \circ d = 0$. Since the wedge product $dx \wedge dy$ is the area 2-form $dA_{\mathbb{R}^2}$ on $\mathbb{R}^2$, we can also write $\omega = s^* dA_{\mathbb{R}^2}$ for $s = (\lambda, \mu)$, giving us

$$\int_\Omega \omega = \int_\Omega s^* dA_{\mathbb{R}^2} = \operatorname{Area}(s(\Omega)). \tag{3}$$

This equation tells us that the vorticity flux over some surface $\Omega \subset M$ can be found by computing the (signed) area (with multiplicity) of the image of Ω under s.

Importantly, having rewritten Clebsch' original proposal using pullbacks we now see the underlying geometry. A Clebsch map describes a variable 1-form through a *fixed* 1-form and a *varying* function. This leads to

Definition 2.1. **Clebsch Representation** Let C and Σ be manifolds with $\alpha \in \Omega^1(C; \mathbb{R})$ a fixed 1-form and $\sigma \in \Omega^2(\Sigma; \mathbb{R})$ a fixed 2-form, and $\pi \colon C \to \Sigma$ such that $\pi^* \sigma = d\alpha$. We call a map

- $\psi \colon M \to C$ a Clebsch map for $\eta \in \Omega^1(M; \mathbb{R})$ if $\psi^* \alpha = \eta$.
- $s \colon M \to \Sigma$ a Clebsch map for $\omega \in \Omega^2(M; \mathbb{R})$ if $s^* \sigma = \omega$.

A Clebsch map $\psi \colon M \to C$ for η automatically implies that $s \coloneqq \pi \circ \psi$ is a Clebsch map for $d\eta = \omega$. Conversely, in many situations it is true that given a Clebsch map $s \colon M \to \Sigma$ for $\omega = d\eta$ there will always exist a Clebsch map $\psi \colon M \to C$ for η such that $s = \pi \circ \psi$ (Thm. 3.1).

Def. 2.1 tells us that Clebsch maps are distinguished by their choice of manifolds (C, Σ), 1- resp. 2-forms (α, σ), and the functions (ψ, s) connecting them:

- Clebsch' original proposal [1859] uses $\psi \colon M \to \mathbb{R}^3$, $\psi = (\lambda, \mu, \phi)$ with $\alpha = x dy - dz$ to get a Clebsch map for η, and $\pi \colon \mathbb{R}^3 \to \mathbb{R}^2$ with $\sigma = dA_{\mathbb{R}^2}$ for ω. Such maps can only represent fields with vanishing helicity (Thm. 4.1).
- A natural generalization of Clebsch' original proposal uses n-dimensional λ and μ [Cartes et al. 2007; Graham and Henyey 2000; Zakharov and Kuznetsov 1997]. Now $\psi = (\boldsymbol{\lambda}, \boldsymbol{\mu}, \phi) : M \to \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}$ for $n = 2, 3$ and $\alpha = \mathbf{x} \cdot d\mathbf{y} - dz$ using the coordinate functions $(\mathbf{x}, \mathbf{y}, z)$ on $\mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}$, yields a Clebsch map for η. In turn ω has a Clebsch map $s : M \to \mathbb{R}^n \times \mathbb{R}^n$, $s = (\boldsymbol{\lambda}, \boldsymbol{\mu})$ with $\sigma = d\mathbf{x} \wedge d\mathbf{y}$. Since $\dim(\Sigma) = 2n > 2$ there is no accessible representation of vortex lines.
- Chern *et al.* [2016] gave a Clebsch map $\psi \colon M \to \mathbb{S}^3$ for η. Using the embedding q of $\mathbb{S}^3$ into the space of quaternions $\mathbb{H}$, we can write $\alpha = \hbar \langle dq, iq \rangle$ where $\hbar > 0$ is a dimensional constant. A Clebsch map for ω with $\sigma = \frac{\hbar}{2} dA_{\mathbb{S}^2}$ is obtained with the Hopf map $\pi \colon \mathbb{S}^3 \to \mathbb{S}^2$, $\pi = \bar{q} i q$ [Hopf 1931; Lyons 2003]. Their Clebsch map s for ω alone was also considered earlier [Kuznetsov and Mikhailov 1980] though only Chern *et al.* used it in a computational context.

The particular form of Clebsch' original proposal was motivated by the fact that the expression must be non-integrable to be able

to describe flows with non-zero vorticity. In fact, α is a *contact form* whose kernel is nowhere Frobenius integrable [Geiges 2006] (see [Delphenich 2017] for its relation to Clebsch' proposal).

For completeness we note that Clebsch maps using three-component $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ can recover the Cauchy-Weber transform [Cauchy 1815; Weber 1868] of the Euler equations for incompressible as well as barotropic, ideal (inviscid) fluids [Euler 1757]. These representations have been extended in the *Eulerian-Lagrangian* approach [Constantin 2001a,b] to include incompressible viscid (Navier-Stokes) fluids. For an historical account tracing these developments back to Cauchy [1815] see [Frisch and Villone 2014].

In Computer Graphics the classic Clebsch maps (λ, μ) have made an appearance as representations of divergence free vector fields for purposes of geometric modeling [Angelidis and Singh 2007; von Funck et al. 2006]. In this setting (λ, μ) are control variables for the deformation field and no inverse problem, finding the variables if the field is given, needs solving.

Summary. The only known Clebsch maps which maintain the compact representation of Clebsch' original proposal, can recover vortex lines as pre-images of points, and approximate velocity fields of arbitrary helicity, are those of Chern and co-workers [2016] and Kuznetsov and Mikhailov [1980]. Hence we will focus on spherical, *i.e.*, $\mathbb{S}^3$- resp. $\mathbb{S}^2$-valued, velocity and vorticity Clebsch maps in the remainder of this paper.

3 SPHERICAL CLEBSCH MAPS

Before diving into the algorithms we need to clarify the relationship between velocity and vorticity Clebsch maps. In the classical setting this was quite straightforward. In the spherical case it is very similar but it needs a little more work to see this.

In this paper we follow [Chern et al. 2016] and work with $C = \mathbb{S}^3$, $\alpha = \hbar\langle dq, iq\rangle$ (for the choice of $\hbar > 0$ see Sec. 5), $\Sigma = \mathbb{S}^2$, $\sigma = \frac{\hbar}{2} dA_{\mathbb{S}^2}$ and π the Hopf map, $\pi = \bar{q}iq$. The fact that $\pi^*\sigma = d\alpha$ follows from [Chern et al. 2016, Thm. 1, Sec. 4.3] applied to the special case $\psi = q$.

Conversely, if we have a vorticity Clebsch map $s\colon M \to \mathbb{S}^2$ then there exists an almost unique Clebsch map for η:

Theorem 3.1. ***Existence of Lifts*** *Suppose we have a Clebsch map $s\colon M \to \mathbb{S}^2$ for an exact 2-form ω. Then there is a Clebsch map $\psi\colon M \to \mathbb{S}^3$ for some η with $\omega = d\eta$ and $\pi \circ \psi = s$ (see App. A for a simplified proof).*

Note that η is not required to be divergence free, though $d\eta$ is of course always divergence free.

Corollary 3.2. ***Non-Uniqueness*** *If ψ is a Clebsch map for a velocity 1-form η, it cannot be unique. If $f\colon \mathbb{S}^2 \to \mathbb{S}^2$ is any area-preserving map, then $\tilde{s} := f \circ \pi \circ \psi$ is also a Clebsch map for $\omega = d\eta$. By Thm. 3.1 there is another Clebsch map $\tilde{\psi}$ for η with $\pi \circ \tilde{\psi} = \tilde{s}$.*

3.1 Finding a Spherical Clebsch Map

For a given η_0 an approximating velocity Clebsch map ψ can be found as a minimizer of

$$E(\psi) = \int_M \tfrac{1}{\hbar^2} |\eta - \eta_0|^2 \tag{4}$$

for $\eta := \psi^*\alpha$, yielding the vorticity Clebsch map $s = \pi \circ \psi$ as well.

Suppose we have found an optimal approximation ψ for the given η_0. By Cor. 3.2 such a ψ gives rise to infinitely many more minimizers $\tilde{\psi}$ through area preserving diffeomorphisms of $\mathbb{S}^2$. To fix this degree of freedom we add the Dirichlet energy of s as a regularizer to Eq. (4), seeking now a minimizer as $\epsilon \to 0$ for

$$\operatorname*{argmin}_{\psi\colon M \to \mathbb{S}^3} E^\epsilon(\psi) \quad \text{where} \quad E^\epsilon(\psi) := \int_M \tfrac{\epsilon}{4} |ds|^2 + \tfrac{1}{\hbar^2} |\eta - \eta_0|^2. \tag{5}$$

This leads to our main algorithm:

Algorithm 1 Optimization scheme for Eq. (5)

Input: $\eta_0 \in \Omega^1(M; \mathbb{R})$ and $\hbar > 0$.
  Randomize $\psi^{(0)}\colon M \to \mathbb{S}^3$;
  $\epsilon^{(0)} = 1$;
  for $k = 0, 1, 2, \ldots$ **do**
    $\psi^{(k+1)} \leftarrow \operatorname{argmin}_\psi E^{\epsilon^{(k)}}(\psi)$ with initial guess $\psi^{(k)}$.
    $\epsilon^{(k+1)} \leftarrow \epsilon^{(k)}/10$
  end for

3.2 The Energy E^ϵ

In this section we show that the energy of Eq. (5) can be rewritten as the Dirichlet energy of ψ under a suitably chosen *connection* and *metric*. To facilitate this derivation we represent elements ψ of $\mathbb{S}^3$ as *unit quaternions*. This allows us to use straightforward quaternion algebra to arrive at the final expression (Eq. (7)) of our energy.

Given η_0, define the connection (covariant derivative) ∇^{η_0} for $\mathbb{S}^3$-valued functions ψ

$$\nabla^{\eta_0}\psi := d\psi - i\tfrac{\eta_0}{\hbar}\psi. \tag{6}$$

For a given $\mathbf{r} \in M$, $(\nabla^{\eta_0})_\mathbf{r}$ is a linear map from $T_\mathbf{r}M$ to $T_{\psi(\mathbf{r})}\mathbb{S}^3 \subset \mathbb{H}$. Choose $\{\psi, i\psi, j\psi, k\psi\}_\mathbf{r}$ as an orthonormal basis for $\mathbb{H}$, where $\{i, j, k\}$ denote the quaternionic imaginary units, and note that $T_{\psi(\mathbf{r})}\mathbb{S}^3 = \operatorname{Span}\{i\psi, j\psi, k\psi\}_\mathbf{r}$. Now split $\mathbb{H}$ into the orthogonal sum of the two subspaces

$$\mathbb{C}\psi_\mathbf{r} := \operatorname{Span}\{\psi, i\psi\}_\mathbf{r} \quad \text{and} \quad \mathbb{C}j\psi_\mathbf{r} := \operatorname{Span}\{j\psi, k\psi\}_\mathbf{r},$$

and decompose $(\nabla^{\eta_0}\psi)_\mathbf{r}$ accordingly

$$(\nabla^{\eta_0}\psi)_\mathbf{r} = P_{\mathbb{C}\psi_\mathbf{r}}(\nabla^{\eta_0}\psi) + P_{\mathbb{C}j\psi_\mathbf{r}}(\nabla^{\eta_0}\psi),$$

where $P_{\mathbb{C}\psi_\mathbf{r}}$ and $P_{\mathbb{C}j\psi_\mathbf{r}}$ denote the corresponding projection operators. It is not difficult to work out (App. B) that

$$P_{\mathbb{C}\psi}(\nabla^{\eta_0}\psi) = \tfrac{1}{\hbar}(\eta - \eta_0)i\psi \quad \text{and} \quad P_{\mathbb{C}j\psi}(\nabla^{\eta_0}\psi) = -\tfrac{1}{2} i\psi ds.$$

Thus we can recover the integrand of Eq. (5) with the aid of an anisotropic metric on $\mathbb{S}^3$. At each point $\psi(\mathbf{r}) \in \mathbb{S}^3$ define a norm on tangent vectors $X \in T_{\psi(\mathbf{r})}\mathbb{S}^3 \subset \mathbb{H}$ as

$$|X|_\epsilon^2 := |P_{\mathbb{C}\psi}(X)|^2 + \epsilon |P_{\mathbb{C}j\psi}(X)|^2.$$

Such a metric on $\mathbb{S}^3$, with an anisotropy along the $i\psi$ direction, is known as a *Berger metric* and $\mathbb{S}^3$ with such a metric as the *Berger sphere*. Using this metric, Eq. (5) is seen to be the connection Dirichlet energy of ψ

$$E^\epsilon(\psi) = \int_M |\nabla^{\eta_0}\psi|_\epsilon^2. \tag{7}$$

Summary. We cast the problem of finding spherical Clebsch maps, which approximate user supplied velocity fields η_0, as a minimization problem. The objective is the Dirichlet energy of ψ using the connection induced by η_0 employing a Berger metric. The metric anisotropy parameter ϵ controls how much of the regularizer $\|ds\|^2$ is added to the fidelity term $\|\eta - \eta_0\|^2$. For each fixed ϵ, in a decreasing sequence, the energy is then minimized.

3.3 Discussion

We have reduced the problem of finding an approximate Clebsch map ψ for a given η_0 to finding a minimizer of the connection Dirichlet energy of ψ (Eq. (7)). While we can hope for good approximations there are theoretical obstacles to finding exact solutions: (1) Isolated zeros in the vorticity field cannot be recovered by a spherical Clebsch map (Sec. 4.2); (2) vortex lines for $\mathbb{S}^2$ valued Clebsch maps are always closed (or begin and end on the domain boundary) while generically they are open (Sec. 4.3); (3) spherical Clebsch maps have quantized helicity, while general velocity fields can have arbitrary helicity (Sec. 4.4); and (4) owing to the non-linearity of the energy it is possible to get stuck in local minima.

In Sec. 4 we demonstrate that these theoretical limitations do not stand in the way of the practical utility of the approximate maps we compute.

There are also practical issues which are independent of the continuous setup: (1) finite resolution of the discretized system impacts the topology of the field; (2) the connection Laplacian can be very ill-conditioned; and (3) boundary conditions need custom tailoring when M "cuts" the field.

In our proof of concept implementation we minimized Eq. (7) through its L^2 gradient flow, $\dot{\psi} = -\operatorname{grad} E^\epsilon(\psi)$. This is a rather simple first order method whose main virtue is its straightforward implementation. To discretize this PDE in time we used the backward Euler method to allow for large, stable time steps. This requires the solution of $\mathbb{C}$-valued Poisson problems for which we used a standard conjugate gradient solver. For production use the latter would need an effective preconditioning strategy to address the ever increasing ill-conditioning under refinement, common to discretizations of second order elliptic operators. Implementation details can be found in App. C. Our implementation in Houdini 15 is included with the ancillary materials in the ACM Digital Library.

Table 2. Performance statistics. Time is given in units of minutes on a MacBook Pro, while "steps" gives the number of time discrete steps taken by the minimizer.

Fig.	Resolution	#(steps)	time	$\frac{\|\eta-\eta_0\|_{L^2}}{\|\eta_0\|_{L^2}}$	$\frac{\|\eta-\eta_0\|_{L^\infty}}{\|\eta_0\|_{L^\infty}}$
1	64^3	120	9:55	1.8×10^{-1}	1.05
2	128^2	100	1:08	1.8×10^{-4}	0.0012
4	50^3	130	6:10	1.9×10^{-2}	0.23
5	64^3	150	20:00	5.4×10^{-2}	0.15
7	$64^2 \times 48$	30	2:12	1.6×10^{-1}	0.3
8	81×41^2	100	6:04	3.0×10^{-2}	0.86
10	64^3	100	8:00	8.4×10^{-2}	1.21
12	64×32^2	120	1:25	1.3×10^{-2}	0.047

4 VALIDATION

Before discussing benchmark problems we briefly describe our approach to visualization.

To visualize Clebsch maps we use pre-images of sets on $\mathbb{S}^2$ under s. For individual points $p \in \mathbb{S}^2$, $s^{-1}(\{p\}) \subset M$ gives vortex lines, while pre-images of regions $\Omega \subset \mathbb{S}^2$ yield vortex tubes.

To construct $s^{-1}(\{p\})$ we represent s as a complex function $\zeta : M \to \mathbb{C}$ through stereographic projection from the antipode $-p$. The set $\{\mathbf{r} \in M \mid s(\mathbf{r}) = p\}$ is then given by the zeros of ζ and extracted as in [Weißmann et al. 2014, Sec. 3].

To visualize a vortex tube $s^{-1}(\Omega) \subset M$ we take some level set function on $\mathbb{S}^2$, $\chi : \mathbb{S}^2 \to \mathbb{R}$ with $\chi = 0$ on $\partial\Omega$. Extracting the zero iso-contour of $\chi \circ s$ then yields the vortex tube surface in M. For the images of vortex tubes shown in this paper we used $s = (s_1, s_2, s_3)^\intercal$ directly, drawing $s_1 = 0.8$ in gold and $s_1 = -0.8$ in blue. The opacity of the vortex tube was set proportional to vorticity magnitude, making it inversely proportional to the cross-sectional area of the vortex tube. Treating $p = (\pm 1, 0, 0)^\intercal$ as the north resp. south pole of $\mathbb{S}^2$, the longitude angle $\arg(\psi_1) - \arg(\psi_2)$ for $\psi = (\psi_1, \psi_2)^\intercal$ was used as a texture coordinate on the vortex tube surfaces for a brush texture, making the brush "lines" tangent to vortex lines.

4.1 2D Example

Fig. 2 shows a velocity field and its $\mathbb{S}^2$-valued vorticity Clebsch map. The disk, rotating rigidly, has constant, non-zero vorticity. This forces the Clebsch map to wrap the central disk multiply around a small ($\hbar = 0.15$) 2-sphere in an area-preserving manner. Simultaneously, the map is as conformal as possible, since the Dirichlet energy of s is the same as the conformal energy, up to a constant area term [Hutchinson 1991]. Outside the central disk, due to vanishing vorticity, the map is (nearly) constant, *i.e.*, it "covers" no area. This shows that our minimizers are not necessarily smooth. However, they appear (empirically) to be Lipschitz and smooth away from finitely many points. Fig. 3 shows the convergence plot for Fig. 2 which is typical of all the examples shown in this paper.

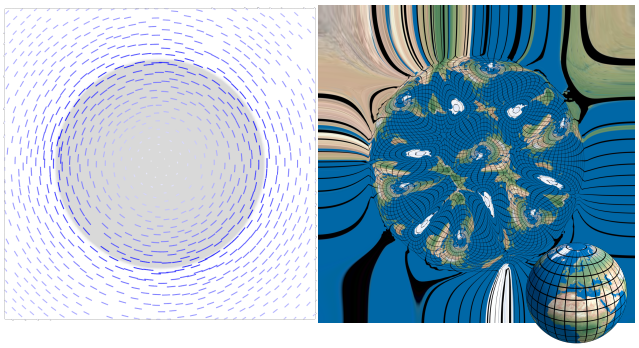

Fig. 2. A velocity field on a square (left) and its $\mathbb{S}^2$-valued vorticity Clebsch map (right; visualized via an earth texture). The velocity field has non-zero, constant vorticity in the gray disk and zero vorticity outside.

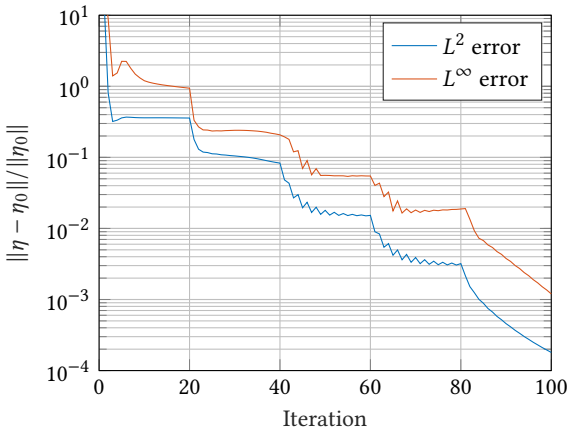

Fig. 3. Convergence of the Clebsch map in Fig. 2 as a function of gradient descent iterations with $\Delta t = 1$ and $\epsilon^{(0)} = 1$ (0-20), $\epsilon^{(1)} = 0.1$ (20-40), $\epsilon^{(2)} = 0.01$ (40-60), $\epsilon^{(3)} = 0.001$ (60-80), and $\epsilon^{(4)} = 0$ (>80).

4.2 Presence of Isolated Zeros

Graham and Henyey [2000] proved that a vorticity field $\mathbf{w}$ with an isolated zero at some $\mathbf{r} \in M$ does not admit an $\mathbb{R}^2$ valued Clebsch representation in the vicinity of $\mathbf{r}$. Their argument also applies to $\mathbb{S}^2$ valued Clebsch maps. In our case the presence of the regularizer ensures that an approximate solution is found nevertheless. Fig. 4 shows a Clebsch map for a linear vorticity field on $\mathbb{R}^3$, and we see that our algorithm deals gracefully with the isolated zero at the origin.

Even though no exact Clebsch map (at least not a smooth one) can exist for this field at the origin, we still get meaningful *approximations* of the field, as visualized with the vorticity integral curves.

4.3 Non-Closed Vortex Lines

In a generic flow most vortex lines are not closed. But Clebsch maps with target space $\mathbb{R}^2$ or $\mathbb{S}^2$ possess only closed (or beginning and ending on the boundary) vortex lines. Hence flows with vortex lines that neither close nor hit the boundary do not possess such a Clebsch map [Hadamard 1903, §68].

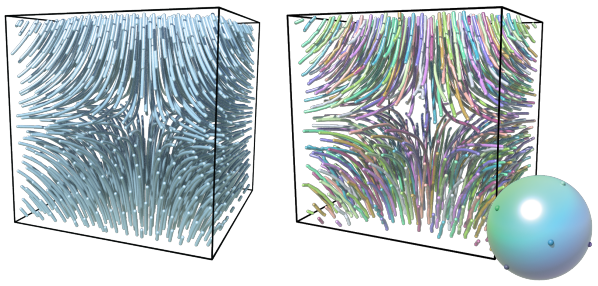

Fig. 4. On the left vortex lines traced in the vorticty vector field $\mathbf{w} = (x, y, -2z)^\intercal$ on a cube with side length three. On the right the pre-image of ten points on $\mathbb{S}^2$ for our Clebsch map approximating $\mathbf{w}$ ($\hbar = 0.05$).

By the Poincaré recurrence theorem [1890], almost every such vortex line will return arbitrarily closely to its initial point, making it "almost closed." Our Clebsch maps approximate these almost closed vortex lines by closed ones. Note that this "closure" depends on the resolution of the data. Fig. 5 gives an example of such a situation where most vortex lines of the original field do not close up, while the vortex lines of the field corresponding to the Clebsch maps are all closed.

Even though no exact Clebsch map exists for such fields, we still get a meaningful *approximation* with our method.

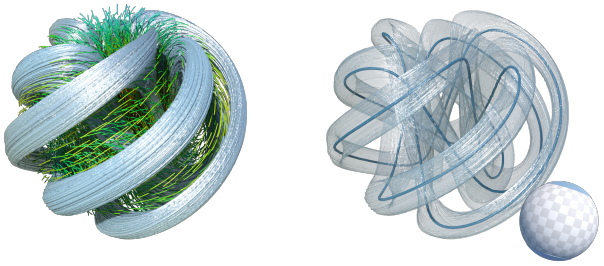

Fig. 5. Clebsch map for a vorticity field $\mathbf{w} = 1.5\mathbf{w}_1 + \mathbf{w}_2$, confined to a ball, with $\mathbf{w}_1 = (xz, yz, 1 - 2(x^2 + y^2) - z^2)^\intercal$ and $\mathbf{w}_2 = (-y, x, 0)^\intercal$. Both $\mathbf{w}_1, \mathbf{w}_2$ have closed vortex lines, but the ratio of their respective periods is irrational, ensuring that most vortex lines of $\mathbf{w}$ are non-closed. The field generated by the Clebsch map, visualized as pre-images of areas on $\mathbb{S}^2$ (left, blue tubes), approximates the field $\mathbf{w}$ (left, green trails) in a way that each vortex line is closed (right, two such vortex lines, *i.e.*, pre-images of two points on $\mathbb{S}^2$).

4.4 Helicity (Linked Vortex Lines)

Consider a vorticity vector field $\mathbf{w}$ on $\mathbb{R}^3$, vanishing outside a closed ball M. Then there is a (non-unique) vector field $\mathbf{u}$ on M such that $\operatorname{curl} \mathbf{u} = \mathbf{w}$. The quantity

$$\mathcal{H}(\mathbf{w}) := \int_M \mathbf{u} \cdot \mathbf{w} = \int_M \eta \wedge \omega \tag{8}$$

is independent of $\mathbf{u}$ (resp. η) and called the *helicity* of $\mathbf{w}$ (resp. ω) [Arnold and Khesin 1998, Ch. 3].

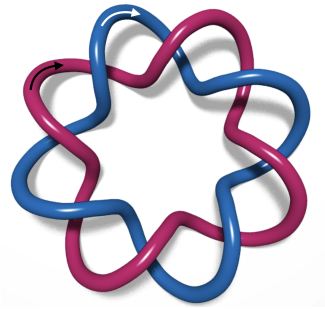

Fig. 6. Example of two linked curves (linking number −4).

To understand the geometric meaning of helicity consider two closed vortex lines, one through $\mathbf{r} \in M$ and the other through $\tilde{\mathbf{r}} \in M$. Let $\mathrm{lk}(\mathbf{r}, \tilde{\mathbf{r}})$ denote their *linking number* [Gauß 1833, p. 605], an integer indicating how often one vortex line winds around the other (Fig. 6). For arbitrary $\mathbf{r}, \tilde{\mathbf{r}} \in M$, whether the vortex lines through $\mathbf{r}$ and $\tilde{\mathbf{r}}$ close up or not, one can still define $\mathrm{lk}(\mathbf{r}, \tilde{\mathbf{r}})$ even though it is no longer an integer [Arnold and Khesin 1998, Chapter III, §1]. As before, $\mathrm{lk}(\mathbf{r}, \tilde{\mathbf{r}})$ measures the amount by which the two curves "spiral around each other."

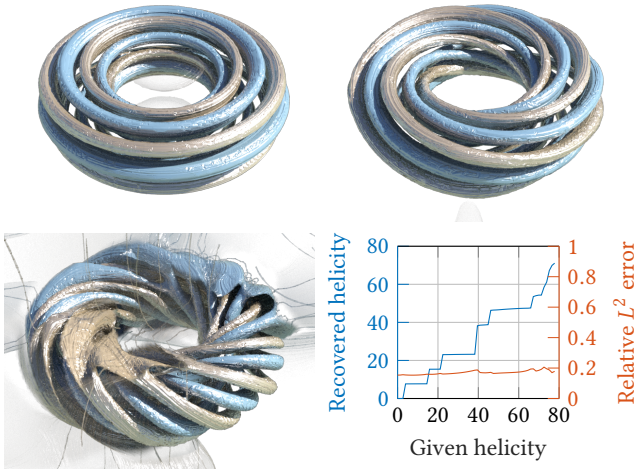

Fig. 7. Clebsch maps for vorticity fields, supported in the interior of a torus, with increasing helicity and a fixed $\hbar = 0.2$. The graph shows the helicity recovered by the Clebsch map and the L^2 error in velocity. Each cross section of the torus is passed by 5 vortex tubes for each of the two colors. In this case the linking number between any two tubes must be a multiple of 5 [Bush et al. 2017], giving rise to the gaps $5(2\pi\hbar)^2 \approx 7.9$ in each jump discontinuity in the plot of helicity. Note that vortex lines connecting to the boundary allow continuous helicity; such events occur at large given helicity. See the video at 00:36.

The relevant fact for us is that helicity can be expressed in terms of these (generalized) linking numbers

$$\mathcal{H}(\mathbf{w}) = \int_M \int_M \mathrm{lk}(\mathbf{r}, \tilde{\mathbf{r}})$$

So we see that helicity measures the average linking of vortex lines in a field. This reveals the main reason that classical Clebsch maps are of limited use:

Theorem 4.1. ***Vanishing Helicity*** *If* $\mathbf{w}$ *can be represented by an* $\mathbb{R}^2$*-valued Clebsch map then* $\mathcal{H}(\mathbf{w}) = 0$ *(see App. D for a proof).*

On the other hand, for $\mathbb{S}^2$-valued Clebsch maps helicity can be non-zero, but is quantized:

Theorem 4.2. ***$\mathbb{S}^2$-valued Clebsch Maps*** *If* $\mathbf{w}$ *can be represented by a Clebsch map that takes values in* $\mathbb{S}^2$ *and* $\omega = \frac{\hbar}{2} s^* dA_{\mathbb{S}^2}$*, then there is an integer* n *such that* $\mathcal{H}(\mathbf{w}) = n \cdot h^2$ *for* $h := 2\pi\hbar$ *(see App. D for a proof).*

Even though for $\mathbb{S}^2$-valued Clebsch maps the helicity is quantized to integer multiples of a (usually small) number, it does not present a practical obstacle to the approximation of fields with arbitrary helicity (Fig. 7).

Summary. In this section we demonstrated that spherical Clebsch maps perform well in practice even in the presence of otherwise challenging input fields. Specifically, they yield good approximations in the presence of isolated zeros (Sec. 4.2), gracefully approximate non-closed vortex lines (Sec. 4.3), and work well with fields carrying non-zero helicity (Sec. 4.4).

5 APPLICATIONS

After considering specific benchmark configurations we now turn to brief application vignettes to suggest possible uses for Clebsch maps.

5.1 Vorticity Visualization

Vorticity fields are of great interest in the study of low Mach number flows because the evolution of vorticity reveals much about the underlying flow dynamics [Saffman 1992]. One way to visualize vorticity is through vortex lines. A challenge in using vortex lines is the difficulty of selecting seed points so that the vortex lines have a spatial density proportional to vorticity magnitude. For $\mathbb{S}^2$-valued vorticity Clebsch maps this is achieved by picking points on $\mathbb{S}^2$ equidistributed with respect to area.

More specialized methods try to identify *regions* associated with vortices (for a comprehensive review see [Jiang et al. 2005]). One such method visualizes level sets of the vorticity norm $|\mathbf{w}|$ to capture regions of high vorticity (Fig. 8, upper left). Since directional information is ignored the resulting surfaces are misaligned in general, no longer being tangent to the vorticity field. A more sophisticated approach, the so-called λ_2-method [Jeong and Hussain 1995], analyses the velocity gradient tensor, but can still yield results which do not correctly represent the underlying vortex lines (Fig. 8; *cf.*

Fig. 9). In contrast, with a vorticity Clebsch map one can draw the pre-image of one or several closed curves on $\mathbb{S}^2$, yielding proper vortex tubes (Figs. 8 and 1).

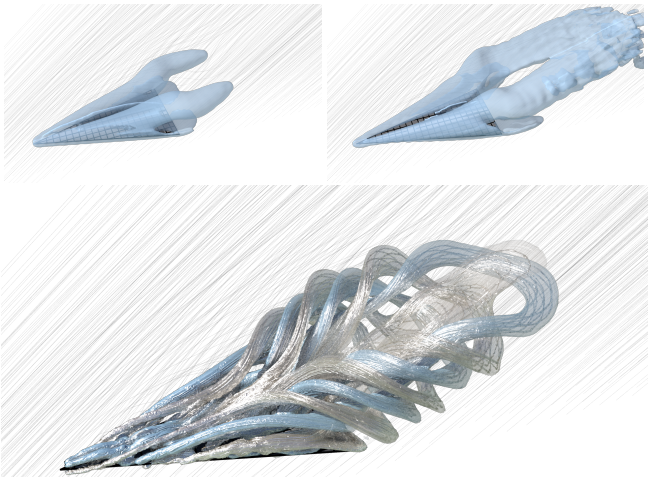

Fig. 8. Vorticity visualized through a $|\mathbf{w}|$ iso-surface (top-left), the λ_2-method (top-right), and a Clebsch map (bottom) for the Delta Wing data set [Ekaterinaris and Schiff 1990] (40° angle of attack, Mach 0.3, Reynolds number 10^6). Gray lines in the background are integral curves of the *velocity field* to give an overall sense of the flow. See the video at 01:39.

Fig. 9. Comparison of vortex tubes from the vorticity Clebsch map with traced integral curves of the original vorticity field (seeded at the wing edges), verifying the Clebsch map result (*cf.* Fig. 8; bottom).

The Parameter $\hbar$. Since the parameter $\hbar$ shows up as the factor in $\omega = \frac{\hbar}{2} s^* dA_{\mathbb{S}^2}$, the vortex tubes drawn by the pre-image $s^{-1}(\Omega)$ of a region $\Omega \subset \mathbb{S}^2$ have strength $\frac{\hbar}{2}\,\mathrm{Area}(\Omega)$. For a given flow and a region Ω, the parameter $\hbar$ controls the spatial frequency of visualized vortex tubes, increasing them as $\hbar$ is lowered. To avoid aliasing artifacts $\hbar$ must be chosen relative to the sampling rate $\Delta\mathbf{r}$ (inverse grid resolution) and velocity magnitude $|\mathbf{u}|_\infty$ to satisfy $\pi\hbar > 2\Delta\mathbf{r}|\mathbf{u}|_\infty$ ($\hbar$ has physical dimension $\mathrm{m}^2\,\mathrm{s}^{-1}$).

Time Coherence. We visualize pre-images of fixed sets in $\mathbb{S}^2$. However, the vorticity Clebsch map s arising from the minimizers ψ of Eq. (7) are only unique up to a global rotation of $\mathbb{S}^2$. Hence, the locations of the selected vortex geometries depend on the initial guess $\psi^{(0)}$ in Alg. 1. For a discrete time sequence of flow data sets one can ensure time coherence by initializing Alg. 1 with the time advected minimizer from the previous time step. For an example see the video of the Hummingbird flapping.

5.2 Initial Data for Incompressible Schrödinger Flow

With the fluid state represented in a Clebsch map ψ one might ask what time evolution should ψ obey for fluid simulation? A straightforward approach is to have ψ passively advected $\frac{\partial}{\partial t}\psi + d\psi(\mathbf{u}) = 0$ by its associated pressure projected velocity $\mathbf{u}$. This yields dynamics in which all vortex lines are advected by $\mathbf{u}$ and therefore $\mathbf{u}$ satisfies the Euler equation. Unfortunately this type of Eulerian-Lagrangian approach leads to chaotic mixing with initially small $\|d\psi\|$ and $\|ds\|$ amplified unboundedly. The necessary resampling of ψ involves averaging based on the grid resolution and hence erases velocity information encoded in the derivative of ψ (an even more severe version of the numerical diffusion seen in backward advection schemes [Fedkiw et al. 2001; Stam 1999]).

Surprisingly there is a variant for the evolution of ψ called incompressible Schrödinger flow (ISF) [Chern et al. 2016] which avoids Lagrangian chaos by including a small amount of the Dirichlet energy of the Clebsch map s in the energy. The resulting time evolution is Hamiltonian, implying that $\|d\psi\|$ and $\|ds\|$ are both uniformly bounded in time under ISF.

One drawback of the approach in [Chern et al. 2016] was that the initial ψ_0 data had to be carefully handcrafted for every simulation. While the authors mention the possibility of using filament sets [Weißmann et al. 2014] the approximation quality of such an approach is unclear. Our method directly searches for an optimally

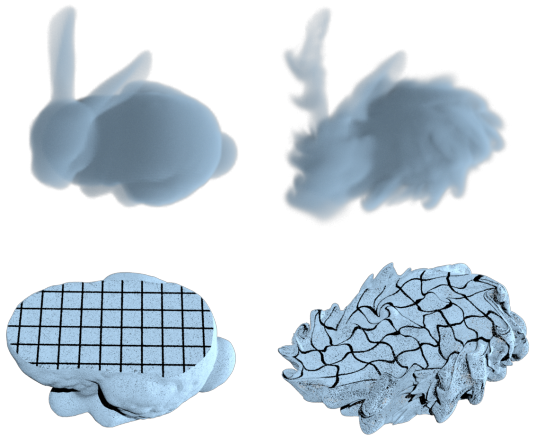

Fig. 10. Fluid simulation by ISF with initial ψ_0 derived from a general $\mathbf{u}_0$. Here $\mathbf{u}_0$ is the divergence-free projection of a rigid-body rotation interior to the bunny (and zero velocity else). The corresponding initial vorticity is uniform in the bunny driving a rigid rotation, and has a concentrated vortex sheet on the surface giving rise to a Kelvin-Helmholtz instability. See the video at 01:16.

close representation of an arbitrary input field which can then serve as initial condition for ISF (Fig. 10).

5.3 Flow Processing

A vorticity Clebsch map encodes vortex lines as level sets of s. This makes it easy to "rearrange" the vortex lines through manipulating ψ (and thus s), enabling a whole new class of flow processing approaches.

An example of this is post-composition $\tilde{\psi} := \Xi \circ \psi$ with a map $\Xi\colon \mathbb{S}^3 \to \mathbb{S}^3$. Here we only consider maps Ξ with a corresponding $\xi\colon \mathbb{S}^2 \to \mathbb{S}^2$ satisfying $\pi \circ \Xi = \xi \circ \pi$. Under this assumption the modification of ψ corresponds to a modification $\tilde{s} := \xi \circ s$ of s.

An interesting example of $\xi\colon \mathbb{S}^2 \to \mathbb{S}^2$ "wraps" $\mathbb{S}^2$ multiple times over itself (App. E gives an explicit example of a multiple, branched covering). The number of times ξ wraps around $\mathbb{S}^2$ is called the *mapping degree*

$$m_\xi := \frac{\mathrm{Area}(\xi(\mathbb{S}^2))}{\mathrm{Area}(\mathbb{S}^2)}$$

where Area(.) computes the signed area with multiplicity. The mapping degree m_ξ shows up as a renormalization when reading off the processed velocity $\tilde{\eta} := \frac{\hbar}{m_\xi}\langle d\tilde{\psi}, i\psi\rangle$, ensuring the overall vorticity flux remains unchanged.

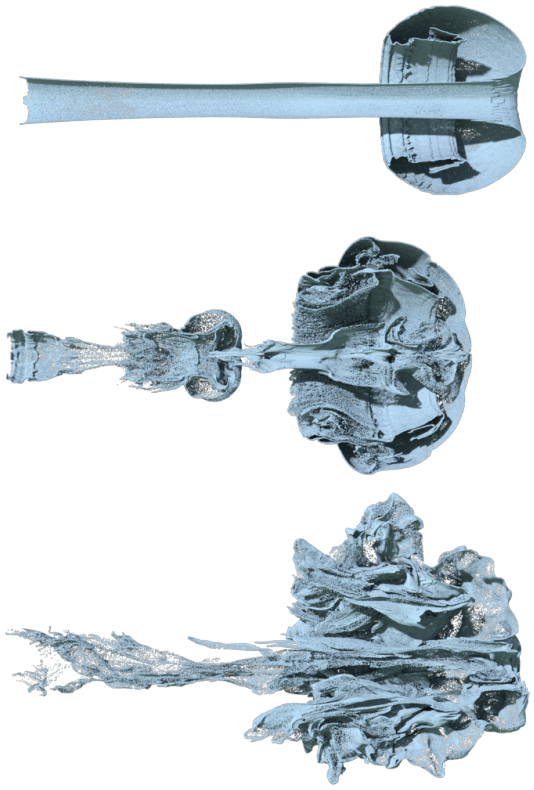

Fig. 11. Top: Jet simulation using RK4 backward advection and MacCormack time marching. Middle: the same scheme with an additional flow processing step (Sec. 5.3) at an earlier time to represent finer vortical structures aligned with coarse ones. Bottom: the result of adding isotropic turbulence at an earlier step. All simulations on a $256 \times 128 \times 128$ grid while the Clebsch map finder used a resolution of $64 \times 32 \times 32$ (due to performance reasons) which was interpolated to the high resolution grid. See the video at 00:54.

Algorithm 2 Flow Processing

Input: $\eta_0, \hbar, \Xi$
  $\psi \leftarrow$ Find Clebsch map $(\eta_0, \hbar)$
  $\tilde{\psi} \leftarrow \Xi \circ \psi$
  $\tilde{\eta} \leftarrow \frac{\hbar}{m_\xi}\langle d\tilde{\psi}, i\tilde{\psi}\rangle$
  return PressureProject $(\tilde{\eta})$

In Fig. 11 we modify a velocity field at a single frame of a jet simulation by Alg. 2 with a degree four map $\xi = \xi_{\text{mobius}} \circ \xi_{\text{tetra}}$ (App. E). This particular modification cascades each vortex into four finer ones, and then concentrates the vorticity in the fine scale as shown in Fig. 12. Effectively this adds small eddies aligned with the large scale vortex direction in a vorticity conserving fashion (Fig. 11, middle) as predicted by the Kelvin-Helmholtz instability that is otherwise not captured at the current resolution of the simulation (top) or by adding isotropic turbulence [Bridson et al. 2007] (bottom).

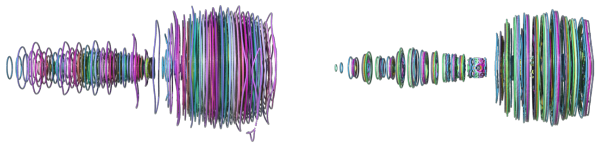

Fig. 12. A Clebsch map s (left) obtained from the velocity data of a jet simulation is modified (right) by $\tilde{s} = \xi \circ s$ with a sphere map $\xi\colon \mathbb{S}^2 \to \mathbb{S}^2$ that effectively concentrates vorticity. It is used for Fig. 11 middle.

6 OUTLOOK

The appearance of spherical Clebsch maps in the recently introduced method for incompressible fluid simulation using wave functions ψ [Chern et al. 2016] motivated us to study Clebsch maps more closely and seek for a method to find such maps for a given flow field.

While there are theoretical limitations to such spherical Clebsch maps these do not stand in the way of producing useful approximations (Sec. 4). Because these Clebsch maps encode the velocity resp. vorticity field they contain all information about the field in a novel form. This can be useful for visualization and flow processing, not just simulation (Sec. 5).

There are a number of theoretical questions which deserve further study. Can we develop methods which are guaranteed to find a global minimum of Eq. (7)? While we can't expect to find exact Clebsch maps for generic flow fields (Sec. 4.3) it would be highly

desirable to derive approximation bounds. How well *can* we do? The role of $\hbar$ also deserves further scrutiny. How does the approximation quality depend on this parameter? For applications which require only a vorticity Clebsch map, is there a simpler algorithm?

At present a signficant obstacle to practical deployment is the computational cost of minimizing Eq. (7). So far we have used only a semi-implicit gradient descent approach. The biggest challenge is the efficient solution of (screened) Poisson problems involving the connection Laplacian. In our experience these systems can be very ill conditioned (and more so as the grid is refined). An effective preconditioning strategy, which we wish to develop in future work, appears essential to make the method practical.

We are most excited about developing the idea of flow processing further. Manipulating the flow field through manipulation of its Clebsch representation as we suggested in Alg. 2 and App. E, ensures important properties such as conservation of the amount of vorticity in a given simulation cell. It also ensures alignment of the finer level structures with the coarser vorticity [Pfaff et al. 2010], something often lost in purely local manipulations.

ACKNOWLEDGMENTS

This work was supported in part by the DFG Collaborative Research Center TRR 109 "Discretization in Geometry and Dynamics." Additional support was provided by SideFX software. The Bunny model is courtesy Stanford Computer Graphics laboratory.

REFERENCES

Ralph Abraham, Jerrold E. Marsden, and Tudor Ratiu. 2001. *Manifolds, Tensor Analysis and Applications.* Number 75 in Appl. Math. Sci. Springer.

Alexis Angelidis and Karan Singh. 2007. Kinodynamic Skinning using Volume-Preserving Deformations. In *Proc. Symp. Comp. Anim.* ACM SIGGRAPH/Eurographics, 129–140.

Vladimir I. Arnold and Boris A. Khesin. 1998. *Topological Methods in Hydrodynamics.* Springer.

Axel Brandenburg. 2010. Magnetic Field Evolution in Simulations with Euler Potentials. *MNRAS* 401, 1 (2010), 347–354.

Francis P. Bretherton. 1970. A Note on Hamilton's Principle for Perfect Fluids. *J. Fl. Mech.* 44, 1 (1970), 19–31.

Robert Bridson, Jim Houriham, and Marcus Nordenstam. 2007. Curl-Noise for Procedural Flow. *ACM Trans. Graph.* 26, 3, Article 46 (2007), 3 pages.

Michael A. Bush, Katelyn R. French, and Joseph R. H. Smith. 2017. Total Linking Numbers of Torus Links and Klein Links. *RH Underg. Math. J.* 15, 1 (2017), 72–92.

Carlos Cartes, Miguel D. Bustamente, and Marc E. Brachet. 2007. Generalized Eulerian-Lagrangian Description of Navier-Stokes Dynamics. *Phys. Fluids* 19, 7, Article 077101 (2007), 7 pages.

Augustin-Louis Cauchy. 1815. Théorie de la Propagation des Ondes a la Surface d'un Fluide Pesant d'une Profondeur Indéfinie. In *Oeuvres Complètes d'Augustin Cauchy.* Vol. 1. Imprimerie Royale. Presented to the French Academy in 1815 (publ. 1827).

Albert Chern. 2017. *Fluids Dynamics with Incompressible Schrödinger Flow.* Ph.D. Dissertation. Caltech.

Albert Chern, Felix Knöppel, Ulrich Pinkall, Peter Schröder, and Steffen Weißmann. 2016. Schrödinger's Smoke. *ACM Trans. Graph.* 35, 4, Article 77 (2016), 13 pages.

Alfred Clebsch. 1859. Ueber die Integration der hydrodynamischen Gleichungen. *J. Reine Angew. Math.* 56 (1859), 1–10. English translation by David H. Delphenich, http://www.neo-classical-physics.info/uploads/3/4/3/6/34363841/clebsch_-_clebsch_variables.pdf.

Peter Constantin. 2001a. An Eulerian-Lagrangian Approach for Incompressible Fluids: Local Theory. *J. Amer. Math. Soc.* 14, 2 (2001), 263–278.

Peter Constantin. 2001b. An Eulerian-Lagrangian Approach to the Navier-Stokes Equations. *Comm. Math. Phys.* 216, 3 (2001), 663–686.

Keenan Crane, Fernando de Goes, Mathieu Desbrun, and Peter Schröder. 2013. Digital Geometry Processing with Discrete Exterior Calculus. In *Courses.* ACM SIGGRAPH.

David Delphenich. 2017. The Role of Integrability in a large Class of Physical Systems. (March 2017). https://arxiv.org/abs/1210.4976.

Mathieu Desbrun, Eva Kanso, and Yiying Tong. 2008. Discrete Differential Forms for Computational Modeling. In *Discrete Differential Geometry,* Alexander I. Bobenko, Peter Schröder, John M. Sullivan, and Günther M. Ziegler (Eds.). Oberwolfach Seminars, Vol. 38. Birkhäuser Verlag.

John A. Ekaterinaris and Lewis B. Schiff. 1990. Vortical Flows over Delta Wings and Numerical Prediction of Vortex Breakdown. In *28th Aerosp. Sc. Meet.* AIAA, 90–102.

Leonhard Euler. 1757. Continuation des Recherches sur la Théorie due Mouvement des Fluides. *Mémoires de l'Académie des Sciences de Berlin* (1757), 316–361. Presented to the Berlin Academy in 1755. For an English translation see [Frisch 2008].

Ronald Fedkiw, Jos Stam, and Henrik Wann Jensen. 2001. Visual Simulation of Smoke. In *Proc. ACM/SIGGRAPH Conf.* ACM, 15–22.

Theodore Frankel. 2012. *The Geometry of Physics: An Introduction* (3rd ed.). Cam. U. P.

Uriel Frisch. 2008. Translation of Leonhard Euler's: General Principles of the Motion of Fluids. (2008). https://arxiv.org/abs/0802.2383v1.

Uriel Frisch and Barbara Villone. 2014. Cauchy's almost Forgotten Lagrangian Formulation of the Euler Equation for 3D Incompressible Flow. *Eu. Phy. J. H* 39, 3 (2014), 325–351.

Carl Friedrich Gauß. 1833. *Zur mathematischen Theorie der elektrodynamischen Wirkungen.* Werke, Vol. 5. Königliche Gesellschaft der Wissenschaften, Göttingen, Chapter Handschriftlicher Nachlass. Publ. 1877.

Hanjörg Geiges. 2006. *Handbook of Differential Geometry.* Vol. 2. North-Holland, Chapter Contact Geometry, 315–382.

C. Robin Graham and Frank S. Henyey. 2000. Clebsch Representation near Points where the Vorticity Vanishes. *Phys. Fluids* 12, 4 (2000), 744–746.

Jacques Hadamard. 1903. *Leçons sur la Propagation des Ondes et les Équations de l'Hydrodynamique.* Librairie Scientifique A. Hermann, Paris.

Pengyu He and Yue Yang. 2016. Construction of Initial Vortex-Surface Fields and Clebsch Potentials for Flows with High-Symmetry using First Integrals. *Phys. Fluids* 28, 3, Article 037101 (2016), 13 pages.

Heinz Hopf. 1931. Über die Abbildungen der Dreidimensionalen Sphäre auf die Kugelfläche. *Math. Ann.* 104, 1 (1931), 637–665.

John E. Hutchinson. 1991. Computing Conformal Maps and Minimal Surfaces. In *Theoretical and Numerical Aspects of Geometric Variational Problems (Proc. Cent. Math. and Appl.),* Gerd Dziuk, Gerhard Huisken, and John Hutchinson (Eds.), Vol. 26. 140–161.

Jinhee Jeong and Fazle Hussain. 1995. On the Identification of a Vortex. *J. Fl. Mech.* 285 (1995), 69–94.

Ming Jiang, Raghu Machiraju, and David Thompson. 2005. *The Visualization Handbook.* Elsevier, Chapter Detection and Visualization of Vortices, 295–312.

P. Robert Kotiuga. 1991. Clebsch Potentials and the Visualization of Three-Dimensional Solenoidal Vector Fields. *IEEE Trans. Magn.* 27, 5 (1991), 3986–3989.

Evgenii A. Kuznetsov and Aleksandr V. Mikhailov. 1980. On the Topological meaning of Canonical Clebsch Variables. *Phys. Lett. A* 77, 1 (1980), 37–38.

David W. Lyons. 2003. An Elementary Introduction to the Hopf Fibration. *Math. Mag.* 76, 2 (2003), 87–98.

Nicolas Magot and Alexander Zvonkin. 2000. Belyi Functions for Archimedean Solids. *Disc. Math.* 217, 1–3 (2000), 249–271.

John W. Milnor. 1965. *Topology from the Differentiable Viewpoint.* U. P. Virginia.

H. Keith Moffatt. 1969. The Degree of Knottedness of Tangled Vortex Lines. *J. Fluid Mech.* 35, 1 (1969), 117–129.

Tobias Pfaff, Nils Thuerey, Jonathan Cohen, Sarah Tariq, and Markus Gross. 2010. Scalable Fluid Simulation using Anisotropic Turbulence Particles. *ACM Trans. Graph.* 29, 6, Article 174 (2010), 8 pages.

Henri Poincaré. 1890. Sur le Problème des Trois Corps et les Équations de la Dynamique. *Acta Math.* 13 (1890), 1–270.

Yan Ren, Haibo Dong, Xinyan Deng, and Bret Tobalske. 2016. Turning on a Dime: Asymmetric Vortex Formation in Hummingbird Maneuvering Flight. *Phys. R. Fl.* 1, 5, Article 050511 (2016), 3 pages.

Philip Geoffrey Saffman. 1992. *Vortex Dynamics.* Cam. U. P.

Jos Stam. 1999. Stable Fluids. In *Proc. ACM/SIGGRAPH Conf.* ACM, 121–128.

Wolfram von Funck, Holger Theisel, and Hans-Peter Seidel. 2006. Vector Field Based Shape Deformations. *ACM Trans. Graph.* 25, 3 (2006), 1118–1125.

Heinrich Weber. 1868. Ueber eine Transformation der hydrodynamischen Gleichungen. *J. Reine Angew. Math.* 68 (1868), 286–292.

Steffen Weißmann, Ulrich Pinkall, and Peter Schröder. 2014. Smoke Rings from Smoke. *ACM Trans. Graph.* 33, 4, Article 140 (2014), 8 pages.

Peter Woit. 2017. Quantum Theory, Groups and Representations: An Introduction. (Jan. 2017). https://www.math.columbia.edu/~woit/QM/qmbook.pdf.

Lodewijk Woltjer. 1958. A Theorem on Force-Free Magnetic Fields. *Proc. Nat. Acad. Sci.* 44, 6 (1958), 489–491.

Yue Yang and Dale I. Pullin. 2010. On Lagrangian and Vortex-Surface Fields for Flows with Taylor-Green and Kida-Pelz Initial Conditions. *J. Fl. Mech.* 661 (2010), 446–481.

Vladimir E. Zakharov and Evgenii A. Kuznetsov. 1997. Hamiltonian Formalism for Nonlinear Waves. *Physics-Uspekhi* 40, 11 (1997), 1087–1116.

A VELOCITY FROM VORTICITY CLEBSCH MAPS

Here we prove a simplified version of Thm. 3.1, *i.e.*, we show that a spherical Clebsch map s for ω can be lifted to a Clebsch map ψ for a velocity field η with $d\eta = \omega$, under the additional assumption that M is star-shaped with respect to the origin. Note that the theorem remains valid for multiply connected domains, but the proof becomes more elaborate [Chern 2017].

Choose any $\psi_0 \in \mathbb{S}^3$ with $\bar{\psi}_0 i \psi_0 = s(\mathbf{0})$. Define $s_{\mathbf{r}} \colon [0,1] \to \mathbb{S}^2$ by $s_{\mathbf{r}}(t) = s(t\mathbf{r})$ for $\mathbf{r} \in M$. Similarly, for $\mathbf{r} \in M$ define $\psi_{\mathbf{r}} \colon [0,1] \to \mathbb{S}^3$ as the solution of the initial value problem

$$\frac{d}{dt}\psi_{\mathbf{r}} = -\frac{i}{2}\psi_{\mathbf{r}} \frac{d}{dt} s_{\mathbf{r}} \qquad \text{and} \qquad \psi_{\mathbf{r}}(0) = \psi_0.$$

Then

$$\frac{d}{dt}\left(\bar{\psi}_{\mathbf{r}} i \psi_{\mathbf{r}}\right) = \frac{d}{dt} s_{\mathbf{r}} \qquad \text{and} \qquad \left(\bar{\psi}_{\mathbf{r}} i \psi_{\mathbf{r}}\right)(0) = s_{\mathbf{r}}(0)$$

and therefore $\bar{\psi}_{\mathbf{r}} i \psi_{\mathbf{r}} = s_{\mathbf{r}}$. Now we can define $\phi \colon M \to \mathbb{S}^3$ by $\phi(\mathbf{r}) \coloneqq \psi_{\mathbf{r}}(1)$. We have already achieved $\bar{\phi} i \phi = s$. Defining $\tilde{\eta} \coloneqq \hbar \langle d\phi, i\phi \rangle$ we get $d\tilde{\eta} = d\eta$. Since M is simply connected we therefore can find a function $\theta \colon M \to \mathbb{R}$ (unique up to an additive constant) such that $d\theta = \eta - \tilde{\eta}$. Then $\psi \coloneqq e^{i\theta}\phi$ is the Clebsch map for η that we were looking for.

B ORTHOGONAL SPLITTING OF $\nabla^{\eta_0}\psi$

Since $\mathbb{S}^3$ can be represented through unit quaternions we will avail ourselves of quaternionic algebra in this section to derive the orthogonal splitting of $\nabla^{\eta_0}\psi$.

Using the definition $\nabla^{\eta_0} = d - i\frac{\eta_0}{\hbar}$ and Eq. (8) of [Chern et al. 2016] we see that

$$-\bar{\psi} i \nabla^{\eta_0}\psi = \tfrac{1}{\hbar}(\eta - \eta_0) - \tfrac{1}{2} ds$$

and thus

$$\nabla^{\eta_0}\psi = \tfrac{1}{\hbar}(\eta - \eta_0) i\psi - \tfrac{1}{2} i\psi ds, \tag{9}$$

where we used $\bar{\psi}\psi = |\psi|^2 = 1$. Since the first summand is a real multiple of $i\psi$ it takes values in $\mathbb{C}\psi = \operatorname{Span}\{\psi, i\psi\}$. That the second summand takes values in $(\mathbb{C}\psi)^{\perp} = \mathbb{C}j\psi$ follows from

$$\langle \psi, i\psi ds \rangle = \operatorname{Re}(\bar{\psi} i \psi ds) = \operatorname{Re}(s ds) = -\tfrac{1}{2} d|s|^2 = 0,$$
$$\langle i\psi, i\psi ds \rangle = \operatorname{Re}(ds) = 0.$$

Here we used $\langle a, b \rangle = \operatorname{Re}(\bar{a}b)$.

To summarize, Eq. (9) is the orthogonal splitting of $\nabla^{\eta_0}\psi$ onto the subspaces $\mathbb{C}\psi$ and $\mathbb{C}j\psi$.

C IMPLEMENTATION DETAILS FOR E^ϵ

In this section we give all the necessary details for the numerical implementation of Alg. 1. Above we used unit quaternions ψ to realize $\mathbb{S}^3$. For the derivations in this section it turns out to be convenient to realize quaternions $\psi \in \mathbb{H}$ as elements $(\psi_1, \psi_2)^{\intercal} \in \mathbb{C}^2$ through the relation $\psi = \psi_1 + \psi_2 j$. For numerical libraries which support a complex data type the final expressions translate directly. If only real data types are supported, pairs of complex numbers become elements of $\mathbb{R}^4$ and a complex entry $x + iy = z \in \mathbb{C}$ in the matrix is realized as a real 2×2 matrix $\left(\begin{smallmatrix} x & -y \\ y & x \end{smallmatrix}\right)$.

Discretization. For the energy, Eq. (7), we use a standard discrete exterior calculus (DEC) [Crane et al. 2013; Desbrun et al. 2008] discretization. It requires a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{V}$ is the set of vertices and $\mathcal{E}$ the set of directed edges. All examples in this paper take $\mathcal{G}$ as the 1-skeleton of a 2D or 3D lattice. Each vertex $\mathrm{i} \in \mathcal{V}$ has a vertex weight $w_{\mathrm{i}} > 0$ and each edge $\mathrm{ij} \in \mathcal{E}$ an edge weight $w_{\mathrm{ij}} > 0$. Functions (0-forms) are discretized as values per vertex, *e.g.*, $\psi_{\mathrm{i}} = \psi(\mathbf{r}_{\mathrm{i}})$ while 1-forms ξ are discretized as values per directed edge $\xi_{\mathrm{ij}} \coloneqq \int_{\mathrm{ij}} \xi$.

The vertex weights and edge weights are designed to give the approximation

$$\int_M |f|^2 \approx \sum_{\mathrm{i} \in \mathcal{V}} w_{\mathrm{i}} |f_{\mathrm{i}}|^2, \quad \int_M |\xi|^2 \approx \sum_{\mathrm{ij} \in \mathcal{E}} w_{\mathrm{ij}} |\xi_{\mathrm{ij}}|^2.$$

On a regular 3D lattice $w_{\mathrm{i}} = V_{\mathrm{i}}$ and $w_{\mathrm{ij}} = \frac{A_{\mathrm{ij}}}{\ell_{\mathrm{ij}}}$ where V_{i} is the volume of the dual cell surrounding $\mathbf{r}_{\mathrm{i}}$, A_{ij} the area of the face dual to ij, and ℓ_{ij} the length of the edge $\mathbf{r}_{\mathrm{j}} - \mathbf{r}_{\mathrm{i}}$.

Covariant Derivative. Given a $\mathbb{C}^2$-valued function $\psi = (\psi_1, \psi_2)^{\intercal}$ on vertices and a $\mathbb{R}$-valued 1-form η_0 on edges, the $\mathbb{C}^2$-valued 1-form $\nabla^{\eta_0}\psi$ is given [Weißmann et al. 2014, Eq. 10] by the difference

$$(\nabla^{\eta_0}\psi)_{\mathrm{ij}} = e^{-i\frac{(\eta_0)_{\mathrm{ij}}}{2\hbar}} \psi_{\mathrm{j}} - e^{i\frac{(\eta_0)_{\mathrm{ij}}}{2\hbar}} \psi_{\mathrm{i}}.$$

Projection Operators. Using $(\psi_1, \psi_2)^{\intercal}$ as our representation the projectors $P_{\mathbb{C}\psi_{\mathbf{r}}}$ resp. $P_{\mathbb{C}j\psi_{\mathbf{r}}}$ can be expressed in terms of Pauli matrices $\sigma_i \in \mathbb{C}^{2\times 2}$

$$\sigma_1 \coloneqq \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad \sigma_2 \coloneqq \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad \sigma_3 \coloneqq \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

Even though the projectors appear to be functions of ψ, they are in fact only functions of $s = \bar{\psi} i \psi = s_1 i + s_2 j + s_3 k$ where

$$s_1 = |\psi_1|^2 - |\psi_2|^2, \quad s_2 = 2\operatorname{Re}(i\bar{\psi}_1\psi_2), \quad s_3 = 2\operatorname{Re}(\bar{\psi}_1\psi_2).$$

The coefficients are given by the Pauli spin vector $z = (z_1, z_2, z_3)^{\intercal} \in \mathbb{R}^3$ with

$$z_i \coloneqq \begin{bmatrix} \bar{\psi}_1 & \bar{\psi}_2 \end{bmatrix} \sigma_i \begin{bmatrix} \psi_1 \\ \psi_2 \end{bmatrix}, \qquad i = 1, 2, 3$$

as

$$s_1 = z_3, \quad s_2 = -z_2, \quad s_3 = z_1.$$

To find the $\mathbb{C}^{2\times 2}$ representation of $P_{\mathbb{C}\psi_{\mathbf{r}}}$ we note that the underlying subspace is spanned by complex multiples of $(\psi_1, \psi_2)^{\intercal}$ and hence

$$P_{\mathbb{C}\psi_{\mathbf{r}}} = \begin{bmatrix} \psi_1 \\ \psi_2 \end{bmatrix}_{\mathbf{r}} \begin{bmatrix} \bar{\psi}_1 & \bar{\psi}_2 \end{bmatrix}_{\mathbf{r}} = \begin{bmatrix} |\psi_1|^2 & \psi_1\bar{\psi}_2 \\ \bar{\psi}_1\psi_2 & |\psi_2|^2 \end{bmatrix}_{\mathbf{r}}$$
$$= \tfrac{1}{2}\left(I + z_1\sigma_1 + z_2\sigma_2 + z_3\sigma_3\right),$$

where I is the $\mathbb{C}^{2\times 2}$ identity matrix. Due to their orthogonality, $P_{\mathbb{C}j\psi_{\mathbf{r}}}$ is just the residual

$$P_{\mathbb{C}j\psi_{\mathbf{r}}} = I - P_{\mathbb{C}\psi_{\mathbf{r}}} = \tfrac{1}{2}\left(I - z_1\sigma_1 - z_2\sigma_2 - z_3\sigma_3\right).$$

In the implementation it is convenient to combine both projectors

$$P^\epsilon \coloneqq P_{\mathbb{C}\psi} + \sqrt{\epsilon} P_{\mathbb{C}j\psi}$$

and rewrite the integrand of E^ϵ as

$$|\nabla^{\eta_0}\psi|^2_\epsilon = |P^\epsilon(\nabla^{\eta_0}\psi)|^2.$$

Since $\nabla^{\eta_0}\psi$ is a 1-form on edges, we need P^ϵ_{ij} for $s_{ij} = \frac{s_i+s_j}{|s_i+s_j|}$ the edge midpoint value of s.

The Discrete Energy. With these definitions we arrive at the discrete energy in terms of the vector $\boldsymbol{\psi} = (\psi_i)_{i\in\mathcal{V}}$

$$\mathsf{E}^\epsilon(\boldsymbol{\psi}) = \sum_{ij\in\mathcal{E}} w_{ij}\left|P^\epsilon_{ij}\left(\left(\nabla^{\eta_0}\boldsymbol{\psi}\right)_{ij}\right)\right|^2 = \bar{\boldsymbol{\psi}}^\mathsf{T}\mathsf{L}\boldsymbol{\psi}.$$

Where the quadratic form L is a $|\mathcal{V}|\times|\mathcal{V}|$ sparse matrix of 2×2 complex blocks. Each edge $ij\in\mathcal{E}$ gives rises to four blocks which are accumulated into the global L, summing over all edges

$$\mathsf{L}_{ii} = w_{ij}(P^\epsilon_{ij})^2 \qquad \mathsf{L}_{ij} = -w_{ij}e^{-i\frac{(\eta_0)_{ij}}{\hbar}}(P^\epsilon_{ij})^2$$
$$\mathsf{L}_{ji} = -w_{ij}e^{i\frac{(\eta_0)_{ij}}{\hbar}}(P^\epsilon_{ij})^2 \qquad \mathsf{L}_{jj} = w_{ij}(P^\epsilon_{ij})^2.$$

Minimization. To minimize the discrete energy for a given ϵ, we use a semi-implicit gradient descent step followed by pointwise normalization of $\boldsymbol{\psi}$. Each gradient step is the solution $\boldsymbol{\psi}^{(k+1)}$ to

$$(M_\mathcal{V} + \Delta t\mathsf{L}^{(k)})\boldsymbol{\psi}^{(k+1)} = M_\mathcal{V}\boldsymbol{\psi}^{(k)},$$

with $M_\mathcal{V}$ the vertex mass matrix and step size $\Delta t > 0$. We typically take 20 steps with $0.1 < \Delta t < 1$ before decreasing ϵ (Alg. 1).

D HELICITY

The helicity of a divergence-free vector field was first introduced in [Woltjer 1958] and received its name in [Moffatt 1969]. The result that helicity vanishes for fields coming from a classical Clebsch map is due to [Bretherton 1970]. Here is a proof in the language of differential forms.

As in section Sec. 4.4 we consider an exact 2-form $\omega = d\eta$ on $\mathbb{R}^3$ which vanishes outside of a large closed ball M. In terms of η and ω the helicity is given by (*cf.* Eq. (8))

$$\mathcal{H}(\omega) = \int_M \eta\wedge\omega.$$

Suppose ω comes from a classical Clebsch map, *i.e.*, $\eta = \lambda d\mu - d\phi$, $\omega = d\lambda\wedge d\mu$ then

$$\mathcal{H}(\omega) = \int_M -d\phi\wedge d\lambda\wedge d\mu = \int_M -d(\phi\, d\lambda\wedge d\mu) = \int_{\partial M} -\phi\omega = 0,$$

proving Thm. 4.1.

The result that helicity is quantized for fields coming from $\mathbb{S}^2$-valued Clebsch maps is due to [Kuznetsov and Mikhailov 1980]. They assume fields that tend to zero at infinity, instead we want a version that also works in a bounded region. Additionally they assume that the Clebsch map itself has a limit at infinity. We now turn to the proof of Thm. 4.2 which does not need such assumptions.

Since the pullback ω of the area form of $\mathbb{S}^2$ vanishes on ∂M, the derivative of the map

$$\hat{s} = s|_{\partial M}\colon \partial M\to\mathbb{S}^2$$

does not have full rank anywhere. Sard's Lemma [Milnor 1965, §2] then asserts that there is a point $p\in\mathbb{S}^2$ that is not in the image of $\hat{s}$. Therefore we can smoothly homotope $\hat{s}$ to a constant map from ∂M to $\mathbb{S}^2$. Thus $s\colon M\to\mathbb{S}^2$ can be extended to a continuous map $\tilde{s}\colon\mathbb{R}^3\cup\{\infty\}\to\mathbb{S}^2$ and the rest of the argument is the same as in [Kuznetsov and Mikhailov 1980].

E RATIONAL MAPS FOR FLOW PROCESSING

Any map $\xi\colon\mathbb{S}^2\to\mathbb{S}^2$ can be written as a function on the extended complex plane by identifying $\mathbb{S}^2\cong\mathbb{C}\cup\{\infty\}$ through stereographic projection $s\mapsto\zeta = \frac{s_2}{1-s_1} + \frac{s_3}{1-s_1}i$ and its inverse $s = (1+|\zeta|)^{-1}(-1+|\zeta|, 2\operatorname{Re}\zeta, -2\operatorname{Re}(i\zeta))$. Similarly, each map $\Xi\colon\mathbb{S}^3\to\mathbb{S}^3$ is a map $\Xi\colon\mathbb{C}^2\to\mathbb{C}^2$, $\Xi(\zeta_1,\zeta_2) = (Q_1(\zeta_1,\zeta_2), Q_2(\zeta_1,\zeta_2))$ where $\zeta_1,\zeta_2\in\mathbb{C}$. Thus designing a sphere map for Sec. 5.3 amounts to finding functions ξ on the extended complex plane and Ξ on $\mathbb{C}^2$ so that $\pi\circ\Xi = \xi\circ\pi$.

For our purpose good candidates for ξ are the rational functions $\xi(\zeta) = P_1(\zeta)(P_2(\zeta))^{-1}$ where P_1, P_2 are polynomials without common divisor. Rational functions give all the conformal maps $\mathbb{S}^2\to\mathbb{S}^2$. The mapping degree m_ξ is given by $\max\{\deg(P_1),\deg(P_2)\}$. For each rational function ξ a function $\Xi = (Q_1, Q_2)$ with $\pi\circ\Xi = \xi\circ\pi$ is easily found with both $Q_1, Q_2\colon\mathbb{C}^2\to\mathbb{C}$ homogeneous polynomials: $\xi(\zeta) = Q_1(\zeta,1)(Q_2(\zeta,1))^{-1}$. Simple examples are scaling, $\xi_{\text{mobius}}(\zeta) = a\zeta$, corresponding to $\Xi_{\text{mobius}}(\zeta_1,\zeta_2) = (a\zeta_1,\zeta_2)$, and squaring $\xi_{\text{sq}}(\zeta) = \zeta^2$ with $\Xi_{\text{sq}}(\zeta_1,\zeta_2) = (\zeta_1^2,\zeta_2^2)$. Belyi gave a gallery of functions that have particular discrete symmetries. For instance,

$$\Xi_{\text{tetra}}(\zeta_1,\zeta_2) = \begin{pmatrix}2\sqrt{2}\zeta_2^4 + 8\zeta_1^3\zeta_2\\ 2\sqrt{2}\zeta_1^4 - 8\zeta_1\zeta_2^3\end{pmatrix}$$

gives ξ_{tetra} that maps the four vertices of a tetrahedron on $\mathbb{S}^2$ to the same point [Magot and Zvonkin 2000]. The inset figure shows the result of applying ξ_{tetra} to a world map on $\mathbb{S}^2$.

To randomize the choice of the scrambling map with rotational symmetry, given a desired mapping degree m, one can choose $\Xi = (Q_1, Q_2)$ with

$$Q_j(\zeta_1,\zeta_2) = \sum_{k=0}^{m}\frac{Z_{jk}}{\sqrt{k!(m-k)!}}\zeta_1^k\zeta_2^{m-k}$$

where $Z_{jk}\in\mathbb{C}$ are independent, normally distributed random numbers for $j = 1, 2$ and $k = 0\ldots m$ [Woit 2017, Sec. 8.2].