Synchronizing Fields with Singularities

NATALIA PACHECO-TALLAJ, Massachusetts Institute of Technology, USA
MATTÉO COUPLET and EDWARD CHIEN, Boston University, USA
DAVID PALMER, Institute of Science and Technology Austria, Austria

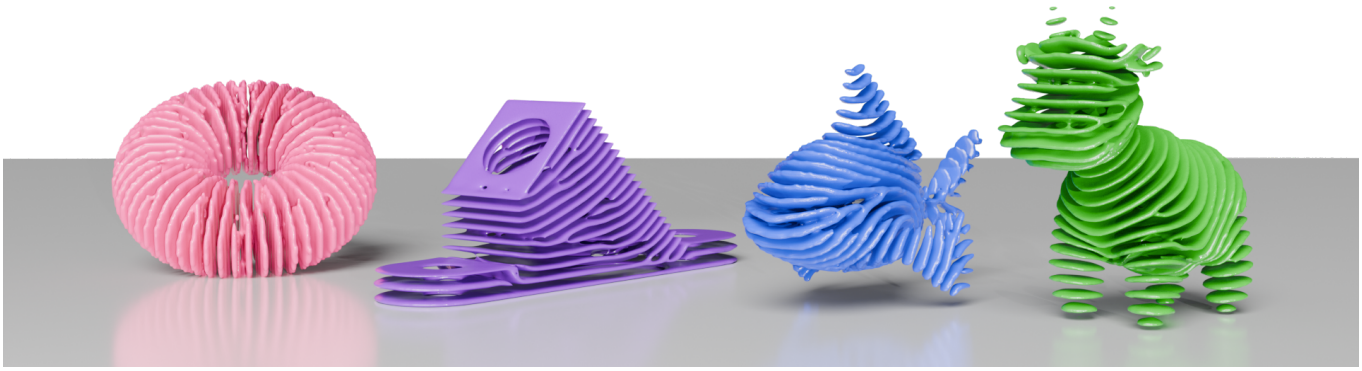

Fig. 1. Many geometry processing tasks can be formulated as synchronization problems. Here, U(1) synchronization over a volumetric domain produces evenly spaced volumetric stripe patterns with explicit boundary control, e.g., for curved-layer 3D printing.

A variety of problems in geometry processing boil down to finding the most parallel field relative to a connection. Instances of this prototypical problem show up in computing direction fields and stripe patterns, quadrilateral meshing, and visualization of fluid flows. When the class of allowed fields includes those with topological defects, a relaxation is required to make the problem well-posed. We observe that these problems can be viewed as *synchronization problems*, which admit a natural semidefinite relaxation. We propose a unified method of solving all these problems via the efficient Burer-Monteiro factorization method. Geometrically, this amounts to lifting the field values to a higher-dimensional manifold, naturally resolving the singular nature of defects. Practically, we show that our convex relaxation method achieves better and more reliable optima than previous work employing alternative relaxations.

CCS Concepts: • **Computing methodologies** → **Shape analysis**; • **Mathematics of computing** → **Semidefinite programming**; *Convex optimization*; *Nonconvex optimization*; *Algebraic topology*.

ACM Reference Format:
Natalia Pacheco-Tallaj, Mattéo Couplet, Edward Chien, and David Palmer. 2026. Synchronizing Fields with Singularities. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '26), July 19–23, 2026, Los Angeles, CA, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3799902.3811225

Authors' Contact Information: Natalia Pacheco-Tallaj, nataliap@mit.edu, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA; Mattéo Couplet, mcouplet@bu.edu; Edward Chien, edchien@bu.edu, Boston University, Boston, Massachusetts, USA; David Palmer, dpalmer@ist.ac.at, Institute of Science and Technology Austria, Klosterneuburg, Austria.

SIGGRAPH Conference Papers '26, Los Angeles, CA, USA
© 2026 Copyright held by the owner/author(s).
This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '26), July 19–23, 2026, Los Angeles, CA, USA*, https://doi.org/10.1145/3799902.3811225.

1 Introduction

Many optimization problems in graphics and geometry processing amount to finding a "most parallel" section of a bundle, i.e., one that minimizes a covariant Dirichlet energy with respect to a connection on the bundle. This prototypical problem often turns up when computing sections that feature *topological defects*, and it can be viewed as a way of *rounding* the curvature of the connection to a discrete set of defects.

Such problems, once discretized, can be recast as instances of a so-called *synchronization problem*. Synchronization problems are non-convex but admit relaxations to convex problems. Previous work in geometry processing has invariably resorted to either non-convex optimization or *spectral relaxation*. We propose to use the natural relaxation to a semidefinite program (SDP) instead. Semidefinite relaxation has several advantages. First, it is tighter than spectral relaxation (dropping fewer constraints), a fact we confirm empirically. Second, the solution to the relaxation provides both a feasible solution to the un-relaxed problem and a lower bound on its value, thus quantifying how close it is to global optimality. Third, combining relaxation with a randomized rounding procedure provides a whole family of near-optimal solutions with different topological defect configurations.

Finally, there is a convenient and efficient factorization algorithm, known as Burer-Monteiro factorization, that amounts to partially reversing the relaxation. Geometrically, the bundle fibers are lifted to higher-dimensional analogues. The factorization problem is then treated by manifold optimization on this lifted space. Remarkably, this low-dimensional non-convex problem is equivalent to the full SDP, without ever constructing the full semidefinite matrix. Moreover, in the lifted setting, topological defects are no longer an issue, making the energy well-defined. The factorization problems provide a sequence of well-defined geometric relaxations that have the advantage of implicit convexity.

We demonstrate the application of these ideas to several problems in geometry processing, from computing optimal direction fields, to stripe and quad patterns on surfaces and in volumes, to Clebsch map regression for visualizing vorticity in fluids. In comparison with the spectral relaxations of previous work, our SDP-based approach produces results with lower energy and better-placed defects.

Contributions. In summary, we:

- unify a host of geometry processing problems featuring topological defects in a common synchronization framework;
- solve them with SDP relaxation and connect its geometry to that of defects; and
- enable new applications such as volumetric stripe patterns by allowing flexible imposition of constraints in a convex relaxation framework.

2 Related Work

2.1 Synchronization

Group synchronization problems, also referred to as generalized phase retrieval, non-unique games [Bandeira et al. 2020], or the little Grothendieck problem [Bandeira et al. 2016], seek to recover unknown group elements from noisy or incomplete relative measurements. They are widely used in areas such as sensor network localization [Cucuringu et al. 2012a]; computer vision [Tzeneva 2011]; structural biology [Bandeira et al. 2020; Cucuringu et al. 2012b; Singer and Shkolnisky 2011]; robotics [Rosen et al. 2019]; and machine learning [Singer and Wu 2012].

Singer [2011] proposed spectral and SDP relaxations of angular synchronization (the U(1) case). Bandeira et al. [2013] established bounds on the performance of spectral relaxation. Bandeira et al. [2016] generalize this bound to synchronization over the higher-dimensional orthogonal and unitary groups. For U(1) synchronization, SDP relaxation with randomized rounding is known to give a $(\pi/4)$–approximation ratio in expectation under a positivity assumption on the cost matrix [Ben-Tal et al. 2003; So et al. 2007; Zhang and Huang 2006]. Bandeira et al. [2017] show tightness of an SDP relaxation, and Wang and Singer [2013] show exactness of an SDP relaxation for a "robust" version of the problem, both under statistical assumptions.

2.2 Semidefinite Programming

Semidefinite programs (SDPs) are efficiently-solvable convex cone optimization problems generalizing linear programs. We refer the reader to Blekherman et al. [2013]; Gärtner and Matousek [2012] for general background material.

SDPs have enjoyed outsized success in approximating the solutions to NP-hard problems, most famously in the case of MAX-CUT [Goemans and Williamson 1995]. MAX-CUT can be viewed as a synchronization problem over the group $\mathbb{Z}_2$, presaging the use of SDP relaxation for other synchronization problems.

In computer graphics, SDP relaxation powers methods for rigid [Dym and Lipman 2017; Maron et al. 2016] and non-rigid [Huang and Guibas 2013] shape matching, hexahedral mesh correction [Marschner et al. 2020], and polynomial geometry kernels [Marschner et al. 2021]. Closest to our work is that of Palmer et al. [2020], which also features ideas of topological defects and semidefinite programming.

The first provably polynomial-time algorithms for solving SDPs employed ellipsoid [Grötschel 1993; Khachiyan 1979] and interior-point methods [Alizadeh 1995; Nesterov and Nemirovskii 1994], extending earlier work on solving linear programs. The factorization technique of Burer and Monteiro [2003] (BM), further developed by Boumal [2016]; Boumal et al. [2020]; Cifuentes [2021]; Cifuentes and Moitra [2022], has led to an alternative polynomial-time method for solving SDPs that produces low-rank solutions. It is this method we exploit in our work. In doing so, we introduce a new perspective on the geometry of the BM factorizations in the case of synchronization-type problems.

2.3 Topological Defects

The theory of topological defects has a long history in condensed matter physics, tracing its roots to the seminal work of Landau [1936, 1937] on symmetry breaking and phase transitions. For an introduction, we recommend the excellent expository paper of Mermin [1979]. Canonical examples include magnetic monopoles in gauge fields [Castelnovo et al. 2008], vortices in superfluids, and disclinations in liquid crystals [Machon and Alexander 2016]. An explanation of phase transitions in terms of statistics of defects [Kosterlitz and Thouless 1973] won the Nobel Prize in physics in 2016. Other applications of defects range from quantum error correction [Kitaev 2003; Pretko et al. 2019] to biology [Copenhagen et al. 2021; Fard et al. 2022; Guillamat et al. 2022; Maroudas-Sacks et al. 2021]. This sampling only hints at the ubiquity of topological defects in the natural world.

In geometry processing, *optimization* of topological defects is important to applications including direction field design [Couplet et al. 2024; Diamanti et al. 2014; Knöppel et al. 2013; Palmer et al. 2020] and surface parameterization [Bommes et al. 2009; Fang et al. 2018]. These problems are intimately tied to quadrilateral and hexahedral meshing, in which the placement of disclination defects (also known as *singularities*) is of paramount importance. Optimization techniques have included spectral relaxation [Knöppel et al. 2013]; nonconvex optimization based on the Ginzburg-Landau functional [Beaufort et al. 2017; Couplet et al. 2024; Golovaty et al. 2021; Viertel and Osting 2019]; explicit optimization of defect positions [Farchi and Ben-Chen 2018; Ma et al. 2024; Myles and Zorin 2013]; and lifting [Palmer et al. 2024]. The above citations only sample the extensive literature [Bommes et al. 2013; Pietroni et al. 2022; Vaxman et al. 2016].

Periodic patterns have been used in graphics for texturing [Knöppel et al. 2015] and quad meshing [Fang et al. 2018; Ray et al. 2006]. Computational fabrication applications like curved-layer 3D printing [Zhang et al. 2022], machine knitting [Mitra et al. 2025; Narayanan et al. 2018], and topology optimization [Maestre et al. 2023; Stutz et al. 2022] also seek striped or laminar structures on curved surfaces or volumetric domains optimized for geometric, structural, or functional criteria. These prominently feature dislocation defects that must be placed intelligently for optimal object performance.

3 Geometric Synchronization Problems

We introduce the language of synchronization first in the prototypical case of optimal unit direction fields on surfaces. We then walk through several additional examples from graphics to build intuition before describing the general case.

3.1 Smooth Direction Field Design

A unit vector field $\mathbf{v}$ on a smooth Riemannian manifold M is a section of the tangent bundle $\mathbf{v} \in \Gamma(TM)$ that additionally satisfies the normalization constraint $|\mathbf{v}|^2 = 1$ everywhere. Equivalently, we can describe $\mathbf{v}$ as a section of the *unit tangent bundle* $\mathrm{S}TM$, consisting of the unit vectors in each tangent space. The fibers of $\mathrm{S}TM$ are spheres. In particular, when M is a two-dimensional surface, $\mathrm{S}TM$ is a circle bundle. The *covariant Dirichlet energy* $E_\nabla[\mathbf{v}] \coloneqq \|\nabla\mathbf{v}\|^2_{L^2(M)}$ measures the smoothness of $\mathbf{v}$ relative to the Levi-Civita connection ∇. A typical field design problem seeks the globally smoothest direction field:

$$\min_{\mathbf{v}\in\Gamma(\mathrm{S}TM)} \|\nabla\mathbf{v}\|^2_{L^2(M)}. \tag{1}$$

We refer to this as a *smooth synchronization problem*.

3.1.1 Topological Defects. Unfortunately, the formulation (1) is ill-posed on most surfaces: the Poincaré-Hopf index theorem rules out non-vanishing tangent vector fields on surfaces of nonzero Euler characteristic.

Thus, we are led to consider *singular* unit fields $\mathbf{v} \in \Gamma(\mathrm{S}TM|_{M\setminus\Sigma})$, defined away from a set Σ of *singularities* or *disclinations*, also known as *topological defects*. A defect is an isolated point x around which $\mathbf{v}$ winds in an irreducible way, such that it cannot be extended continuously over x. This winding is characterized by an integer index associated to x, the same index that appears in the Poincaré-Hopf theorem. Beyond these topological considerations, all the "interesting" behavior of a direction field is concentrated at its singularities, and positioning these singularities is the main challenge of direction field design.

3.1.2 Ill-posedness and Relaxation. Consider a singular unit vector field $\mathbf{v}$ with a singularity at a point $x \in M$. On a small neighborhood $U \ni x$, we may choose a complex coordinate so that locally $\mathbf{v}\,|_{U\setminus\{x\}}$ may be represented as $e^{i\theta}$, and the connection as $\nabla \coloneqq \mathrm{d} - i\eta$, for $\eta \in \Omega^1(U)$ a local connection 1-form (in particular, $\mathrm{d}\eta = \kappa$, the Gauss curvature 2-form). The phase $\theta \in \mathbb{R}$ is multi-valued on $U\setminus\{x\}$, winding by an integer multiple of 2π on any loop enclosing the defect. Now consider a loop γ of length $2\pi\epsilon$ bounding a region $V \subset U$ about x. Then on γ,

$$\begin{aligned}
2\pi\epsilon \int_\gamma |\nabla\mathbf{v}|^2 \,\mathrm{d}s &= \left(\int_\gamma \mathrm{d}s\right)\left(\int_\gamma |\mathrm{d}\theta - \eta|^2 \,\mathrm{d}s\right) \\
&\geq \left|\int_\gamma \mathrm{d}\theta - \eta\right|^2 && \text{(Cauchy-Schwarz)} \\
&= \left|2\pi - \int_V \kappa\right|^2 \xrightarrow{\epsilon\to 0} 4\pi^2 && \text{(Stokes; smooth } \kappa\text{)}
\end{aligned}$$

Integrating over the radial direction, one finds that $\nabla\mathbf{v}$ is not square-integrable in the neighborhood of a defect. To get a well-posed problem, one must change either the objective functional or the constraint in (1). This is known as *relaxation*. We will discuss relaxation in more detail in Section 4.2. First, we discuss the discretized version of (1) as a synchronization problem.

Fig. 2. Dislocations in 2D and 3D.

3.2 Discrete Direction Field Design

We now recall the standard approach to discretizing smooth synchronization problems of the form (1).

To discretize (1) on a triangulated surface, one represents $\mathbf{v}$ by its values on either vertices or faces. The connection is represented by parallel transport operators along or across edges, respectively. We summarize this state of affairs by a graph $\Gamma = (V, E)$, which will be either the vertex-edge or face–dual edge graph of the triangle mesh. After making an arbitrary choice of frame at each graph node $i \in V$, unit tangent vectors can be represented by unit complex numbers v_i. The connection is represented by parallel transport operators $u_{ij} \in \mathrm{SO}(2) \cong \mathrm{U}(1)$, i.e., also unit complex numbers. With these choices, we can immediately define the discrete covariant derivative and Dirichlet energy as:

$$E_\nabla[\mathbf{v}] \coloneqq \sum_{i\sim j} w_{ij}\|(\nabla\mathbf{v})_{ij}\|^2_F, \qquad (\nabla\mathbf{v})_{ij} \coloneqq v_j - u_{ij}v_i. \tag{2}$$

Here $w_{ij} > 0$ are geometric weights, such as the standard cotangent weights in the case where Γ is the primal graph.

Thus, we arrive at the discrete synchronization problem

$$\min_v \left\{v^\dagger L_u v : |v_i|^2 = 1 \;\forall i \in V\right\}, \tag{3}$$

where L_u is the *connection Laplacian* associated to the weights w_{ij} and connection u_{ij}, and $v \in \mathbb{C}^n$ is a vector whose ith entry is v_i. The constraint manifold, being a product of n copies of $\mathrm{U}(1) \cong \mathrm{S}^1$, is nonlinear and nonconvex, and in practice, (3) is difficult to solve to global optimality. This is a second reason why relaxation of the problem is essential.

3.3 Stripe Patterns

In stripe pattern generation as formulated by Knöppel et al. [2015], one aims to produce a periodic stripe pattern over a surface M guided by an input vector field $\mathbf{v}$. Just as direction fields feature disclination defects, stripe patterns in the real world include topological defects called *dislocations* at which it is impossible to choose a well-defined phase θ (see Figure 2). Allowing automatic placement of dislocations is essential for flexible design of stripe patterns.

Formally, away from dislocations, a circle-valued field $\varphi : M \to \mathrm{S}^1$ encodes the phase of the periodic pattern. Ideally, we want to find φ whose phase rotates at a rate $\hat{\mathbf{v}} = \mathbf{v}/\hbar$, where $\hbar$ is a global parameter controlling the period of the stripes. In other words, we want $\nabla\theta \approx \hat{\mathbf{v}}$,

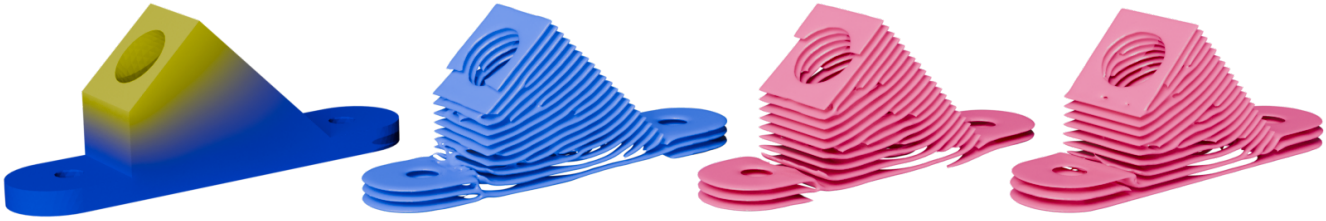

Fig. 3. The gradient vector field of a harmonic function (left) guides the formation of volumetric stripes with spectral relaxation (blue) and SDP relaxation (pink). In the SDP framework, we may impose boundary conditions that give us better control over alignment of the layers (right).

where locally $\varphi = e^{i\theta}$. With an input vector field $\mathbf{v}$ of constant norm, the level sets of φ (stripes) will then be approximately evenly spaced.

Relaxing this requirement to a least-squares objective yields an optimization problem of the form

$$\min_{\varphi: M \to \mathrm{S}^1} \|\nabla\theta - \hat{\mathbf{v}}\|^2_{L^2(M)}. \tag{4}$$

It is convenient to work with the 1-form $\hat{\eta} = \hat{\mathbf{v}}^\flat$. The objective can then be rewritten as follows:

$$E_{\hat{\eta}}[\varphi] \coloneqq \|\nabla\theta - \hat{\mathbf{v}}\|^2 = \|\mathrm{d}\theta - \hat{\eta}\|^2 = \|(i\mathrm{d}\theta - i\hat{\eta})\varphi\|^2 = \|\mathrm{D}_{\hat{\eta}}\varphi\|^2, \tag{5}$$

where $\mathrm{D}_{\hat{\eta}} \coloneqq \mathrm{d} - i\hat{\eta}$ is the covariant derivative associated to the U(1)–connection form $\hat{\eta}$. Like (1), this is an instance of a U(1)-synchronization problem.

Discretely, the underlying graph comprises the vertices and edges of a triangulation. The discrete connection values are calculated as

$$u_{ij} = \exp\left(i \int_{ij} \hat{\eta}\right) = \exp\left(\frac{i}{\hbar} \sum_{t \sim ij} \alpha^t_{ij} \mathbf{v}_t \cdot \mathbf{e}^t_{ij}\right), \tag{6}$$

where $\mathbf{v}_t$ is the (constant) value of $\mathbf{v}$ on face t, $\mathbf{e}^t_{ij}$ is the vector representing the edge (i, j) in the local frame associated to the face t, and α^t_{ij} is an averaging weight.

3.4 Volumetric Stripe Patterns

It is straightforward to generalize the stripe pattern problem to a volumetric setting, with synchronization taking place on the vertices and edges of a tetrahedral mesh. The resulting evenly-spaced laminar structures could, e.g., model layer deposition in curved-layer 3D prints [Zhang et al. 2022]. For this application, optimal surface finish requires that layers align to boundary surfaces, necessitating boundary constraints (see discussion in Sections 4.4 and 5). An example is shown in Figure 3.

To visualize volumetric stripe patterns, we generalize the interpolant from Knöppel et al. [2015] to singular tetrahedra, pictured on the right of Figure 2. When a tetrahedron is punctured by a dislocation curve, any cross-section through the pattern transverse to the curve features a 2D dislocation like the one on the left of Figure 2. Details on this volumetric interpolant are provided in Supplementary Section 2.

3.5 Quad Patterns

Fang et al. [2018] compute bidirectionally periodic patterns, driven by cross fields, for application in quad meshing (see Figure 4). This generalizes the synchronization problems we have examined so far to higher-dimensional fibers and groups.

These "quad patterns" have two independent directions of periodicity at each point, encoded in two independent phases. So locally, they can be represented by functions valued in $\mathrm{S}^1 \times \mathrm{S}^1 = \mathbb{T}^2$. However, around a disclination of the cross field, these two phases cannot be treated independently. As a cross field experiences a rotation of $\pi/2$ upon encircling an index-1 disclination, so must the quad pattern phases experience a corresponding transformation taking $(e^{i\theta_x}, e^{i\theta_y}) \mapsto (e^{i\theta_y}, e^{-i\theta_x})$, whose matrix is

$$A = \begin{pmatrix} 0 & I \\ K & 0 \end{pmatrix}, \tag{7}$$

where K is the complex conjugation map.

To set up the discrete synchronization problem on a triangulated surface M, given a cross field χ with disclinations Σ, the surface must first be cut to separate $M \setminus \Sigma$ into a topological disk $\bar{M}$. On $\bar{M}$, χ can be lifted to a pair of unit vector fields, which exchange places and signs according to the gluing relations along the seams. Finally, $\mathrm{S}^1 \times \mathrm{S}^1$ fibers are attached to each vertex of $\bar{M}$, with constraints of the form $\varphi_i = A^n \varphi_j$ for each pair of vertices $i \sim j$ identified across a seam. The synchronization problem can then be constructed as for a pair of ordinary stripe patterns driven by a pair of vector fields on each triangle. Since A preserves the product of real parts of the phases, a globally-defined real scalar oscillatory field can be extracted, as visualized in Figure 4. Note, however, that due to the anti-unitarity of K, we cannot use complex number computations and must instead represent the circular phases in real coordinates.

Similarly, computing a non-orientable stripe pattern guided by a line field amounts to synchronization on a cover branched at disclinations of the line field [Knöppel et al. 2015] (see Figure 6). Both quad and stripe patterns optimize dislocations against a background of *fixed* disclinations. Automatically placing both kinds of defects remains an open problem.

3.6 Fluid Visualization: Clebsch Map Regression

In fluid dynamics, a generalized Clebsch map is a map $\varphi : \Omega \to X$ from the fluid domain $\Omega \subseteq \mathbb{R}^3$ to a model manifold X, such that the vorticity 2-form ω of the fluid on Ω is the pullback of a fixed 2-form

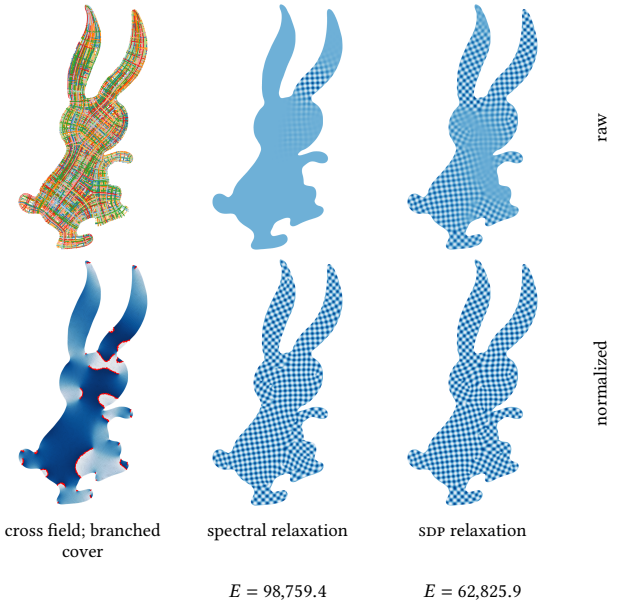

Fig. 4. Synchronization on a nontrivial torus bundle produces oscillatory fields suitable for quad meshing.

β on X,

$$\omega = \varphi^* \beta. \tag{8}$$

In computer graphics, Clebsch maps are used to visualize vortex structures such as filaments and tubes [Chern et al. 2017; Weißmann et al. 2014]. First, given the velocity 1-form η and corresponding vorticity $\omega = \mathrm{d}\eta$ of a fluid flow, one finds a map φ approximately satisfying (8). Next, vortex structures are extracted from φ. We refer to such approaches as *Clebsch map regression* methods.

3.6.1 Vortex Filaments. Weißmann et al. [2014] set $X = \mathbb{C}$ and $\beta = h\delta_0$, where h is a quantization parameter and δ_0 is the Dirac delta measure at the origin. With these choices, a Clebsch map has the interpretation of cutting out *vortex filaments*: given a test surface patch $S \subset \Omega$,

$$\int_S \omega = \int_S \varphi^* \beta = \int_{\varphi(S)} \beta = \#(\varphi(S) \cap \{0\}) = \#(S \cap \varphi^{-1}(0)), \tag{9}$$

where # denotes oriented count. In other words, the vorticity through any surface is determined by its intersection with a set of oriented curves comprising $\varphi^{-1}(0)$, each carrying vorticity h. One can think of the Clebsch map regression as a way of *quantizing* or *rounding* ω to such a singular vorticity field. As such, we refer to this problem as *singular Clebsch mapping*.

To find φ, one might naïvely pose the optimization problem

$$\min_{\varphi:\Omega\to\mathbb{C}} \quad \|\varphi^*\beta - \omega\|_{L^2}^2. \tag{10}$$

However, since β is singular, it is more natural to work with its integral, $\mathrm{d}\theta$, where θ is the multiple-valued argument function on $\mathbb{C}$. Recall that $\mathrm{d}^2\theta \neq 0$, but is instead equal to $2\pi\delta_0$. Now we want

$$\mathrm{d}\eta = \omega = \varphi^*\beta = \frac{h}{2\pi}\mathrm{d}(\varphi^*\mathrm{d}\theta) =: \hbar\mathrm{d}(\varphi^*\mathrm{d}\theta). \tag{11}$$

It is equivalent to seek φ such that

$$\hat{\eta} := \frac{\eta}{\hbar} = \varphi^*\mathrm{d}\theta = \mathrm{d}(\theta \circ \varphi) \tag{12}$$

φ now appears only via its composition with θ, so we might as well treat it (away from vortex filaments) as valued in $\mathrm{S}^1 \subset \mathbb{C}$. Thus, writing $\varphi = e^{i\theta}$, observing that $\mathrm{d}\varphi = i\varphi\,\mathrm{d}\theta$, and enforcing (12) in least squares, we obtain the optimization problem

$$\min_{\varphi:\Omega\to\mathrm{S}^1} \quad \|\mathrm{d}\varphi - i\hat{\eta}\varphi\|_{L^2}^2, \tag{13}$$

whose objective is identical to (5). For direction fields and stripe patterns, the topological defects were *disclinations* and *dislocations*, respectively; here they encode the vortex filaments (see Figure 5).

3.6.2 Vortex Tubes. Chern et al. [2017] use Clebsch map regression to visualize *vortex tubes*. They set $X = \mathrm{S}^2$ and β the area form on X, scaled by $\hbar$. Because the model manifold is a sphere, they refer to these as *spherical Clebsch maps*. To see that spherical Clebsch maps encode vortex tubes, let $R \subseteq \mathrm{S}^2$ and consider again a test surface $S \subset \Omega$.

$$\int_{S\cap\varphi^{-1}(R)} \omega = \int_{S\cap\varphi^{-1}(R)} \varphi^*\beta = \int_{\varphi(S)\cap R} \beta = \hbar\mathrm{area}(\varphi(S) \cap R), \tag{14}$$

counted with orientation and multiplicity. That is, wherever the tube $\varphi^{-1}(R)$ passes through S, it carries a constant vorticity flux $\hbar\mathrm{area}(R)$.

Unlike in the singular Clebsch case, β is not exact. But it is equal to the (scaled) curvature 2-form of the sphere. So it becomes exact after lifting to the unit tangent bundle of S^2, i.e., the Hopf fibration $\pi_H : \mathrm{S}^3 \to \mathrm{S}^2$:

$$\pi_H^*\beta = \hbar\mathrm{d}\alpha, \tag{15}$$

where α is the Levi-Civita connection form on S^3 viewed as $S\mathrm{T}\mathrm{S}^2$. Thus Chern et al. [2017] compute a lifted Clebsch map $\psi : \Omega \to \mathrm{S}^3$, with $\varphi := \pi_H \circ \psi$. With this choice, $\varphi^*\beta = \omega$ if and only if $\psi^*\alpha = \hat{\eta}$.

After adding a regularization term, Chern et al. [2017] arrive at the objective:

$$\|\mathrm{d}\psi - i\hat{\eta}\psi\|_\epsilon^2 \tag{16}$$

where ϵ is a regularization parameter, and $\|\cdot\|_\epsilon$ denotes the L^2 metric incorporating the Berger metric on the sphere. We restrict to the case $\epsilon = 1$, which corresponds to the ordinary isotropic metric. Note the extreme similarity to (13), (5), and (1). The only difference is that the U(1) connection $\hat{\eta}$ now acts on ψ valued in S^3 rather than S^1.

3.7 General Synchronization Problems

It should be evident that the previous examples fit a template, which we summarize in the following definition:

Definition 3.1. Let $\Gamma = (V, E, w)$ be a weighted graph. Let $F \subseteq \mathbb{R}^n$ and let G be a subgroup of $\mathrm{O}(n)$. Given a discrete G-connection u, consisting of one element $u_{ij} \in G$ for each edge $(i, j) \in E$, the *G–synchronization problem* on Γ with fiber F and connection u is defined as follows:

$$\begin{aligned} \min_{\varphi} \quad & \sum_{ij} w_{ij}\|\varphi_j - u_{ij}\varphi_i\|^2 \\ \text{subject to} \quad & \varphi_i \in F \quad \forall i. \end{aligned} \tag{17}$$

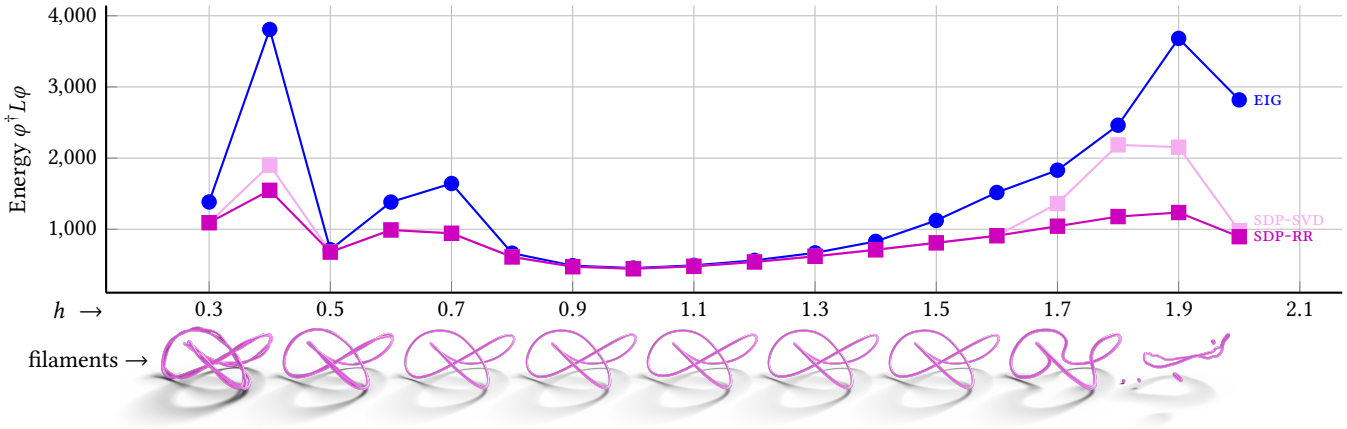

Fig. 5. Weißmann et al. [2014] recover vortex filaments by solving a synchronization problem. Here we compare several relaxations of their covariant Dirichlet energy on a $40 \times 40 \times 20$ grid. **EIG** shows the energy attained by the smallest eigenvector of L after normalization (spectral relaxation). **SDP-SVD** corresponds to the energy of the (normalized) leading eigenvector extracted from the SDP relaxation via singular value decomposition. **SDP-RR** shows the further reduction in energy obtained by randomized rounding: when the SDP solution has rank greater than one, we sample multiple random projections from the SDP solution space and retain the lowest-energy realization.

Note that the right action of G, $\varphi_i \mapsto \varphi_i g$, preserves the objective function. So the best one can hope for is to find the orbit under this symmetry, i.e., a point in F^n/G.

In the literature, synchronization is often given a statistical interpretation. The u_{ij} are viewed as noisy observations of relative orientation or phase, and the goal of synchronization is to stitch together these partial, incompatible observations to infer global information about absolute orientations [Singer 2011; Singer and Sigworth 2020; Wang and Singer 2013]. However, the more geometric language of connections has also shown up in the literature [Bandeira et al. 2013; Gao et al. 2021], in line with concepts of discrete connections in geometry processing [Berwick-Evans et al. 2023; Braune et al. 2025; Knöppel et al. 2025] and lattice gauge theory in physics.

In the smooth setting, we view φ as a (possibly singular) section of a fiber bundle $\pi : E \to M$ with fiber F and connection D. For example, cross fields on surfaces are sections of $\mathrm{S}(TM)^{\otimes 4}$, the sphere bundle of the fourth tensor power of the tangent bundle, expressing the well-known fact that a cross field can be represented by its complex fourth power in an appropriate basis.

4 Method

In this section, we describe and motivate our use of SDP relaxation and factorization to solve synchronization problems like those detailed above. We focus here on the case of U(1) synchronization for convenience, though the constructions are substantially similar in the more general setting. The discrete U(1) synchronization problem with connection Laplacian L

$$s^*(L) = \min_{\varphi \in \mathbb{C}^n} \left\{ \varphi^\dagger L \varphi : |\varphi_i|^2 = 1 \ \forall i \right\} \tag{18}$$

encompasses the cases of directional fields, stripe patterns, and singular Clebsch maps from Section 3.

4.1 Detour: Spectral Relaxation

We have seen in Sections 3.1.2 and 3.2 that synchronization problems require relaxation. Knöppel et al. [2013, 2015]; Weißmann et al. [2014] all employ *spectral relaxation*, also known as *eigenvalue relaxation*. In this approach, one replaces the vertex-wise normalization constraint in (18) with a global L^2 normalization constraint:

$$\lambda_1(L, M) \coloneqq \min_{\varphi} \left\{ \varphi^\dagger L \varphi : \varphi^\dagger M \varphi = 1 \right\}, \tag{19}$$

where M is a diagonal mass matrix, and L is a discrete connection Laplacian derived from a triangle mesh.

The resulting problem is equivalent to finding the minimum generalized eigenvalue and corresponding eigenspace of the pair (L, M):

$$L\varphi_1 = \lambda_1(L, M) M \varphi_1 \tag{20}$$

This is indeed a relaxation—after rescaling by a factor $\sqrt{\operatorname{tr} M}$, any feasible solution of (18) is a solution of (19). After solving the eigenproblem (20), vertex-wise normalization recovers an approximate solution to the original problem (18). Unfortunately, this is a loose relaxation, replacing n constraints with a single constraint.

4.2 Semidefinite Relaxation

To motivate a tighter *semidefinite relaxation* of (18), we first rephrase the spectral relaxation (19) as an SDP:

$$\lambda_1(L, M) = \min_{\Phi} \left\{ \langle L, \Phi \rangle : \Phi \succeq 0, \langle M, \Phi \rangle = 1 \right\} \tag{21}$$

where Φ is Hermitian, and $\langle L, \Phi \rangle \coloneqq \operatorname{tr}(L^\dagger \Phi)$ denotes the Frobenius inner product beteween Hermitian matrices. We state the following result and defer the (short) proof to the supplement:

Proposition 4.1. *The optimal value of* (21) *is the minimum generalized eigenvalue of* (L, M). *The columns of the optimal* Φ^* *sit in the corresponding eigenspace. Thus* (21) *is equivalent to* (19).

A tighter semidefinite relaxation, due to Singer [2011], is

$$S^*(L) \coloneqq \min_{\Phi} \left\{ \langle L, \Phi \rangle : \Phi \succeq 0, \ \Phi_{ii} = 1 \ \forall i \right\} \tag{22}$$

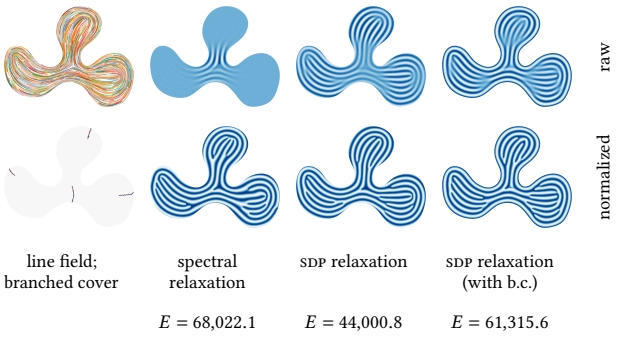

Fig. 6. While eigenfunctions decay to zero over much of the domain, SDP relaxation is more faithful to the original constraints, leading to lower objective values after normalization. Unlike with spectral relaxation, it is easy to impose boundary conditions.

To see that (22) is a relaxation, note that if Φ is a feasible solution of rank 1, we can decompose it as $\Phi =: \varphi\varphi^\dagger$, whence $1 = \Phi_{ii} = |\varphi_i|^2$ for all i. In other words, the rank-one feasible solutions of (22) are exactly the feasible solutions of (18).

On the other hand, (22) is tighter than (21): any feasible solution of (22) is, after rescaling by $\operatorname{tr} M$, a feasible solution of (21). As such,

$$(\operatorname{tr} M)\,\lambda_1(L, M) \le S^*(L) \le s^*(L). \tag{23}$$

Moreover, if the optimum Φ^* of (22) has rank 1, the second inequality becomes an equality, and its factor φ^* is a global optimum of (18). This certifiable optimality is known as *exact recovery* (see [Blekherman et al. 2013]).

4.3 Burer-Monteiro Factorization

SDPs like (22) can be solved to global optimality by a variety of efficient methods. As we have seen in Section 4.2, when an SDP arises from a relaxation, a low-rank solution is desirable. It is thus convenient to use an optimization algorithm that exploits low rank.

We employ the factorization approach to solving SDPs, due to Burer and Monteiro [2003, 2005]. The Burer-Monteiro (BM) method searches for a low-rank solution by partially inverting the relaxation procedure. For example, given the problem (22), pick a rank $r \le n$ and restrict to solutions that factor as $\Phi = YY^\dagger$, where $Y \in \mathbb{C}^{n\times r}$. The SDP (22) becomes:

$$\min_Y \left\{\operatorname{tr}(Y^\dagger L Y) : |Y_j|^2 = 1 \;\forall j\right\} \tag{24}$$

where Y_j is the jth row of Y. Note the similarity between (24) and the original, unrelaxed problem (18). The only difference is that Y_j is now a unit complex vector rather than a unit complex scalar. Geometrically, we have lifted a problem valued in $(\mathrm{S}^1)^n$ to one valued in $(\mathrm{S}^{2r-1})^n$.

Though (24) is a nonconvex quadratically constrained quadratic program (QCQP), it turns out that, under mild conditions, it is actually *equivalent* to the SDP (22) for high enough candidate rank r.

To wit, Barvinok [1995]; Pataki [1998] showed that an SDP with a generic objective will have a solution of rank $r^* \in O(\sqrt{m})$, where m is the number of linear constraints. Moreover, any *rank-deficient* local optimum of (24) is actually a global optimum of (22). More recently, it has been established that for a generic objective L and high enough rank $r \ge r^*$, any second-order local optimum will be rank-deficient, making (24) equivalent to (22) [Boumal et al. 2020]. Putting these ideas together, BM factorization provides a polynomial-time algorithm that finds low-rank ($\Theta(\sqrt{m})$) solutions [Cifuentes and Moitra 2022]. In practice, the rank of solutions is often constant.

BM factorization is applicable more generally for SDP relaxations of problems over smooth homogeneous spaces like spheres or Lie groups [Boumal et al. 2020; Cifuentes 2021], which include synchronization problems of the form (17). Specifically, the nonconvex factored problem must satisfy a linear independence constraint qualification (LICQ). For such SDPs featuring block-diagonal constraints, the corresponding BM problem will—as we saw for (24)—also be an optimization problem over a product of low-dimensional homogeneous spaces. The SDP can be solved efficiently by solving BM problems at increasing ranks via Riemannian optimization. This is known as the Riemannian staircase algorithm [Boumal 2016]. We sketch it in Algorithm 1 in the case of U(1). We solve the subproblem (24) by the Riemannian trust regions method (RTR) [Absil et al. 2007] as implemented in MANOPT [Boumal et al. 2014].[1]

Algorithm 1: U(1) Synchronization with Riemannian Staircase

```
function Y = staircase(L):
    for r ∈ 1, 2, ..., R do
        solve (24) for Y via RTR starting from Y_0
        if rank Y < r then // rank deficient - done!
            return Y
        else // increase rank
            Y_0 ← (Y  0)
            // find minimum eigenpair of Riemannian Hessian
            Z ← arg min{tr(Z†LZ) | Z ∈ T_{Y_0} M_r, ||Z||_F^2 = 1}
            λ ← tr(Z†LZ)
            if λ ≥ 0 then // rank deficient - done!
                return Y_0
            else // found a descent direction
                Y_0 ← linesearch(Y_0, Z)
    end
end
function φ = synch(L):
    Y ← staircase(L)
    begin // svd rounding
        USV† ← svd(Y)
        φ_i ← U_{i1}/|U_{i1}|   ∀i
        return φ
    end
end
```

4.4 Boundary Conditions

Our BM formulation easily extends to Dirichlet boundary conditions (see Figure 6). Decomposing φ into interior and fixed boundary terms $\varphi = I\varphi_I + B\varphi_B$, the unrelaxed objective becomes

$$\min_{\varphi_I \in \mathbb{C}^n} \left\{\varphi_I^\dagger L_{II}\varphi_I + 2\Re(b^\dagger\varphi_I) : |\varphi_i|^2 = 1 \;\forall i \in V_I\right\}, \tag{25}$$

[1] Code for this paper is available at https://github.com/dpa1mer/synchfields.

where $b \coloneqq L_{IB}\varphi_B$ and $L_{II} = I^\dagger L I$, etc. The SDP relaxation of this inhomogeneous problem takes the form

$$\begin{aligned} &\min_{\tilde{\Phi}} \left\{ \langle C, \tilde{\Phi} \rangle : \tilde{\Phi} \succeq 0,\ \tilde{\Phi}_{ii} = 1\ \forall i \right\}, \\ &\tilde{\Phi} = \begin{pmatrix} 1 & \varphi^\dagger \\ \varphi & \Phi \end{pmatrix}, \quad C = \begin{pmatrix} 0 & b^\dagger \\ b & L_{II} \end{pmatrix}. \end{aligned} \tag{26}$$

Choosing a factorization of $\tilde{\Phi}$ and dropping constant terms, the BM objective becomes

$$\operatorname{tr} Y_I^\dagger L_{II} Y_I + \frac{2}{\sqrt{r}} \Re(b^\dagger Y_I \mathbf{1}), \tag{27}$$

where $\mathbf{1}$ is the all-ones vector and $Y_I \in \mathbb{C}^{n \times r}$.

4.5 Rounding

Just as the spectral relaxation (19) required entrywise normalization to recover a feasible point of (18), the SDP method requires a *rounding procedure*.

The solution of (22) is an $n \times n$ Hermitian semidefinite matrix Φ. One way to look at such a matrix is as the *moment matrix* of a probability distribution. Suppose $\varphi \sim \mathcal{N}(0, \Phi)$ is a random variable. Then

$$\mathbb{E}[\varphi^\dagger L \varphi] = \operatorname{tr}(L\Phi) = \langle L, \Phi \rangle. \tag{28}$$

Similarly, $\mathbb{E}[|\varphi_i|^2] = \Phi_{ii} = 1$. In other words, a solution to the SDP (22) can be viewed as a Gaussian random variable that optimizes the original problem (18) *in expectation*. This motivates a natural randomized rounding scheme: sample φ and normalize its entries to obtain a feasible solution of (18). Practically, given the BM solution $Y \in \mathbb{C}^{n \times r}$, sampling φ can be done as follows:

$$\mu, \nu \sim \mathcal{N}(0, I_{r\times r}) \implies \tilde{\varphi} = Y(\mu + i\nu) \implies \varphi_k = \tilde{\varphi}_k / |\tilde{\varphi}_k|. \tag{29}$$

Alternatively, a simple deterministic rounding scheme computes the singular value decomposition (SVD) of Y and sets $\tilde{\phi}$ to the dominant singular vector. Note that as Y approaches rank-one, the randomized and SVD approaches produce increasingly similar solutions.

4.6 Geometric Interpretation

The composition of SDP relaxation and BM factorization lifts a synchronization problem to a higher-dimensional manifold. We have already seen how in the $\mathrm{U}(1) \cong \mathrm{S}^1$ case, the discrete problem (18), valued in $(\mathrm{S}^1)^n$, gets transformed into (24), valued in $(\mathrm{S}^{2r-1})^n$. Interpreted in continuous language, this means that we have lifted our circle bundle $E \to M$ to a new, higher-dimensional bundle $\tilde{E} \to M$ and pulled back the connection to $\tilde{E}$. The connection remains a $\mathrm{U}(1)$ connection operating on the higher-dimensional bundle.

A solution to the lifted problem is a section $\tilde{\varphi} \in \Gamma(\tilde{E})$. Rounding converts this into a section $\varphi \in \Gamma(E)$. The rounding procedure ρ consists of projection Π onto a complex vector $\xi \in \mathbb{C}^r$, followed by normalization N. However, the solution φ is defined only up to a global phase (see Section 3.7), so the rounding procedure is parametrized by the complex line $\bar{\xi} = \operatorname{span} \xi \subset \mathbb{C}^r$. Globally, rounding is parametrized by a choice of complex line bundle $\ell \subset \tilde{E}$, i.e., by a section of the projectivized bundle $\operatorname{proj} \tilde{E}$. In the typical case, where $E \simeq M \times \mathrm{S}^1$ is trivial, this reduces to the choice of a constant point $\bar{\xi} \in \mathbb{CP}^{r-1}$.

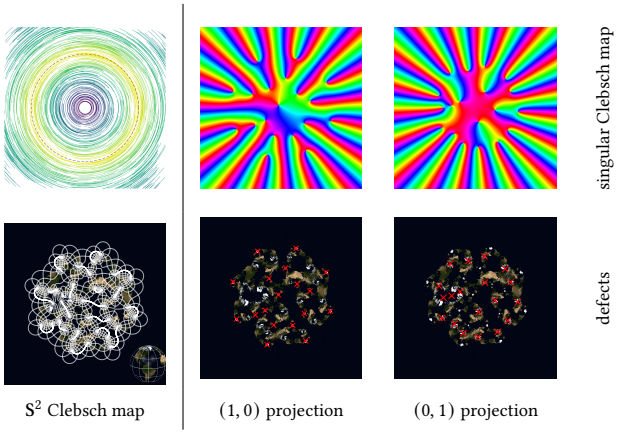

Fig. 7. An S^3-valued Clebsch map yields an S^2-valued Clebsch map via the Hopf fibration, visualized here with an earth texture map. Since S^3 synchronization is the first Burer-Monteiro lift of S^1 synchronization, each point on $\mathrm{S}^2 \cong \mathbb{CP}^1$ also gives a different singular Clebsch map with a corresponding defect configuration.

The Clebsch mapping application supplies geometric insight. The first BM lift of the $\mathrm{U}(1)$-synchronization problem coming from singular Clebsch map regression is an S^3-synchronization problem with a $\mathrm{U}(1)$ connection—precisely the problem of spherical Clebsch map regression!

In the other direction, this means that we can round a spherical Clebsch map $\psi : M \to \mathrm{S}^3$ to a singular Clebsch map $\varphi : M \setminus \Sigma \to \mathrm{S}^1$. The rounding procedure is parametrized by a choice of point $\xi \in \mathbb{CP}^1 \cong \mathrm{S}^2$. Whenever $\tilde{\varphi}(x) \perp \xi \subset \mathbb{C}^2$ for some $x \in M$, the projection is ill-defined, producing a topological defect at x in the resulting φ. This is true exactly when $\pi_H \circ \psi(x)$ is antipodal to ξ in S^2. Hence, every time the spherical Clebsch map $\pi_H \circ \psi$ covers the 2-sphere, one vortex will appear in the singular Clebsch map φ, precisely at the inverse image of $-\xi$. Varying $\xi \in \mathrm{S}^2$ will simply move these defects around the base domain M. Some examples of singular Clebsch maps derived from one spherical Clebsch map are depicted in Figure 7.

Generalizing this perspective, we can view the first BM lift of any $\mathrm{U}(1)$ synchronization problem as optimizing a "spherical Clebsch map." The lifted map is defect-free because $\pi_1(\mathrm{S}^3) = 0$. But it encodes defects in an approximate sense: for any choice of projection, every inverse image of the 2-sphere contains a defect somewhere.

5 Results

Figure 6 demonstrates the faithfulness of SDP relaxation relative to spectral relaxation. Whereas the lowest eigenfunction of the connection Laplacian decays to zero over large regions of the domain, the SDP-based solution has a consistent magnitude across the domain, decaying only where necessary near singularities. Consequently, normalization is much more faithful when applied to the SDP solution compared to the eigenfunction, explaining the lower energy of the normalized SDP solution. To put it another way, the tighter relaxation means that the SDP-relaxed energy measures phase variation more faithfully.

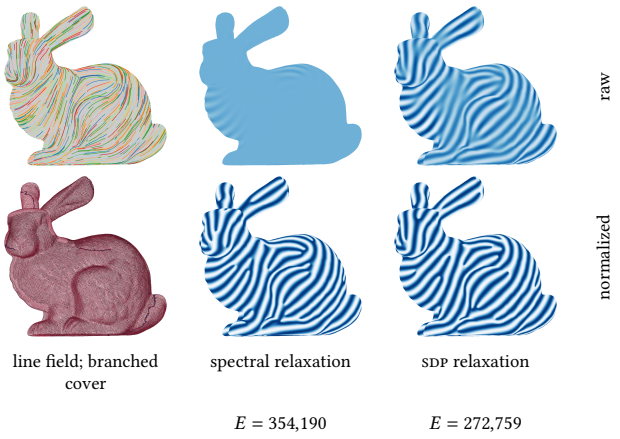

Fig. 8. Aligning to a line field derived from the principal curvature directions, SDP relaxation finds a significantly lower-energy stripe pattern on the bunny.

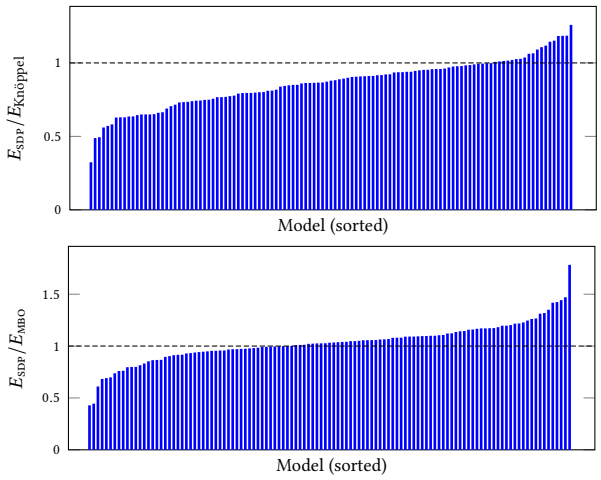

Fig. 9. Our SDP relaxation improves on the method of Knöppel et al. [2013] and performs comparably to the nonconvex MBO method from Viertel and Osting [2019] on computing unit vector fields over a dataset of surface meshes.

Moreover, unlike with spectral relaxation, it is easy to impose additional linear constraints, such as boundary conditions, when solving the SDP. In Figure 6, imposing a constant phase along the boundary results in stripes that precisely match the boundary curve. The line field input was also computed via SDP relaxation.

Results on a curved surface are shown in Figure 8. Note that the line fields displayed are rotated by 90° relative to the ones used for computations, which are always normal to the stripes. This is to aid visual comparison of alignment.

Figure 4 depicts quadrilateral oscillatory patterns computed with SDP and spectral relaxation.

In Figure 3, we show a comparison between spectral and SDP relaxation in computing volumetric stripes. Figure 3 also depicts a volumetric stripe pattern with fixed boundary conditions. Notice how the top surfaces of the model line up exactly with the top stripe layer. This is important for curved-layer 3D printing as aliasing at object boundaries would lead to imperfect surface finish and higher surface friction (see Section 3.4).

In Figure 9 we compare the energy values of SDP relaxation to two other methods on the task of computing unit vector fields on a dataset of 116 surface models from Myles et al. [2014], all rescaled to have unit total area. All three methods were implemented using the same connection Laplacian (and mass matrix), which was also used to measure the covariant Dirichlet energy E. For MBO [Merriman et al. 1994; Viertel and Osting 2019], which takes a timescale parameter, we use the parameter 0.1 for all models. For the method of Knöppel et al. [2013], we use spectral relaxation on closed surfaces and solve a linear system with Dirichlet boundary conditions otherwise since spectral relaxation is incompatible with boundary conditions.

This experiment was conducted on a MacBook Pro (M1 Max, 64 GB RAM). On 67% of meshes, solving our SDP took less than 60 seconds. 39% of the time, our method was faster than MBO. However, on 22% of examples, it was at least an order of magnitude slower than MBO, with the very slowest run taking over 7800 seconds (two orders of magnitude slower than MBO). The timing variation is not wholly explained by mesh size. For example, we observe a negative correlation between runtime and minimum triangle area (as low as 10^{-13} of total area on this dataset). Similarly, we observe positive correlations between runtime and both maximum (up to 5×10^8 on this dataset) and RMS vertex curvature. As usual in geometry processing, quality and performance depend to a large degree on the quality of the input geometry.

Figure 5 compares the two relaxation methods on computing singular Clebsch maps at different values of $\hbar$ for the trefoil knot vorticity field from Weißmann et al. [2014]. The resulting quantized vortices are shown below. SDP relaxation out-competes spectral relaxation with either rounding procedure.

Figure 7 displays a two-dimensional spherical Clebsch map computed from the planar velocity field from Chern et al. [2017], whose curl is constant inside the unit circle and zero elsewhere. As the spherical Clebsch problem is the first Burer-Monteiro factorization of the singular Clebsch problem, in computing the spherical map we implicitly compute a family of singular maps, also shown in Figure 7.

6 Conclusion

Our synchronization framework unifies direction fields, stripe patterns on surfaces and in volumes, quad oscillatory fields, and Clebsch maps. We anticipate our SDP relaxation and BM factorization approach will be adaptable to many other problems that can be phrased as synchronization. For example, the full quad meshing problem (with free disclinations) can be viewed as finding a singular section of a twisted torus bundle (see [Hocking et al. 2022]). Volumetric frame field optimization for hexahedral meshing involves synchronization over a quotient of SO(3) [Palmer et al. 2020].

Another avenue for future work is improving the efficiency of our method. The Riemannian trust regions algorithm that forms the backbone of our method benefits from preconditioning, so one

direction to pursue might be a multigrid preconditioner that works for connection Laplacians [Kalkreuter 1995]. Alternatively, MBO could be used to solve the BM factorizations, fusing the efficiency of MBO with the reliability of a convex method.

Acknowledgments

The authors thank Steven J. Gortler and Nicolas Boumal for interesting discussions.

David Palmer acknowledges the generous support of the NSF MSPRF under award #2303403 during his time at Harvard University. Natalia Pacheco-Tallaj was supported by the NSF DGE-2141064. Mattéo Couplet was supported by Wallonie-Bruxelles International.

References

Pierre-Antoine Absil, Christopher G. Baker, and Kyle A. Gallivan. 2007. Trust-Region Methods on Riemannian Manifolds. *Foundations of Computational Mathematics* 7, 3 (July 2007), 303–330. doi:10.1007/s10208-005-0179-9

Farid Alizadeh. 1995. Interior Point Methods in Semidefinite Programming with Applications to Combinatorial Optimization. *SIAM Journal on Optimization* 5, 1 (Feb. 1995), 13–51. doi:10.1137/0805002

Afonso S. Bandeira, Nicolas Boumal, and Amit Singer. 2017. Tightness of the Maximum Likelihood Semidefinite Relaxation for Angular Synchronization. *Math. Program.* 163, 1-2 (May 2017), 145–167. doi:10.1007/s10107-016-1059-6

Afonso S Bandeira, Yutong Chen, Roy R Lederman, and Amit Singer. 2020. Non-Unique Games over Compact Groups and Orientation Estimation in Cryo-EM. *Inverse Problems* 36, 6 (June 2020), 064002. doi:10.1088/1361-6420/ab7d2c

Afonso S. Bandeira, Christopher Kennedy, and Amit Singer. 2016. Approximating the Little Grothendieck Problem over the Orthogonal and Unitary Groups. *Mathematical Programming* 160, 1 (Nov. 2016), 433–475. doi:10.1007/s10107-016-0993-7

Afonso S. Bandeira, Amit Singer, and Daniel A. Spielman. 2013. A Cheeger Inequality for the Graph Connection Laplacian. *SIAM J. Matrix Anal. Appl.* 34, 4 (Jan. 2013), 1611–1630. doi:10.1137/120875338

Alexander I. Barvinok. 1995. Problems of Distance Geometry and Convex Properties of Quadratic Maps. *Discrete & Computational Geometry* 13, 2 (March 1995), 189–202. doi:10.1007/BF02574037

Pierre-Alexandre Beaufort, Jonathan Lambrechts, François Henrotte, Christophe Geuzaine, and Jean-François Remacle. 2017. Computing Cross Fields A PDE Approach Based on the Ginzburg-Landau Theory. *Procedia Engineering* 203 (Jan. 2017), 219–231. doi:10.1016/j.proeng.2017.09.799

Aharon Ben-Tal, Arkadi Nemirovski, and Cornelis Roos. 2003. Extended Matrix Cube Theorems with Applications to μ-Theory in Control. *Mathematics of Operations Research* 28, 3 (Aug. 2003), 497–523. doi:10.1287/moor.28.3.497.16392

Daniel Berwick-Evans, Anil N. Hirani, and Mark D. Schubel. 2023. Discrete Vector Bundles with Connection and the Bianchi Identity. arXiv:2104.10277 [math] doi:10.48550/arXiv.2104.10277

Grigoriy Blekherman, Pablo A. Parrilo, and Rekha R. Thomas (Eds.). 2013. *Semidefinite Optimization and Convex Algebraic Geometry*. Society for Industrial and Applied Mathematics : Mathematical Programming Society, Philadelphia.

David Bommes, Bruno Lévy, Nico Pietroni, Enrico Puppo, Claudio Silva, Marco Tarini, and Denis Zorin. 2013. Quad-Mesh Generation and Processing: A Survey. *Computer Graphics Forum* 32, 6 (2013), 51–76. doi:10.1111/cgf.12014

David Bommes, Henrik Zimmer, and Leif Kobbelt. 2009. Mixed-Integer Quadrangulation. *ACM Trans. Graph.* 28, 3 (July 2009), 77:1–77:10. doi:10.1145/1531326.1531383

Nicolas Boumal. 2016. A Riemannian Low-Rank Method for Optimization over Semidefinite Matrices with Block-Diagonal Constraints. arXiv:1506.00575 [math] doi:10.48550/arXiv.1506.00575

Nicolas Boumal, Bamdev Mishra, Pierre-Antoine Absil, and Rodolphe Sepulchre. 2014. Manopt, a Matlab Toolbox for Optimization on Manifolds. *Journal of Machine Learning Research* 15, 42 (2014), 1455–1459. http://jmlr.org/papers/v15/boumal14a.html

Nicolas Boumal, Vladislav Voroninski, and Afonso S. Bandeira. 2020. Deterministic Guarantees for Burer-Monteiro Factorizations of Smooth Semidefinite Programs. *Communications on Pure and Applied Mathematics* 73, 3 (March 2020), 581–608. doi:10.1002/cpa.21830

Theo Braune, Yiying Tong, François Gay-Balmaz, and Mathieu Desbrun. 2025. A Discrete Exterior Calculus of Bundle-valued Forms. arXiv:2406.05383 [math] doi:10.48550/arXiv.2406.05383

Samuel Burer and Renato D.C. Monteiro. 2003. A Nonlinear Programming Algorithm for Solving Semidefinite Programs via Low-Rank Factorization. *Mathematical Programming* 95, 2 (Feb. 2003), 329–357. doi:10.1007/s10107-002-0352-8

Samuel Burer and Renato D.C. Monteiro. 2005. Local Minima and Convergence in Low-Rank Semidefinite Programming. *Mathematical Programming* 103, 3 (July 2005), 427–444. doi:10.1007/s10107-004-0564-1

Claudio Castelnovo, Roderich Moessner, and Shivaji L. Sondhi. 2008. Magnetic Monopoles in Spin Ice. *Nature* 451, 7174 (Jan. 2008), 42–45. doi:10.1038/nature06433

Albert Chern, Felix Knöppel, Ulrich Pinkall, and Peter Schröder. 2017. Inside Fluids: Clebsch Maps for Visualization and Processing. *ACM Transactions on Graphics* 36, 4 (July 2017), 1–11. doi:10.1145/3072959.3073591

Diego Cifuentes. 2021. On the Burer–Monteiro Method for General Semidefinite Programs. *Optimization Letters* 15, 6 (Sept. 2021), 2299–2309. doi:10.1007/s11590-021-01705-4

Diego Cifuentes and Ankur Moitra. 2022. Polynomial Time Guarantees for the Burer-Monteiro Method. *Advances in Neural Information Processing Systems* 35 (Dec. 2022), 23923–23935. https://proceedings.neurips.cc/paper_files/paper/2022/hash/9708c7d3a0fef3710f33ba05a74e10b3-Abstract-Conference.html

Katherine Copenhagen, Ricard Alert, Ned S. Wingreen, and Joshua W. Shaevitz. 2021. Topological Defects Promote Layer Formation in Myxococcus Xanthus Colonies. *Nature Physics* 17, 2 (Feb. 2021), 211–215. doi:10.1038/s41567-020-01056-4

Mattéo Couplet, Alexandre Chemin, and Jean-François Remacle. 2024. Integrable Frame Fields Using Odeco Tensors. In *Proceedings of the 2024 International Meshing Roundtable (IMR) (Proceedings)*. Society for Industrial and Applied Mathematics, 53–65. doi:10.1137/1.9781611978001.5

Mihai Cucuringu, Yaron Lipman, and Amit Singer. 2012a. Sensor Network Localization by Eigenvector Synchronization over the Euclidean Group. *ACM Trans. Sen. Netw.* 8, 3 (Aug. 2012), 19:1–19:42. doi:10.1145/2240092.2240093

Mihai Cucuringu, Amit Singer, and David Cowburn. 2012b. Eigenvector Synchronization, Graph Rigidity and the Molecule Problem. *Information and Inference: A Journal of the IMA* 1, 1 (Dec. 2012), 21–67. doi:10.1093/imaiai/ias002

Olga Diamanti, Amir Vaxman, Daniele Panozzo, and Olga Sorkine-Hornung. 2014. Designing N-PolyVector Fields with Complex Polynomials. *Computer Graphics Forum* 33, 5 (2014), 1–11. doi:10.1111/cgf.12426

Nadav Dym and Yaron Lipman. 2017. Exact Recovery with Symmetries for Procrustes Matching. *SIAM Journal on Optimization* 27, 3 (Jan. 2017), 1513–1530. doi:10.1137/16M1078628

Xianzhong Fang, Hujun Bao, Yiying Tong, Mathieu Desbrun, and Jin Huang. 2018. Quadrangulation through Morse-Parameterization Hybridization. *ACM Transactions on Graphics* 37, 4 (Aug. 2018), 1–15. doi:10.1145/3197517.3201354

Nahum Farchi and Mirela Ben-Chen. 2018. Integer-Only Cross Field Computation. *ACM Transactions on Graphics* 37, 4 (Aug. 2018), 1–13. doi:10.1145/3197517.3201375

Golnar Gharooni Fard, Daisy Zhang, Francisco López Jiménez, and Orit Peleg. 2022. Crystallography of Honeycomb Formation under Geometric Frustration. *Proceedings of the National Academy of Sciences* 119, 48 (Nov. 2022), e2205043119. doi:10.1073/pnas.2205043119

Tingran Gao, Jacek Brodzki, and Sayan Mukherjee. 2021. The Geometry of Synchronization Problems and Learning Group Actions. *Discrete & Computational Geometry* 65, 1 (Jan. 2021), 150–211. doi:10.1007/s00454-019-00100-2

Bernd Gärtner and Jiri Matousek. 2012. *Approximation Algorithms and Semidefinite Programming*. Springer Berlin Heidelberg, Berlin, Heidelberg. doi:10.1007/978-3-642-22015-9

Michel X. Goemans and David P. Williamson. 1995. Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems Using Semidefinite Programming. *J. ACM* 42, 6 (Nov. 1995), 1115–1145. doi:10.1145/227683.227684

Dmitry Golovaty, Jose Alberto Montero, and Daniel Spirn. 2021. A Variational Method for Generating n-Cross Fields Using Higher-Order Q-Tensors. *SIAM Journal on Scientific Computing* 43, 5 (Jan. 2021), A3269–A3304. doi:10.1137/19M1287857

Martin Grötschel. 1993. *Geometric Algorithms and Combinatorial Optimization* (2nd ed ed.). Number v.2 in Algorithms and Combinatorics Ser. Springer Berlin / Heidelberg, Berlin, Heidelberg.

Pau Guillamat, Carles Blanch-Mercader, Guillaume Pernollet, Karsten Kruse, and Aurélien Roux. 2022. Integer Topological Defects Organize Stresses Driving Tissue Morphogenesis. *Nature Materials* 21, 5 (May 2022), 588–597. doi:10.1038/s41563-022-01194-5

Brook J. Hocking, Helen S. Ansell, Randall D. Kamien, and Thomas Machon. 2022. The Topological Origin of the Peierls-Nabarro Barrier. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 478, 2258 (Feb. 2022), 20210725. doi:10.1098/rspa.2021.0725

Qi-Xing Huang and Leonidas Guibas. 2013. Consistent Shape Maps via Semidefinite Programming. *Computer Graphics Forum* 32, 5 (Aug. 2013), 177–186. doi:10.1111/cgf.12184

Thomas Kalkreuter. 1995. Multigrid Methods for Propagators in Lattice Gauge Theories. *J. Comput. Appl. Math.* 63, 1 (Nov. 1995), 57–68. doi:10.1016/0377-0427(95)00049-6

Leonid G. Khachiyan. 1979. A Polynomial Algorithm in Linear Programming. *Doklady Akademii Nauk SSSR* 244, 5 (1979), 1093–1096.

Alexei Kitaev. 2003. Fault-Tolerant Quantum Computation by Anyons. *Annals of Physics* 303, 1 (Jan. 2003), 2–30. doi:10.1016/S0003-4916(02)00018-0

Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2013. Globally Optimal Direction Fields. *ACM Transactions on Graphics* 32, 4 (July 2013), 59:1–59:10. doi:10.1145/2461912.2462005

Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2015. Stripe Patterns on Surfaces. *ACM Transactions on Graphics* 34, 4 (July 2015), 1–11. doi:10.1145/2767000

Felix Knöppel, Ulrich Pinkall, Peter Schröder, and Yousuf Soliman. 2025. Rolling Spheres and the Willmore Energy. *Discrete & Computational Geometry* (Nov. 2025). doi:10.1007/s00454-025-00751-4

John M. Kosterlitz and David J. Thouless. 1973. Ordering, Metastability and Phase Transitions in Two-Dimensional Systems. *Journal of Physics C: Solid State Physics* 6, 7 (April 1973), 1181. doi:10.1088/0022-3719/6/7/010

Lev Davidovich Landau. 1936. The Theory of Phase Transitions. *Nature* 138, 3498 (Nov. 1936), 840–841. doi:10.1038/138840a0

Lev Davidovich Landau. 1937. On the Theory of Phase Transitions. *Zhurnal Éksperimentalnoi i Teoreticheskoi Fiziki* 7, 19-32 (1937), 926.

Long Ma, Ying He, Jianmin Zheng, Yuanfeng Zhou, Shiqing Xin, Caiming Zhang, and Wenping Wang. 2024. Computing Smooth and Integrable Cross Fields via Iterative Singularity Adjustment. *IEEE Transactions on Visualization and Computer Graphics* 31, 9 (2024), 4850–4867. doi:10.1109/TVCG.2024.3418892

Thomas Machon and Gareth P. Alexander. 2016. Global Defect Topology in Nematic Liquid Crystals. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 472, 2191 (July 2016), 20160265. doi:10.1098/rspa.2016.0265

Juan Montes Maestre, Yinwei Du, Ronan Hinchet, Stelian Coros, and Bernhard Thomaszewski. 2023. Differentiable Stripe Patterns for Inverse Design of Structured Surfaces. *ACM Transactions on Graphics* 42, 4 (Aug. 2023), 1–14. doi:10.1145/3592114

Haggai Maron, Nadav Dym, Itay Kezurer, Shahar Kovalsky, and Yaron Lipman. 2016. Point Registration via Efficient Convex Relaxation. *ACM Transactions on Graphics* 35, 4 (July 2016), 1–12. doi:10.1145/2897824.2925913

Yonit Maroudas-Sacks, Liora Garion, Lital Shani-Zerbib, Anton Livshits, Erez Braun, and Kinneret Keren. 2021. Topological Defects in the Nematic Order of Actin Fibres as Organization Centres of Hydra Morphogenesis. *Nature Physics* 17, 2 (Feb. 2021), 251–259. doi:10.1038/s41567-020-01083-1

Zoë Marschner, David Palmer, Paul Zhang, and Justin Solomon. 2020. Hexahedral Mesh Repair via Sum-of-Squares Relaxation. *Computer Graphics Forum* 39, 5 (2020), 133–147. doi:10.1111/cgf.14074

Zoë Marschner, Paul Zhang, David Palmer, and Justin Solomon. 2021. Sum-of-Squares Geometry Processing. *ACM Transactions on Graphics* 40, 6 (Dec. 2021), 253:1–253:13. doi:10.1145/3478513.3480551

N. David Mermin. 1979. The Topological Theory of Defects in Ordered Media. *Reviews of Modern Physics* 51, 3 (July 1979), 591–648. doi:10.1103/RevModPhys.51.591

Barry Merriman, James K. Bence, and Stanley J. Osher. 1994. Motion of Multiple Junctions: A Level Set Approach. *J. Comput. Phys.* 112, 2 (June 1994), 334–363. doi:10.1006/jcph.1994.1105

Rahul Mitra, Mattéo Couplet, Tongtong Wang, Megan Hoffman, Kui Wu, and Edward Chien. 2025. Curl Quantization for Automatic Placement of Knit Singularities. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25)*. Association for Computing Machinery, New York, NY, USA, 1–10. doi:10.1145/3721238.3730715

Ashish Myles, Nico Pietroni, and Denis Zorin. 2014. Robust Field-Aligned Global Parametrization. *ACM Trans. Graph.* 33, 4 (July 2014), 135:1–135:14. doi:10.1145/2601097.2601154

Ashish Myles and Denis Zorin. 2013. Controlled-Distortion Constrained Global Parametrization. *ACM Transactions on Graphics* 32, 4 (July 2013), 1–14. doi:10.1145/2461912.2461970

Vidya Narayanan, Lea Albaugh, Jessica Hodgins, Stelian Coros, and James Mccann. 2018. Automatic Machine Knitting of 3D Meshes. *ACM Trans. Graph.* 37, 3 (Aug. 2018), 35:1–35:15. doi:10.1145/3186265

Yurii Nesterov and Arkadii Nemirovskii. 1994. *Interior-Point Polynomial Algorithms in Convex Programming*. Number 13 in Studies in Applied and Numerical Mathematics. Society for Industrial and Applied Mathematics, Philadelphia, PA. doi:10.1137/1.9781611970791

David Palmer, David Bommes, and Justin Solomon. 2020. Algebraic Representations for Volumetric Frame Fields. *ACM Transactions on Graphics* 39, 2 (April 2020), 16:1–16:17. doi:10.1145/3366786

David Palmer, Albert Chern, and Justin Solomon. 2024. Lifting Directional Fields to Minimal Sections. *ACM Transactions on Graphics* 43, 4 (July 2024), 1–20. doi:10.1145/3658198

Gábor Pataki. 1998. On the Rank of Extreme Matrices in Semidefinite Programs and the Multiplicity of Optimal Eigenvalues. *Mathematics of Operations Research* 23, 2 (1998), 339–358. jstor:3690515 https://www.jstor.org/stable/3690515

Nico Pietroni, Marcel Campen, Alla Sheffer, Gianmarco Cherchi, David Bommes, Xifeng Gao, Riccardo Scateni, Franck Ledoux, Jean Remacle, and Marco Livesu. 2022. Hex-Mesh Generation and Processing: A Survey. *ACM Trans. Graph.* 42, 2 (Oct. 2022), 16:1–16:44. doi:10.1145/3554920

Michael Pretko, Zhengzheng Zhai, and Leo Radzihovsky. 2019. Crystal-to-Fracton Tensor Gauge Theory Dualities. *Physical Review B* 100, 13 (Oct. 2019), 134113. doi:10.1103/PhysRevB.100.134113

Nicolas Ray, Wan Chiu Li, Bruno Lévy, Alla Sheffer, and Pierre Alliez. 2006. Periodic Global Parameterization. *ACM Trans. Graph.* 25, 4 (Oct. 2006), 1460–1485. doi:10.1145/1183287.1183297

David M Rosen, Luca Carlone, Afonso S Bandeira, and John J Leonard. 2019. SE-Sync: A Certifiably Correct Algorithm for Synchronization over the Special Euclidean Group. *The International Journal of Robotics Research* 38, 2-3 (March 2019), 95–125. doi:10.1177/0278364918784361

Amit Singer. 2011. Angular Synchronization by Eigenvectors and Semidefinite Programming. *Applied and Computational Harmonic Analysis* 30, 1 (Jan. 2011), 20–36. doi:10.1016/j.acha.2010.02.001

Amit Singer and Yoel Shkolnisky. 2011. Three-Dimensional Structure Determination from Common Lines in Cryo-EM by Eigenvectors and Semidefinite Programming. *SIAM Journal on Imaging Sciences* 4, 2 (Jan. 2011), 543–572. doi:10.1137/090767777

Amit Singer and Fred J. Sigworth. 2020. Computational Methods for Single-Particle Electron Cryomicroscopy. *Annual Review of Biomedical Data Science* 3 (July 2020), 163–190. doi:10.1146/annurev-biodatasci-021020-093826

Amit Singer and Hau-Tieng Wu. 2012. Vector Diffusion Maps and the Connection Laplacian. *Communications on Pure and Applied Mathematics* 65, 8 (2012), 1067–1144. doi:10.1002/cpa.21395

Anthony Man-Cho So, Jiawei Zhang, and Yinyu Ye. 2007. On Approximating Complex Quadratic Optimization Problems via Semidefinite Programming Relaxations. *Mathematical Programming* 110, 1 (June 2007), 93–110. doi:10.1007/s10107-006-0064-6

Florian Cyril Stutz, Tim Felle Olsen, Jeroen Peter Groen, Tuan Nguyen Trung, Niels Aage, Ole Sigmund, Justin Solomon, and Jakob Andreas Bærentzen. 2022. Synthesis of Frame Field-Aligned Multi-Laminar Structures. *ACM Trans. Graph.* 41, 5 (May 2022), 170:1–170:20. doi:10.1145/3516522

Tzvetelina Tzeneva. 2011. *Global Alignment of Multiple 3-D Scans Using Eigevector Synchronization*. Senior Thesis. Princeton University, Princeton, NJ. http://arks.princeton.edu/ark:/88435/dsp01xp68kj31s

Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommes, Klaus Hildebrandt, and Mirela Ben-Chen. 2016. Directional Field Synthesis, Design, and Processing. *Computer Graphics Forum* 35, 2 (2016), 545–572. doi:10.1111/cgf.12864

Ryan Viertel and Braxton Osting. 2019. An Approach to Quad Meshing Based on Harmonic Cross-Valued Maps and the Ginzburg–Landau Theory. *SIAM Journal on Scientific Computing* 41, 1 (Jan. 2019), A452–A479. doi:10.1137/17M1142703

Lanhui Wang and Amit Singer. 2013. Exact and Stable Recovery of Rotations for Robust Synchronization. *Information and Inference* 2, 2 (Dec. 2013), 145–193. doi:10.1093/imaiai/iat005

Steffen Weißmann, Ulrich Pinkall, and Peter Schröder. 2014. Smoke Rings from Smoke. *ACM Transactions on Graphics* 33, 4 (July 2014), 1–8. doi:10.1145/2601097.2601171

Shuzhong Zhang and Yongwei Huang. 2006. Complex Quadratic Optimization and Semidefinite Programming. *SIAM Journal on Optimization* 16, 3 (July 2006), 871–890. doi:10.1137/04061341X

Tianyu Zhang, Guoxin Fang, Yuming Huang, Neelotpal Dutta, Sylvain Lefebvre, Zekai Murat Kilic, and Charlie C. L. Wang. 2022. S^3 -Slicer: A General Slicing Framework for Multi-Axis 3D Printing. *ACM Transactions on Graphics* 41, 6 (Dec. 2022), 1–15. doi:10.1145/3550454.3555516