Complex Line Bundles over Simplicial Complexes

vorgelegt von
Diplom-Mathematiker
Felix Jakob Knöppel
Potsdam

Von der Fakultät II - Mathematik und Naturwissenschaften
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften
Dr. rer. nat.

vorgelegte Dissertation

Promotionsausschuss:

Vorsitzender:	Prof. Dr. Stefan Felsner
Berichter/Gutachter:	Prof. Dr. Ulrich Pinkall
Berichter/Gutachter:	Prof. Dr. Boris Springborn
Berichter/Gutachter:	Prof. Dr. Johannes Wallner

Tag der wissenschaftlichen Aussprache: 14. September 2015

Berlin 2015

Zusammenfassung

Diese kumulative Dissertation behandelt praktische und theoretische Aspekte diskreter Vektorbündel über Simplizialkomplexen. Dabei stehen hier vor allem hermitesche Linienbündel im Vordergrund, welche in den letzten Jahren eine Reihe bemerkenswerter Anwendungen in der Computergrafik fanden. Zwei dieser Anwendungen sind dabei Teil dieser Dissertation: die Berechnung von optimalen n-Richtungsfeldern auf Flächen (Kapitel 1) und die Berechnung von Streifenmustern auf Flächen (Kapitel 2). Die entwickelten Algorithmen liefern im Vergleich mit anderen modernen Methoden Resultate von gleicher Qualität, sind aber um eine Größenordnung schneller. Motiviert durch ihre Anwendungen werden dann im letzten Kapitel diskrete Vektorbündel aus dem Blickwinkel der diskreten Differentialgeometrie betrachtet (Kapitel 3). Dies beinhaltet eine vollständige Klassifikation der diskreten Vektorbündel mit Zusammenhang und die Klassifikation der hermiteschen Linienbündel mit Krümmung. Weiter wird jedem diskreten hermiteschen Linienbündel mit Krümmung ein eindeutiges stückweise glattes Bündel mit Zusammenhang zugeordnet. Dies führt zu einer Verallgemeinerung des bekannten Kotangens-Laplace-Operators auf beliebige diskrete hermitesche Linienbündel mit Krümmung über endlichen Euklidischen Simplizialkomplexen.

Abstract

This cumulative dissertation treats practical and theoretical aspects of discrete vector bundles over simplicial complexes. Here the main focus is on discrete hermitian line bundles which, in recent years, found a number of remarkable applications in computer graphics. Two of these applications are part of this thesis: The computation of optimal n-direction fields on surfaces (Chapter 1) and the computation of stripe patterns on surfaces (Chapter 2). The developed algorithms yield, in comparison with state-of-the-art methods, results of same quality but are an order of magnitude faster. In the last chapter, motivated by their applications, discrete vector bundles are investigated from the viewpoint of discrete differential geometry (Chapter 3). This includes a complete classification of discrete vector bundles with connection and the classification of discrete hermitian line bundles with curvature. Moreover, to each discrete hermitian line bundle with curvature we assign a unique piecewise-smooth hermitian line bundle with connection. This leads to a generalization of the well-known cotangent-Laplace operator to arbitrary discrete hermitian line bundles with curvature over finite Euclidean simplicial complexes.

Contents

Introduction

Vector bundles are fundamental objects in Differential Geometry. The tangent bundle, the cotangent bundle and other bundles associated to the tangent bundle are omnipresent.

At the beginning of the 20th century Tullio Levi-Civita introduced in context of Riemannian geometry the notion of parallel transport. He also found its relation to curvature [**41**]. Shortly thereafter, inspired by physical considerations, Hermann Weyl put bundles on a more abstract level. By introduction of the "scale" bundle – a real line bundle with connection – he generalized Riemannian geometry to what nowadays is known as Weyl geometry. Moreover, this led to one of the first examples of a unified field theory [**73**].

After the rise of quantum mechanics Weyl realized that electromagnetism should instead be encoded by a connection on a U(1)-bundle – the "phase" bundle. This led to a quantum mechanical description and allowed to give a gauge-theoretical explanation of the conservation of electric charge [**74**]. In contrast to the "scale" bundle, the "phase" bundle can be topologically non-trivial. This causes certain topological effects, considered e.g. by Dirac in his theory of magnetic monopoles [**21**], and led to the prediction and verification of the Aharonov-Bohm effect [**1, 8**].

In [**77**] Yang and Mills then generalized these ideas to SU(2)-bundles. This was the beginning of non-abelian gauge-theories whose predictions are verified today in large supercolliders.

Thus bundle theory is motivated to a large extend by Physics. As bundles play such an important role, the Physics literature is also the main place where discretizations of vector bundles and discrete Laplace operators appear [**12, 13, 25, 37, 64**].

Recently discrete line bundles also found a number of remarkable applications in computer graphics [**11, 33, 34, 72**]. Two of these articles are contained in this thesis.

The first application, described in Chapter 1, is the computation of optimally n-RoSy fields on surfaces. Here an n-RoSy field is just a pointwise normalized n-vector field. Since the tangent bundle of an oriented Riemannian surface is a hermitian line bundle, an n-RoSy field can be regarded as a pointwise normalized section of the n-th tensor power of the tangent bundle. This turns the problem into the search for certain direction fields, i.e. pointwise normalized sections in a hermitian line bundle over the surface (defined up to a set of isolated singularities). The term 'optimal' suggests that there is an underlying variational energy. A natural measure for the smoothness of a vector field is its Dirichlet energy. Unfortunately, in the presence of singularities, the Dirichlet energy of a direction field is infinite. Since, in general, direction fields have singularities, we define the energy of a given field ϕ instead to be the ground state energy of a certain Schrödinger operator associated to ϕ. As a consequence we obtain a remarkable simple algorithm for the computation of an optimal direction field:

The optimal direction fields are exactly the pointwise normalized eigenvectors of the Laplace operator corresponding to the smallest eigenvalue.

From the algorithm it is evident that singularities are placed automatically. By inclusion of an additional linear energy meas the deviance from a given section, a specific alignment can be obtained.

Furthermore, a finite element method for discrete hermitian line bundles with curvature over surface is developed. This finally allows a discretization of the problem and hence the computation of optimal direction fields in arbitrary discrete hermitian line bundles.

In order to compute optimal n-RoSy fields we then assign to each oriented Euclidean simplicial surface a discrete tangent bundle. Therefore we equip the surface with its canonical Riemann surface structure and restrict its tangent bundle to the vertex set. Though the original metric has conical singularities the parallel transport of directions is well-defined and yields a discrete connection turning the discrete bundle into a discrete hermitian line bundle with curvature – the discrete tangent bundle of the Euclidean simplicial surface. In particular, we obtain an efficient method for the computation of optimal n-RoSy fields on discrete surface.

Moreover, we use a formula of Weitzenböck-type to include a parameter which can be used to bias the placement of singularities.

Chapter 2 deals with stripe patterns on surfaces. Here one is interested in a function ϕ taking its values in the unit circle, which 'spirals' along given directions, i.e.

$$d\phi = \imath\omega\phi$$

for some given 1-form ω on the surface M - possibly with isolated singularities. The stripes then arise as the level lines of ϕ and hence are perpendicular to ω. Of course, such ϕ can only exist for special forms ω. Taking a different point of view, one possibility to relax the problem seems to be obvious: If one considers the bundle $\mathrm{M} \times \mathbb{C}$ with connection $\nabla = d - \imath\omega$, then ϕ satisfying the equation above is a non-vanishing parallel section - which does not exist in general. Instead we ask for an optimal direction field in the sence of Chapter 1. Such field always exists and yields, in case the bundle was trivial, the desired solution. The stripe patterns generated by this method show in general, besides the singularities of ω, a characteristic kind of singularities (edge dislocations). These appear in nature on several scales.

If one considers instead of a 1-form ω a quadratic differential q, a corresponding unoriented stripe pattern can be generated by switching to a double cover of M. Though a double cover is now involved, the size of problem stays the same. Moreover we provide a non-linear interpolation scheme to extract an (up to isolated singularities) continuous stripe pattern out of the discrete data.

In comparison with state-of-the-art methods, the developed algorithms yield results of same quality, but are an order of magnitude faster.

Motivated by their applications, Chapter 3 treats discrete vector bundles from the viewpoint of discrete differential geometry. This includes a complete classification of the discrete vector bundles over finite simplicial complexes by their monodromy, which can be regarded as a discrete analogue of a theorem of Shôshichi Kobayashi, as well as a discrete analogue of a theorem of André Weil on the classification of hermitian line bundles with connection by their curvature. Moreover, for each section of a discrete hermitian line bundle with curvature a discrete index form is defined in such a way that a discrete version of the famous Poincaré-Hopf index theorem is valid. Finally, it is shown that each discrete hermitian line bundle with curvature over the simplicial complex extends uniquely to a piecewise-smooth hermitian line bundle with connection - its associated bundle. Each discrete section is then identified with a section of the associated bundle, which is used to define a discrete Dirichlet energy for arbitrary discrete hermitian line bundles with curvature over Euclidean simplicial manifolds. The corresponding Laplace operator is a generalization of the well-known cotangent-Laplace operator.

1 | Globally Optimal Direction Fields

Felix Knöppel, Keenan Crane, Ulrich Pinkall, Peter Schröder
ACM Trans. Graph. 32, 4 (July), 59:1–59:10, 2013.

Abstract. We present a method for constructing smooth n-direction fields (line fields, cross fields, *etc.*) on surfaces that is an order of magnitude faster than state-of-the-art methods, while still producing fields of equal or better quality. Fields produced by the method are globally optimal in the sense that they minimize a simple, well-defined quadratic smoothness energy over all possible configurations of singularities (number, location, and index). The method is fully automatic and can optionally produce fields aligned with a given guidance field such as principal curvature directions. Computationally the smoothest field is found via a sparse eigenvalue problem involving a matrix similar to the cotan-Laplacian. When a guidance field is present, finding the optimal field amounts to solving a single linear system.

1 Introduction

A direction field φ of degree $n \in \mathbb{N}$ associates a collection of n evenly spaced unit tangent vectors to each point of a surface. For instance, $n = 1, 2$ and 4 correspond to direction, line, and cross fields, respectively (Figure 2). In general such fields must have *singularities*, *i.e.*, isolated points where the field fails to vary smoothly.

At first glance, computing the smoothest n-direction field appears to be a difficult combinatorial optimization problem for two reasons. First, we must determine the optimal number, placement, and indices of singularities. Second, we must identify directions that differ in angle by integer multiples of $2\pi/n$. For a *fixed* configuration of singularities, Crane *et al.* **[16]** demonstrate that an optimal solution can be found by solving a pair of sparse linear systems. In many situations, however, it is desirable to place singularities automatically. Historically this task has been formulated in terms of difficult nonconvex optimization problems where little can be said about global optimality (Section 1.1). In this paper we describe a simple quadratic smoothness energy that easily admits a global minimum with respect to all possible configurations of singularities. Fig. 1 shows one example.

Our method has two key ingredients. **First**, we represent n-direction fields by storing the n^{th} power of a complex number at each vertex, together with an arbitrary (but fixed) tangent basis direction. Optimizing the smoothness of such a field does not require period jumps or trigonometric functions as in previous methods (Section 1.1). **Second**, we measure the smoothness of an n-direction field using the ground state energy of an appropriate *Schrödinger operator*. Unlike many methods, this formulation does not require a nonconvex unit-norm constraint on each vector, and is well-defined even for singular n-direction fields (Section 3). In addition, we introduce a continuum of "geometry-aware" smoothness energies that provide a tradeoff between the straightness of field lines and the total number of singularities. Finally, we allow a tradeoff between smoothness and alignment with a guidance field, which in the case of principal curvature alignment leads to a simple, automatic scheme without the need for careful tuning of parameters.

For triangle meshes our algorithm requires the setup of a matrix similar to the cotan-Laplacian, but with complex instead of real variables. To build this system we need a *connection*, *i.e.*, a prescription for mapping tangent vectors between tangent spaces, expressed via unit complex numbers (rotations) associated with edges. Finding an eigenvector corresponding to the smallest eigenvalue of this matrix gives a global minimizer of the smoothness energy. If alignment with

Figure 1. Smoothest unit vector field on the Stanford bunny over all possible configurations of singularities, computed by solving a single eigenvector problem. Red and blue spheres indicate positive and negative singularities, respectively. (471ms, $|T| = 28$k)

a given field is desired, solving a single linear system yields a global minimizer of a weighted smoothness and alignment energy. An outline of these algorithms is given at the end of Section 6.

Throughout we will distinguish between *directions* which have unit length, and *vectors* which can have any length.

1.1 Related Work

Computation of smooth n-direction fields on surfaces, in particular the case $n = 4$, is an essential component in applications ranging from nonphotorealistic rendering **[27]** to texture synthesis **[39]**, parameterization **[59]**, and remeshing **[29, 5, 53]**, to name a few. In these applications one is often interested in smooth fields which approximate directions of *principal curvature*, *i.e.*, the most extreme directions of bending. While it is relatively straightforward to *formulate* energies which encode the desired effects, they are typically nonconvex and often NP-hard to solve. Finding a simple *convex* formulation that yields the globally smoothest n-direction fields – in *any* well-defined sense – has so far eluded researchers in this area.

A principal difficulty is the need to identify directions modulo $2\pi/n$. One possible approach is to simply multiply angles by n. For example, Hertzmann and Zorin **[27]** work with 4-direction fields

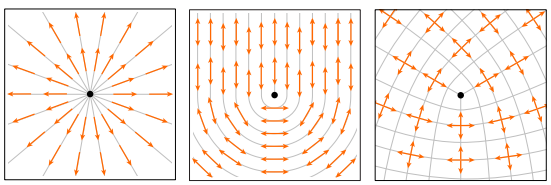

Figure 2. *Left to right:* examples of n-direction fields for $n = 1$ (direction), $n = 2$ (line), and $n = 4$ (cross), near singularities of index $+1$, $+\frac{1}{2}$, and $+\frac{1}{4}$, respectively.

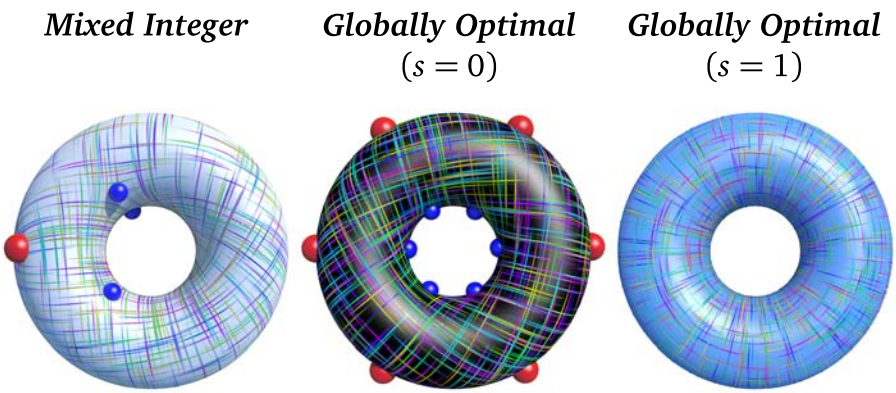

Figure 3. Smoothness energies involving integer variables are easy to formulate but difficult to minimize. *Left:* solution produced by the mixed-integer method of Bommes *et al.* [**5**]. *Center:* solution produced by our method with $s = 0$. *Right:* solution produced by our method with $s = 1$. Both solutions are global minimizers; the parameter s offers a tradeoff between smoothness and number of singularities.

and use an energy

$$E(\phi) = -\sum_{ij \in E} \cos\left(4(\phi_i - \phi_j + \theta_{ij})\right),$$

where angles ϕ_i specify directions at vertices and θ_{ij} is the angle between neighboring tangent frames. A variant of this "cosine" energy was later adopted by Kälberer *et al.* [**29**]. Owing to nonconvexity, little can be said about optimality of solutions found via local descent; moreover, results depend on initialization.

Another possibility is to use the *representation vectors* $(u_i, v_i) := (\cos(n\phi_i), \sin(n\phi_i))$ themselves as variables [**59, 54, 60**]. Now however an additional unit constraint $u_i^2 + v_i^2 = 1$ enters, leading to an NP-hard nonconvex problem [**43**]. In practice the unit constraint is included as a penalty term or relaxed through iterative renormalization. As before, little can be said about global optimality.

A different route was taken by Ray *et al.* [**61**] who work with angles directly and use so-called *period jumps* $p_{ij} \in \mathbb{Z}$ to compare angles across a given edge e_{ij}, leading to a smoothness energy

$$E(\phi, p) = \sum_{e_{ij} \in E} \left(\phi_j - \phi_i + \theta_{ij} + \tfrac{2p_{ij}\pi}{n}\right)^2, \tag{1.1}$$

which is quadratic in ϕ and p. Since the period jumps are integer variables, the feasible set of solutions is nonconvex. In general, mixed-integer optimization problems are NP-hard and a relaxation must be applied. Ray *et al.* apply a direct rounding procedure after first treating the p_{ij} as real valued, while Bommes *et al.* [**5**] use a greedy rounding procedure which tends to produce solutions with lower energy. Again, global optimality of the resulting solutions cannot be asserted.

In all these approaches the user need not provide locations for the singularities (which must be present on surfaces of arbitrary topology). If singularities are prescribed ahead of time, very simple quadratic formulations are possible as demonstrated by Fisher *et al.* [**22**] for vector fields and by Crane *et al.* [**16**] for n-direction fields.

Our representation is perhaps closest to that of Kass and Witkin [**32**] who used squared complex numbers to extract unoriented line fields (*i.e.*, $n = 2$) from images. Palacios and Zhang [**54**] propose a similar idea for arbitrary n but use real representation vectors $(R\cos(n\phi), R\sin(n\phi))$ and encode transformations between tangent spaces via matrices. These representations are attractive, but neither work addresses the question of finding optimally smooth fields on surfaces.

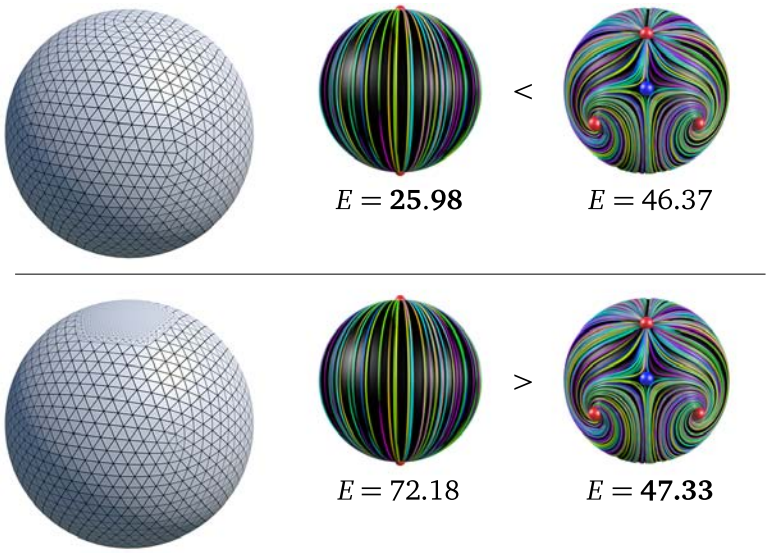

Figure 4. Dirichlet energy of a *unit* n-direction field is not a reliable measure of quality, since under refinement the energy contributed by a singularity grows without bound. *Top:* uniform tessellation yields the expected result: antipodal singularities have lower energy than the configuration on the right. *Bottom:* refining around the poles increases the energy of the antipodal pair, reversing this relationship. This phenomenon is a consequence of the smooth formulation and is not fixed by, *e.g.*, more accurate area weighting.

2 n-Vector Fields on Surfaces

We begin with a description of our representation of n-vector fields on surfaces in the continuous setting; in Section 6 we discuss discretization on triangle meshes. Here and throughout $\imath$ denotes the imaginary unit. We use $|\cdot|$ to denote the magnitude of a complex number and "arg" to denote the angle it makes with the real axis. Recall that a vector $z \in \mathbb{C}$ rotated by an angle $\theta \in \mathbb{R}$ can be expressed as $e^{\imath\theta} z$.

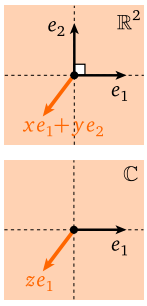

Real Planes and Complex Lines. Traditionally, a tangent space T_pM over a point p of a surface M is viewed as a copy of the real Euclidean plane $\mathbb{R}^2$. In this case, any tangent vector can be expressed as a real linear combination of two basis vectors e_1, e_2, *i.e.*, $xe_1 + ye_2$ for some pair of coefficients $x, y \in \mathbb{R}$. A useful alternative is to think of T_pM as a copy of the complex numbers $\mathbb{C}$, in which case any vector can be written as a *complex* multiple of a *single* basis vector e_1. (For instance, any point in $\mathbb{C}$ can be expressed as some complex number z times the real unit vector 1.) In this sense, $\mathbb{C}$ is a *one*-dimensional complex vector space, also known as the *complex line*.

n-Vector Fields. When viewed as a collection of complex lines, the tangent bundle TM is referred to as a *complex line bundle*, and a choice of unit basis vector $X_p \in T_pM$ at each point $p \in M$ is called a *basis section*. The *complex structure* J represents a 90-degree rotation in each tangent space. For surfaces in $\mathbb{R}^3$, J is induced by a quarter turn around the unit normal N, *i.e.*, for any tangent vector field Z we let $(JZ)_p := N_p \times Z_p$.

As suggested above, a vector field Z can be expressed as $Z = zX$ for some coefficient function $z : M \to \mathbb{C}$ relative to a basis section X. Likewise, any n-vector field ψ can be expressed as a collection of complex functions

$$\{e^{\imath 2k\pi/n} z, k = 0, \ldots, n-1\}$$

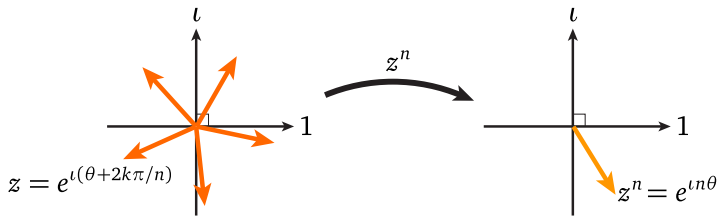

Figure 5. We represent tangent vectors as complex numbers; directions that differ in angle by some multiple of $2\pi/n$ become indistinguishable when raised to the n^{th} power.

which describes n copies of Z, each rotated by some integer multiple of $2\pi/n$. Alternatively, we can raise any of these functions to the n^{th} power yielding the single complex function

$$u := z^n,$$

as depicted in Figure 5. This function provides a concise representation of an n-vector field via $\psi = uX$. Moreover if $b = |u|$ and $\phi = \arg(u)$, we can then recover individual vector fields by computing the n^{th} roots

$$\{b^{1/n} e^{\iota(\phi/n+2k\pi/n)}, k = 0, \ldots, n-1\}.$$

Note that these functions carry meaning *only with respect to the chosen basis section*. In particular, we cannot treat them as ordinary scalar functions, but instead have to consider *parallel transport*, as discussed below. Formally, an n-vector field is a section of the n^{th} order tensor product $L := TM^{\otimes n}$ and the basis section X is naturally identified with a basis section of L.

Parallel Transport. The main benefit of this representation is that, at least within a given tangent space T_pM, we can measure the difference between two n-vectors via the simple quadratic expression $|u_1(p) - u_2(p)|^2$. When comparing vectors from different tangent spaces, however, we must be more careful – in particular, we must first map both vectors into a common tangent space via *parallel transport*. More explicitly, let X_p and X_q be basis vectors for the two tangent spaces T_pM and T_qM, respectively, and let γ be a geodesic from p to q, as depicted in Figure 6. If θ_p, θ_q are the angles made by γ with the two basis vectors and $\theta_{pq} := \theta_q - \theta_p$ is the difference between these angles, then the *parallel transport map* $P_{pq} \colon T_pM \to T_qM$ is given by

$$P_{pq}(z_p X_p) := e^{\iota\theta_{pq}} z_p X_q. \tag{2.1}$$

In other words, the coefficient z_p with respect to X_p gets mapped to the coefficient $e^{\iota\theta_{pq}} z_p$ with respect to X_q. Geometrically, the parallel transport map translates vectors from one tangent

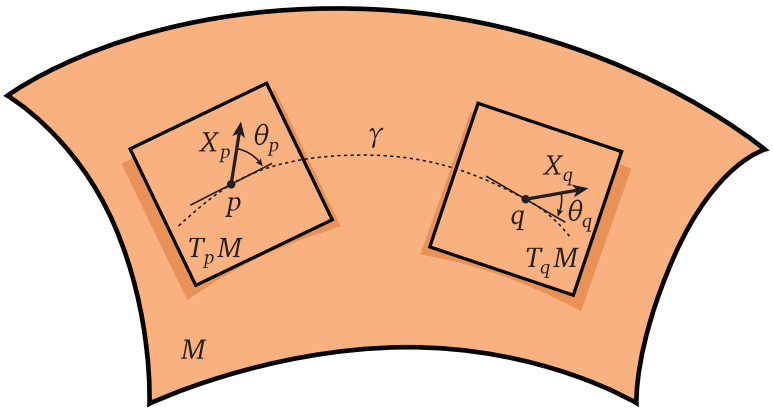

Figure 6. To transport a vector from one tangent space to another it is sufficient to measure the angle it makes against a geodesic connecting the two tangent spaces.

space to another without any in-plane rotation (Levi-Civita). The difference between two vectors is then defined as

$$|e^{\imath\theta_{pq}} z_p - z_q|^2,$$

and more generally we can write the difference between two n-vectors as

$$|(e^{\imath\theta_{pq}} z_p)^n - z_q^n|^2 = |e^{\imath n\theta_{pq}} z_p^n - z_q^n|^2 = |r_{pq} u_p - u_q|^2, \tag{2.2}$$

where for convenience we define the angle

$$\rho_{pq} := n\theta_{pq} \tag{2.3}$$

and the corresponding *transport coefficient*

$$r_{pq} := e^{\imath\rho_{pq}}.$$

Importantly, Equation (2.2) remains quadratic in z since the coefficient r is constant, depending only on the geometry of M, the choice of basis section X, and the degree n of the field. Note that by introducing the Levi-Civita connection, we also introduce a dependence on the Riemannian metric.

3 Smooth n-Direction Fields

We now focus on *n-direction* fields, *i.e.*, unit n-vector fields. The smoothness of a function is often measured via its *Dirichlet energy*; this energy can also be used for n-vector fields $\psi = uX$:

$$E_D(\psi) := \tfrac{1}{2}\int_M |\nabla\psi|^2 dA. \tag{3.1}$$

Here ∇ denotes the covariant derivative, *i.e.*, the Levi-Civita connection on M. For *n-direction* fields there are two main problems with this measure of smoothness. First, we must enforce the pointwise constraint $|u| = 1$, resulting in hard optimization problems as discussed in Section 1.1. Second, the Dirichlet energy of a *singular* n-direction field is not in general well-defined. In particular, consider that any vector field Z can be expressed as the product of a unit vector field $\tilde{Z}$ and a real nonnegative scale factor a, *i.e.*, $Z = a\tilde{Z}$. Since the change in a *unit* vector field is orthogonal to the field itself, we get $\nabla Z = (\nabla a)\tilde{Z} + a\omega J\tilde{Z}$, where ω is the rotation speed of Z. The Dirichlet energy of Z is then

$$\tfrac{1}{2}\int_M |\nabla Z|^2 dA = \tfrac{1}{2}\int_M |\nabla a|^2 + a^2|\omega|^2 dA = \tfrac{1}{2}\langle\langle(\Delta + |\omega|^2)a, a\rangle\rangle$$

where Δ is the positive semidefinite Laplace-Beltrami operator on M and $\langle\langle ., .\rangle\rangle$ denotes the L_2 inner product. For a *direction* field the scale factor is $a \equiv 1$, hence the Dirichlet energy is simply $\int_M |\omega|^2$. The rotation speed ω at a distance r from a singularity is proportional to $1/r$ (as $r \to 0$)

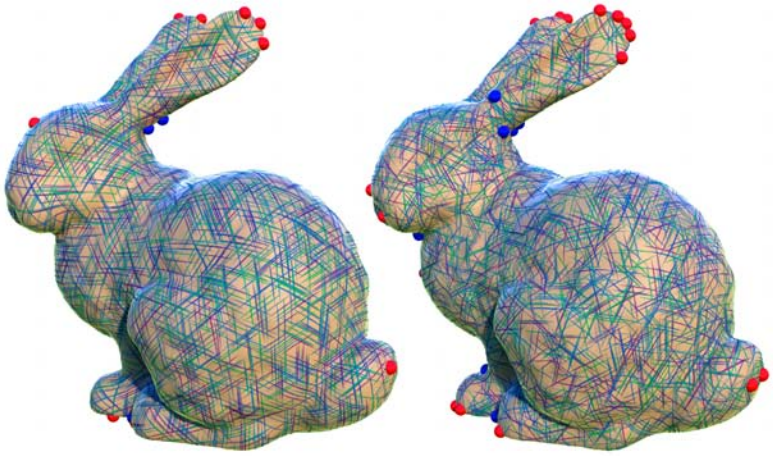

Figure 7. Our algorithm can produce fields of any degree – here $n = 3$ (left) and $n = 5$ (right) using the holomorphic energy (Section 4).

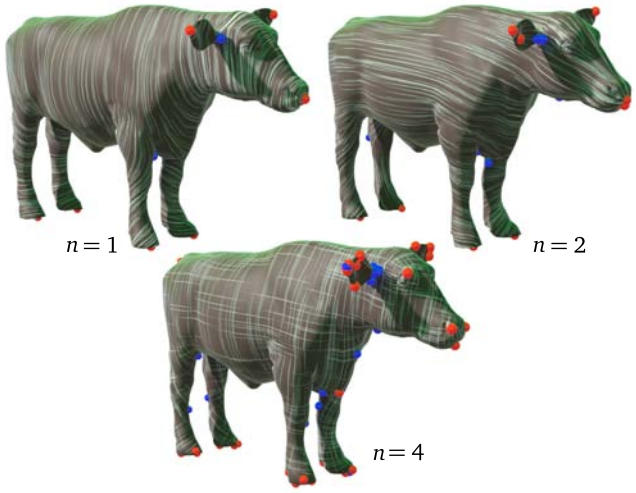

Figure 8. Smoothest direction fields for varying n, as measured by the holomorphic energy. For higher n singularities arise in groups, as if splitting from those for smaller n. (382ms, $|T| = 25$k)

since the total angle of rotation $2k\pi$ is divided by the circumference $2\pi r$ for a circle of radius r. *Consequently, the Dirichlet energy of a direction field with singularities is undefined (infinite).*

On a mesh, Dirichlet energy is finite but blows up under refinement. For instance, even if we add area weights to the energy $E(\phi, p)$ (Equation (1.1)), refinement around singularities increases energy so much that visibly inferior solutions yield lower energy (Figure 4). This observation motivates a different treatment of smoothness. In particular, we define the energy $\hat{E}$ of an n-direction field φ as the lowest Dirichlet energy among all n-vector fields $\psi = a\varphi$ parallel to φ. In other words, to evaluate the energy of a fixed unit field φ we look for an optimal scaling $a \geq 0$:

$$\hat{E}(\varphi) := \min_{a \geq 0,\, \|a\|=1} \int_M |\nabla(a\varphi)|^2 dA, \tag{3.2}$$

where the constraint $||a|| = 1$ prevents the trivial solution $a \equiv 0$.

At first glance this energy seems ill-defined, since the potential $|\omega|^2$ is unbounded for singular n-direction fields. However, minimizing the quadratic form $\langle\!\langle (\Delta + |\omega|^2)a, a \rangle\!\rangle$ over unit-norm functions a is equivalent to solving the eigenvalue problem

$$(\Delta + |\omega|^2)a = \lambda a$$

for the smallest eigenvalue λ and corresponding eigenfunction a. In physics, this equation is known as a *time-independent Schrödinger equation*, where the scalar function $|\omega|^2$ can be interpreted as the potential energy of a particle moving on the surface. The eigenfunction a with smallest eigenvalue is equivalent to the wave function of a free particle with lowest energy, which is always well-defined. From here, the globally optimal n-direction field is obtained by minimizing $\hat{E}(\varphi)$ over all unit fields φ:

$$\min_{|\varphi|=1} \hat{E}(\varphi) = \min_{|\varphi|=1} \left(\min_{a \geq 0,\, \|a\|=1} \int_M |\nabla(a\varphi)|^2 dA \right).$$

But since the set of rescaled n-direction fields is no different from the set of all n-vector fields, we can instead solve

$$\min_{||\psi||=1} \int_M |\nabla\psi|^2 dA,$$

which amounts to a smallest eigenvalue problem $\Delta\psi = \lambda\psi$, and recover the solution to our original problem via $a \leftarrow |\psi|$, $\varphi \leftarrow \psi/a$. In the end, we obtain a smoothest n-direction field (among all possible configurations of singularities) via a surprisingly simple procedure: find the smoothest n-vector field and normalize the resulting vectors. The physical interpretation of $\hat{E}(\varphi)$ also provides some insight into why the globally smoothest n-direction field has a small (though not necessarily minimal) number of singularities: adding too many poles to ω increases the potential energy $|\omega|^2$,

hence it also increases the ground-state energy of the corresponding wave function. We emphasize, however, that we never explicitly build the Schrödinger operator $\Delta + |\omega|^2$ or minimize the energy $\hat{E}$ – these objects are introduced simply to analyze the intermediate variational problem. In practice we need only build the Laplacian Δ used in the final eigenvalue problem.

Fig. 8 visualizes the fields we obtain by varying n on the same shape; Figure 7 shows more exotic examples.

4 Quadratic Smoothness Energies

The most commonly used smoothness energy is the Dirichlet energy E_D. However, we can obtain a richer family of energies by orthogonally splitting the covariant derivative into a sum of Cauchy-Riemann derivatives, *i.e.*, $\nabla\psi = \bar{\partial}\psi + \partial\psi$ where

$$\bar{\partial}_Z\psi := \tfrac{1}{2}\left(\nabla_Z\psi + J\nabla_{JZ}\psi\right), \quad \partial_Z\psi := \tfrac{1}{2}\left(\nabla_Z\psi - J\nabla_{JZ}\psi\right),$$

and Z is an arbitrary vector field (see **[75]** and **[2**, p. 27]). We call an n-vector field *holomorphic* if $\bar{\partial}_Z\psi = 0$ for all vector fields Z, and *anti-holomorphic* if $\partial_Z\psi = 0$ **[52**, Ch. 2.4]. These definitions mirror the standard notion that a function $f : \mathbb{C} \to \mathbb{C}$ is holomorphic (*i.e.*, angle- and orientation-preserving) if it satisfies the Cauchy-Riemann equation $0 = \bar{\partial}f = \frac{1}{2}(\frac{\partial}{\partial x} + \imath\frac{\partial}{\partial y})f = 0$.

Using the orthogonal splitting described above, the Dirichlet energy decomposes into holomorphic and anti-holomorphic terms E_H and E_A, respectively:

$$E_D(\psi) = E_H(\psi) + E_A(\psi) := \tfrac{1}{2}\int_M |\bar{\partial}\psi|^2 dA + \tfrac{1}{2}\int_M |\partial\psi|^2 dA.$$

We therefore define our smoothness energy as

$$E_s := (1+s)E_H + (1-s)E_A = E_D - s(E_A - E_H), \tag{4.1}$$

providing a continuum from anti-holomorphic ($s = -1$) to Dirichlet ($s = 0$) to holomorphic ($s = 1$) energy. Equivalently, the parameter s controls the deviation of E_s from the standard Dirichlet energy by the difference

$$E_A(\psi) - E_H(\psi) = \tfrac{1}{2}\int_M nK|\psi|^2 dA - \tfrac{1}{2}\int_{\partial M} \mathrm{Im}\langle\nabla\psi, \psi\rangle,$$

where K denotes Gaussian curvature (Appendix C). The potential in the Schrödinger operator now becomes $|\omega|^2 - \frac{s}{2}nK$, biasing singularities towards areas of high ($s < 0$) or low ($s > 0$) Gaussian curvature – this is the sense in which our smoothness energy is "geometry aware." More precisely, direction fields obtained for $s = \pm 1$ depend only on the conformal structure; for all other s, results depend on the metric. As shown in Fig. 9, the parameter s provides a useful tradeoff between number of singularities and "straightness" (*i.e.*, geodesic curvature) of integral curves.

Equation to Solve. Let A be the positive semidefinite quadratic form given by $E_s(\psi) = \frac{1}{2}\langle\langle A\psi, \psi\rangle\rangle$. We find a global minimizer ψ of E_s among fields with unit norm by solving the problem

$$A\psi = \lambda\psi \tag{4.2}$$

for the smallest eigenvalue $\lambda \in \mathbb{R}$.

Discussion. The solution ψ to Equation (4.2) is unique up to a complex constant with unit norm – this constant determines the global "phase," *i.e.*, it rotates every tangent vector by the same angle. In the case where λ is a repeated eigenvalue we obtain one of many possible solutions; note that all such solutions are global minimizers of our smoothness energy. Alignment constraints (Section 5) can be used to specify additional criteria.

A 1-vector field Y on M is anti-holomorphic if and only if the associated real-valued 1-form $X \mapsto \langle Y, X\rangle$ is harmonic, where $\langle\cdot,\cdot\rangle$ denotes the Riemannian metric. Likewise, since vector fields on any open set $U \subset \mathbb{C}$ can be naturally identified with complex functions, we can say that Y is holomorphic whenever it can be expressed as a holomorphic function in each local coordinate chart. These relationships help us make the connection with existing methods from computer graphics. For instance, holomorphic *1-forms* (corresponding to anti-holomorphic 1-vector fields)

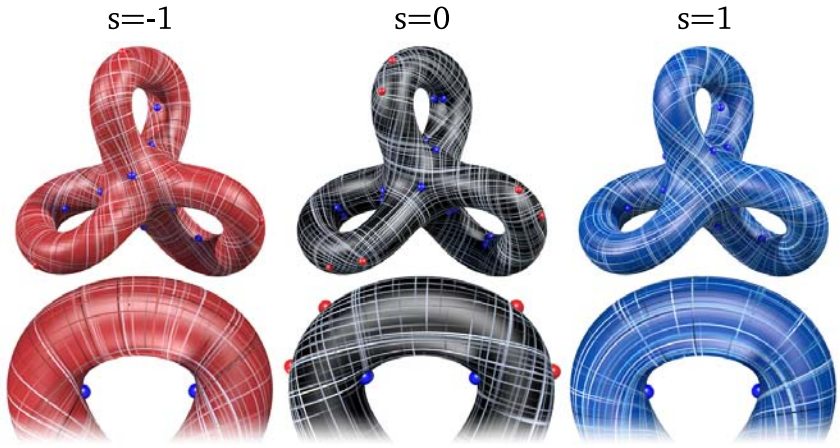

Figure 9. When seeking a smooth vector field, one often faces a choice: fewer singularities, or straighter field lines? Our method offers a tradeoff between the two, determined by a continuously varying parameter s. As seen above, $s = 1$ (holomorphic energy) typically yields fewer singularities, whereas $s = 0$ (Dirichlet energy) gives straighter lines; $s = -1$ (anti-holomorphic energy) provides a compromise between the two. (405ms, $|T| = 25$k)

appear in **[24]** in the context of parameterization, while the holomorphic energy for *functions* was first used in the context of parameterization by Lévy *et al.* **[42]** and Desbrun *et al.* **[19]**, where it was referred to as the conformal energy. Anti-holomorphic energy was used by Fisher *et al.* **[22]** in a DEC **[18]** context to find smoothest 1-vector fields as minimizers of the Hodge-Laplacian $\frac{1}{2}\int_M |\nabla \times Z|^2 + |\nabla \cdot Z|^2 = \int_M |\partial Z|^2 = 2E_A(Z)$. The holomorphic energy of 1-vector fields also appeared in the work of Ben-Chen *et al.* **[4**, Eq. 8**]**, where the Killing energy of a 1-vector field is defined as $E_K(Z) = 4\int_M \frac{1}{2}|\nabla \cdot Z|^2 + |\bar{\partial} Z|^2$, encoding the fact that a Killing 1-vector field is both holomorphic and divergence free.

5 Alignment Energies

It is often desirable to balance smoothness and alignment with a given field ϕ. Alignment is accomplished via the functional

$$E_l(\psi) = \int_M \mathrm{Re}\left(\langle \phi, \psi \rangle\right) dA = \mathrm{Re}\left(\langle\!\langle \phi, \psi \rangle\!\rangle\right).$$

where ϕ is normalized so that $||\phi|| = 1$. We then let

$$E_{s,t}(\psi) := (1-t)E_s(\psi) - tE_l(\psi) \tag{5.1}$$

where $t \in [0, 1]$ controls the strength of alignment. Just as before we minimize $E_{s,t}$ over all fields ψ with $\|\psi\| = 1$. The pointwise magnitude $|\phi|$ determines the local weighting between alignment and smoothness terms. For example, setting $\phi = 0$ on regions with missing or unreliable information will smoothly interpolate data from regions where $|\phi| > 0$.

Equation to Solve. Let A be the quadratic form corresponding to $E_{s,t}$. We find the global minimizer of $E_{s,t}$ by solving

$$(A - \lambda_t I)\widetilde{\psi} = \phi, \tag{5.2}$$

and normalizing the result, *i.e.*, $\psi \leftarrow \widetilde{\psi}/\|\widetilde{\psi}\|$ (Appendix A). Letting λ_1 be the smallest eigenvalue of A, the parameter $\lambda_t \in (-\infty, \lambda_1)$ controls the tradeoff between alignment ($\lambda_t \to -\infty$ for $t \to 1$) and smoothness ($\lambda_t \to \lambda_1$ for $t \to 0$), as illustrated in Figure 10. See Appendix A for further discussion.

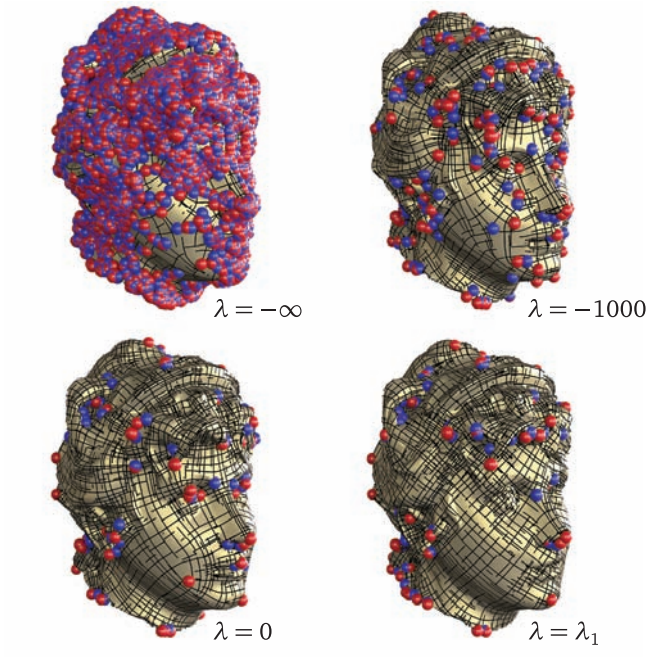

Figure 10. Adding an alignment term, we can interpolate between a curvature aligned field (upper left) and a smooth field (lower right). Here $s = 0$ and λ_1 is the smallest eigenvalue; we use principal curvature as a guidance field. (498ms, $|T| = 65$k)

6 Computation on Triangle Meshes

So far we have described our approach in the smooth setting. In this section we collect all details needed for numerical computation on triangle meshes, which are based on a finite element discretization using piecewise linear (PL) basis sections.

Triangle Meshes. We assume that our input is an oriented 2-manifold triangle mesh of arbitrary topology with or without boundary. We use V, E, and T to denote the set of vertices v_i, edges e_{ij} and triangles t_{ijk}, respectively. Vertices have coordinates $p_i \in \mathbb{R}^3$, with the rest of the surface defined via linear interpolation.

Vertex Tangent Spaces and Parallel Transport. Each vertex v_i carries a unit basis vector X_i (for convenience we use one of the incident edge directions) and a coefficient $u_i \in \mathbb{C}$ representing the associated n-vector with respect to this basis. To measure angles in tangent spaces, we rescale Euclidean angles at each vertex to sum to 2π **[57, 79]**. This rescaling effectively "flattens" the vertices, pushing their curvature into the incident triangles $t_{ijk} \ni i$. Let

$$s_i := \frac{2\pi}{\sum_{t_{ijk} \ni i} \alpha_i^{jk}}, \tag{6.1}$$

where α_i^{jk} denotes the Euclidean angle at v_i in t_{ijk}, opposite e_{jk}. (On the boundary, $s_i := 1$.) For parallel transport from v_i to v_j let

$$\rho_{ij} = n\left(\theta_j(X_j, e_{ij}) - \theta_i(X_i, e_{ij})\right) \tag{6.2}$$

(*cf.* Equation (2.3)), with $\theta_i(X_i, e_{ij})$ denoting the *rescaled Euclidean angle* from X_i to e_{ij} in T_iM (and similarly for θ_j). Each oriented edge of the mesh is assigned the transport coefficient $r_{ij} = e^{\imath\rho_{ij}}$, which is needed for matrix assembly.

Curvature. Transporting an n-vector around the boundary of t_{ijk} computes the *holonomy* $\Omega_{ijk} \in (-\pi, \pi]$, which represents the curvature of our line bundle L over t_{ijk}:

$$e^{\imath\Omega_{ijk}} := r_{ij} r_{jk} r_{ki}. \tag{6.3}$$

For each triangle we compute this value as $\Omega_{ijk} = \arg(r_{ij} r_{jk} r_{ki})$, which is needed for matrix assembly (Section 6.1.1). Note that this definition implicitly assumes that curvature is constant in each triangle, and thus proportional to area – conceptually

$$\Omega_{ijk} = \int_{t_{ijk}} n K dA = n K_{ijk} |t_{ijk}| \tag{6.4}$$

where $|t_{ijk}|$ and K_{ijk} are the area and Gaussian curvature in triangle t_{ijk}.

PL Finite Elements. For purposes of the finite element method the unit basis vectors at vertices are parallel transported into the incident triangles and attenuated linearly with the standard hat function. This gives PL basis sections Ψ_i supported on the triangles t_{ijk} incident to v_i.

From now on all fields ψ are piecewise linear, given as complex linear combinations of basis sections (Appendix D.1)

$$\psi = \sum_{v_i \in V} u_i \Psi_i,$$

with u now denoting the length $|V|$ vector of coefficients of ψ.

6.1 Algorithms

Smoothest Field. The continuous problem of Equation (4.2) turns into the generalized matrix eigenvector problem

$$Au = \lambda M u. \tag{6.5}$$

Here A is now a $|V| \times |V|$ Hermitian matrix representing the energy E_s with respect to the PL basis sections Ψ_i

$$A_{ij} = \langle\langle A\Psi_i, \Psi_j \rangle\rangle,$$

while M is the Hermitian mass matrix

$$M_{ij} = \langle\langle \Psi_i, \Psi_j \rangle\rangle.$$

Since we care about an eigenvector belonging to the smallest eigenvalue, we use an inverse power iteration to solve Equation (6.5) – see Section 7 for further details.

Aligned Field. The continuous problem of Equation (5.2) turns into the matrix problem

$$(A - \lambda_t M)\tilde{u} = Mq, \tag{6.6}$$

with A and M as above, while q is the coefficient vector for the PL version of the guidance field ϕ. The unit vector u minimizing the discretized version of $E_{s,t}$ is given by $u = \tilde{u}/\|\tilde{u}\|$ for $t = (1 + \|\tilde{u}\|)^{-1}$. While λ_t depends on t monotonically, we have no closed form expression for the relationship. In practice we find that $\lambda_t = 0$ is often a good (starting) value. To solve this problem we use Cholesky factorization followed by back substitution.

6.1.1 Matrix Entries

Because the entries of A and M are given as integrals we can assemble them via a loop over all triangles, *i.e.*, a sum of integrals over t_{ijk}. The entries of these *local* 3×3 matrices can be computed in closed form (see Appendix D for their derivation). Writing $\langle\langle ., . \rangle\rangle_{ijk}$ the subscript reflects the integration over t_{ijk} only. Assembly into the global matrices through summation is assumed.

Mass Matrix. On triangle t_{ijk} the local mass matrix induced by the L_2 inner product is given by (Appendix D.2)

$$\begin{aligned} M_{ii} &= \langle\langle \Psi_i, \Psi_i \rangle\rangle_{ijk} = \tfrac{1}{6}|t_{ijk}| \\ M_{jk} &= \langle\langle \Psi_j, \Psi_k \rangle\rangle_{ijk} = \bar{r}_{jk}|t_{ijk}| \frac{6e^{\imath\Omega_{ijk}} - 6 - 6\imath\Omega_{ijk} + 3\Omega_{ijk}^2 + \imath\Omega_{ijk}^3}{3\Omega_{ijk}^4}. \end{aligned} \tag{6.7}$$

The right hand side fraction has a removable singularity in the limit of no curvature ($\Omega_{ijk} \to 0$) where it takes on the value $1/12$ as expected.

Energy Matrix. We begin with the Laplacian, *i.e.*, the Dirichlet energy terms (Appendix D.3)

$$\begin{aligned}\Delta_{ii} &= \langle\!\langle \nabla\Psi_i, \nabla\Psi_i \rangle\!\rangle_{ijk} = \tfrac{1}{4|t_{ijk}|}\left[|p_{jk}|^2 + \Omega_{ijk}^2 \tfrac{|p_{ij}|^2 + \langle p_{ij}, p_{ik}\rangle + |p_{ki}|^2}{90}\right] \\ \Delta_{jk} &= \langle\!\langle \nabla\Psi_j, \nabla\Psi_k \rangle\!\rangle_{ijk} \\ &= \tfrac{\bar{r}_{jk}}{|t_{ijk}|}\left[(|p_{ij}|^2 + |p_{ki}|^2) f_1(\Omega_{ijk}) + \langle p_{ij}, p_{ik}\rangle f_2(\Omega_{ijk})\right],\end{aligned}$$

where $p_{ij} := p_j - p_i$ is the edge vector along e_{ij} and

$$\begin{aligned}f_1(s) &:= \tfrac{1}{s^4}\left(3 + \imath s + \tfrac{s^4}{24} - \tfrac{\imath s^5}{60} + (-3 + 2\imath s + \tfrac{s^2}{2})e^{\imath s}\right) \\ f_2(s) &:= \tfrac{1}{s^4}\left(4 + \imath s - \tfrac{\imath s^3}{6} - \tfrac{s^4}{12} + \tfrac{\imath s^5}{30} + (-4 + 3\imath s + s^2)e^{\imath s}\right).\end{aligned}$$

As before the singularities in f_1 and f_2 are removable and

$$\langle\!\langle \nabla\Psi_j, \nabla\Psi_k \rangle\!\rangle_{ijk} \to -\tfrac{\bar{r}_{jk}\langle p_{ij}, p_{ik}\rangle}{4|t_{ijk}|} = -\bar{r}_{jk}\tfrac{1}{2}\cot\alpha_i^{jk} \quad \text{for} \quad \Omega_{ijk} \to 0,$$

recovering the standard cotan-Laplacian in the flat setting [**45**].

With this, the per triangle energy matrix as a function of the parameter $s \in [-1, 1]$ follows as (Appendix D.4)

$$\begin{aligned}A_{ii} &= \Delta_{ii} - s\tfrac{\Omega_{ijk}}{|t_{ijk}|}M_{ii} \\ A_{jk} &= \Delta_{jk} - s\left(\tfrac{\Omega_{ijk}}{|t_{ijk}|}M_{jk} - \epsilon_{jk}\tfrac{\imath\bar{r}_{jk}}{2}\right),\end{aligned} \tag{6.8}$$

where $\epsilon_{jk} = \pm 1$ according to the orientation (positive or negative) of e_{jk} with respect to t_{ijk}.

Boundaries. Assembling the energy matrix triangle by triangle we automatically account for the boundary of the mesh (Figure 11). In the case of the Dirichlet energy ($s = 0$) this amounts to zero Neumann conditions, while for $s \neq 0$ we have the boundary terms due to $E_A - E_H$ (Appendix C).

Figure 11. Assembly of local matrices per triangle into the global matrix automatically accounts for boundary conditions. Here an example using the anti-holomorphic energy and curvature alignment ($\lambda = 0$) to the minimum (left; $n = 2$) and both (right; $n = 4$) principal curvature directions. (976ms, $|T| = 106$k)

Numerical Evaluation. The expressions given above for the off-diagonal entries of the mass and Dirichlet matrices have removable singularities as $\Omega_{ijk} \to 0$. To avoid numerical difficulties and ensure efficient, reliable, and accurate (to machine precision) evaluation of these expressions, we employ Chebyshev expansions [**23**]. Due to the equiripple property of Chebyshev polynomials, these can guarantee an error on the order of the least significant bit in double precision over an entire interval. Since the expressions for M_{jk} and Δ_{jk} are only evaluated for $\Omega_{ijk} \in (-\pi, \pi]$ and due to the symmetry resp. anti-symmetry of their real resp. imaginary parts, it is sufficient to find Chebyshev expansions for arguments in $[0, \pi]$. We computed these expansions with the aid of Mathematica and include ready to use code in the ancillary material.

6.1.2 Curvature Alignment

An important practical example for the alignment energy $E_{s,t}$, seeks to compute smoothed versions of curvature lines ($n = 2$) and crosses ($n = 4$). In this case the alignment field

$$\phi = \sum_{v_i \in V} q_i \Psi_i$$

is a PL approximation of the Hopf differential, *i.e.*, the trace-free part of the shape operator. In the discrete setting it is only accessible as a distribution ϕ^δ concentrated along edges [**14**]. Pairing it with our basis sections Ψ_i we compute coefficients (Appendix D.5)

$$\tilde{q}_i = \int_{t_{ijk} \ni i} \phi^\delta \Psi_i dA = -\tfrac{1}{4} \sum_{e \ni i} r_{ie} \beta_e |p_e|,$$

where β_e denotes the dihedral angle at e and the transport coefficients $r_{ie} = e^{\imath 2\theta_i(X_i, e)}$ depend on the rescaled Euclidean angle between X_i and e. The coefficients q of ϕ are the solution of

$$Mq = \tilde{q}.$$

Note that $|q|$ is proportional to the squared difference of principal curvatures $(\kappa_1 - \kappa_2)^2$. In umbilic regions, where principal directions are ill-defined, smoothness is therefore automatically favored over alignment (at least for $t < 1$ – see Equation (5.1)).

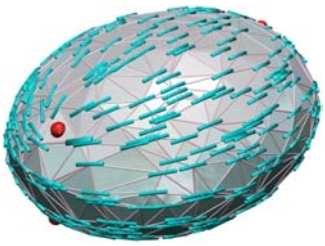

An example of this behavior is demonstrated in the inset image, where q is extremely noisy due to irregular tessellation. Due to the smoothing term we still obtain good results, reliably producing principal curvature directions and the four expected umbilics. In this example we used $-q$ to align with minimum curvature directions; for 4-direction fields we can use q_i^2. Figure 10 shows an example of the global trade off between smoothness and alignment and how it controls spurious singularities. Several results in Fig. 12 demonstrate further examples of alignment.

6.1.3 Index Computation

Given an n-direction field ψ we want to label each triangle t by an integer $\operatorname{index}_t \psi$ (an index of ± 1 indicating the presence of a singularity). Let ψ be given by complex numbers u_i of norm one for each vertex. Then for each edge e_{ij} we define the *rotation angle* of ψ as the unique number $\omega_{ij} \in (-\pi, \pi)$ such that $u_j = e^{\imath \omega_{ij}} r_{ij} u_i$, and define

$$\operatorname{index}_t \psi := \tfrac{1}{2\pi}(\omega_{ij} + \omega_{jk} + \omega_{ki} + \Omega_{ijk}) \in \{-1, 0, 1\}. \tag{6.9}$$

See Appendix B for more details on this definition and a proof of the associated discrete Poincaré-Hopf theorem. Singularities of index $+1$ and -1 are plotted as red and blue spheres, respectively, in all images.

7 Evaluation

To evaluate our algorithm, we compared it to the state-of-the-art *mixed integer* method of Bommes *et al.* [**5**]. We implemented both algorithms in a common C++ framework, using CoMISo [**6**] for the mixed integer problem, and CHOLMOD [**10**] for linear systems required in our method. To avoid factorization issues in case A is rank deficient we add a small multiple ($\varepsilon = 10^{-8}$) of M to

Algorithm 1 Setup

Input: $(\{V, E, T\}, s \in (-1, 1), n \in \mathbb{N})$
Output: (M, A)
Compute s_i at each vertex. ▷ Equation (6.1)
Pick arbitrary unit vector X_i at each vertex. ▷ Section 2
Compute transport at edges $r_{ij} \leftarrow e^{\iota \rho_{ij}}$. ▷ Equation (6.2)
Compute curvatures $\Omega_{ijk} \leftarrow \arg(r_{ij} r_{jk} r_{ki})$. ▷ Equation (6.3)
Assemble M. ▷ Equation (6.7)
Assemble A. ▷ Equation (6.8)
$A \leftarrow A + \varepsilon M$ ▷ Section 7

Algorithm 2 Smoothest Field (Section 6.1)

Input: (M, A)
Output: u
$LL^T \leftarrow \text{Cholesky}(A)$
$u \leftarrow \text{UniformRand}(-1, 1) \in \mathbb{C}^{|V|}$
for $i = 1$ to nPowerIterations **do**
$x \leftarrow \text{BackSubstitute}(LL^T, Mu)$ ▷ Equation (6.5)
$u \leftarrow x / \sqrt{x^T M x}$
end for

A prior to factorization; this shift does not change the eigenvectors. All timings were taken on the same 2.4GHz Intel Core 2 Duo machine. For our eigenvector problem (Equation (6.5)) we

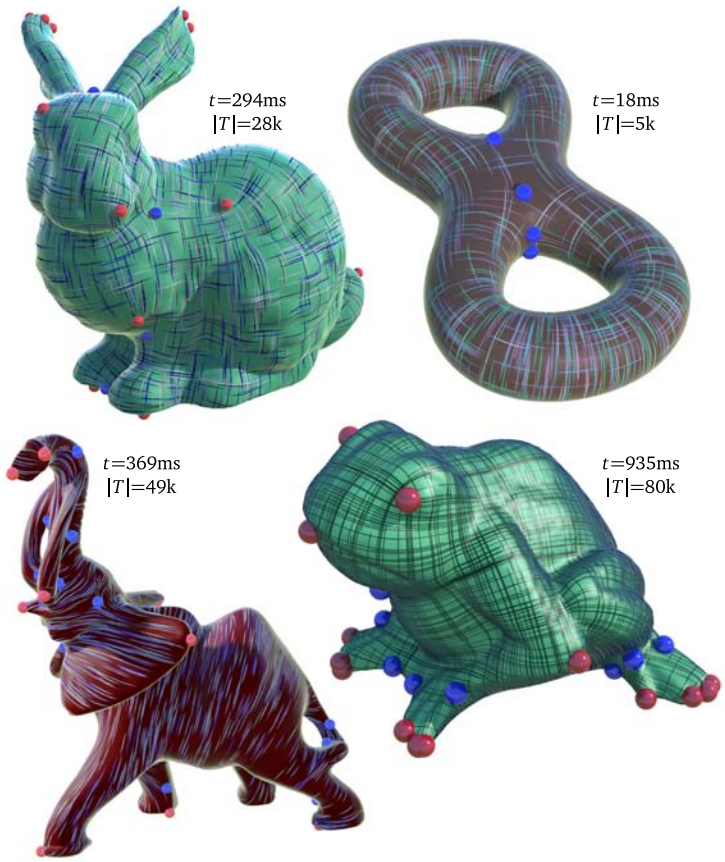

Figure 12. A gallery of examples. The bunny with a 4-direction field aligned ($\lambda = 0$) with (q_i^2) using $s = 1$. The two-holed torus aligning with (q_i^2) using $s = -1$. The frog shows a smoothest ($s = 1$) 4-direction field. The Elephant uses $s = -1$ and alignment ($\lambda = 0$) with the minimum curvature direction $-q$.

Algorithm 3 Aligned Field (Section 6.1)

Input: $(q, \lambda_t \in (-\infty, \lambda_0))$
Output: u
$LL^T \leftarrow \text{Cholesky}(A - \lambda_t M)$
$x \leftarrow \text{BackSubstitute}(LL^T, Mq)$ ▷ Equation (6.6)
$u \leftarrow x/\sqrt{x^T M x}$

factorized A and applied a fixed number of power iterations (20 for all examples shown in this paper). Note that since we need only the *smallest* eigenvalue, we do not require a sophisticated eigensolver like ARPACK, making our method particularly easy to implement. To minimize alignment energy we performed factorization and a single back-substitution. On a mesh of over 350k triangles, total time for our method was 11s for smoothing, 7s for alignment; the mixed integer method required more than three minutes for smoothing and just under three minutes for alignment. Figure 13 gives more detailed timing information; overall speedup was approx. 19x on average.

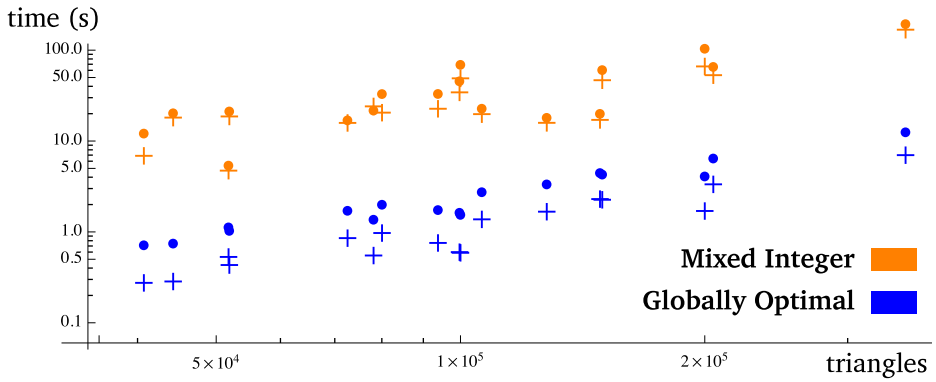

Figure 13. Log-log plot of computation time for a variety of common models. Dots and crosses correspond to smoothest and curvature-aligned fields, respectively. Overall, we observe an average speedup of 19.3x relative to the method of Bommes *et al.*

To assess the quality of our results we compared fields produced by *both* methods using the energy described in Equation (1.1), *i.e.*, the same energy minimized by Bommes *et al.*Note that our method produces smoother fields, even though we do not explicitly minimize this energy – in these examples we minimize our Dirichlet energy ($s = 0$), which is most similar to Equation (1.1). For curvature-aligned fields, we do not attempt a direct comparison of energies due to the large number of parameters involved in the method of Bommes *et al.*; we report only the number of singularities S and the time t.

Fig. 14 shows a rounded cube – here we achieve a lower energy than the mixed integer method both when optimizing smoothness (Dirichlet) and alignment. Even with no guidance field present (top row), our method places the expected 8 singularities at cube corners. A final set of examples is shown in Figure 15. On these more complex shapes the two methods find similar configurations of singularities, yet we achieve our results in *significantly* shorter time.

Robustness. Inspired by [**5**, Fig. 8], Fig. 16 shows the results of our curvature-aligned smoothing algorithm applied to an extremely poor triangulation (left) of an object with sharp features (here we use $s = 0$). Our method places singularities in the expected locations; adding noise *(middle)* or regularization *(right)* only slightly perturbs singularity locations in otherwise flat regions.

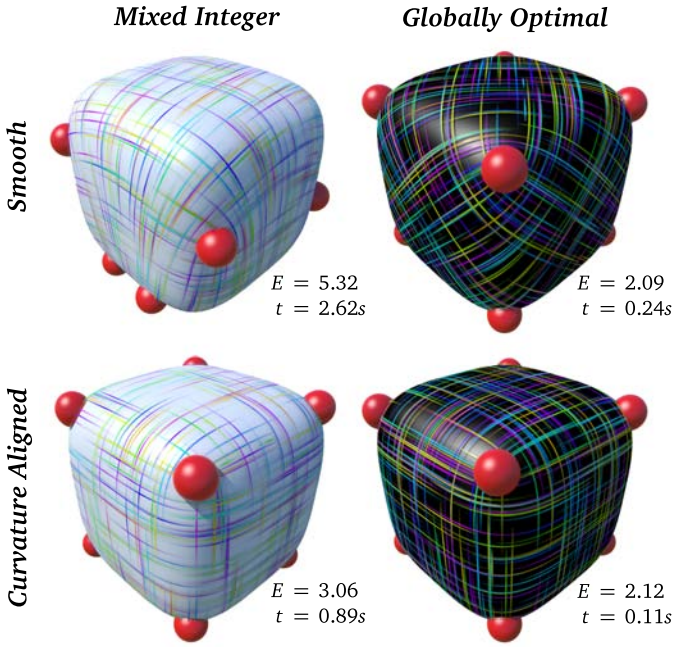

Figure 14. A comparison of results on the rounded cube. Smoothest fields (top row) and curvature aligned fields (bottom row) by method. Optimization of Equation (1.1) (left column) versus our Dirichlet energy (right column). In each case the energy according to Equation (1.1) is reported.

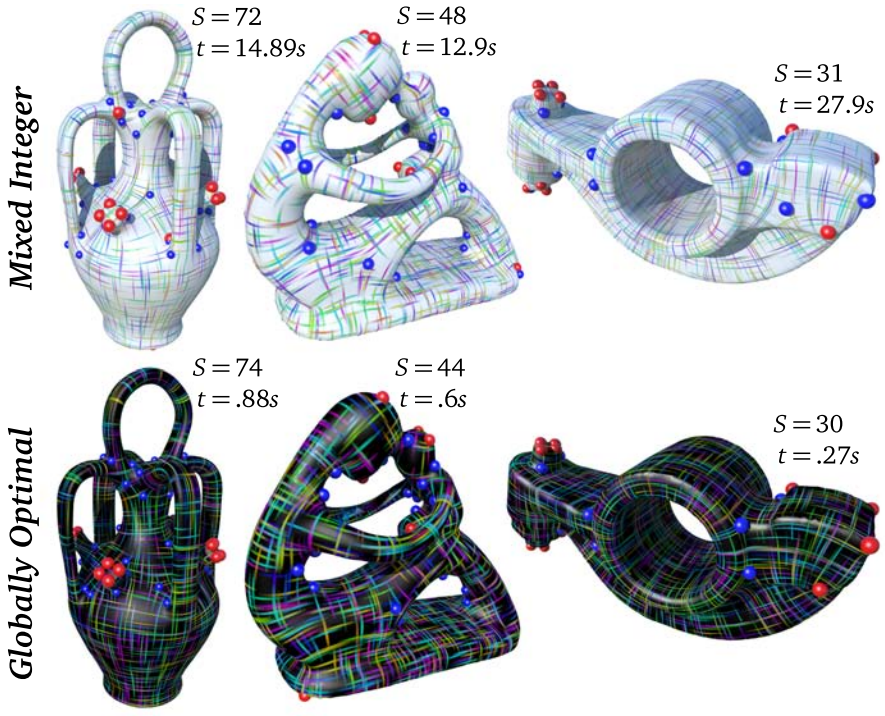

Figure 15. Comparison of results using optimization of Equation (1.1) (top) with minimization of our Dirichlet energy (bottom). Note in particular the difference in timings.

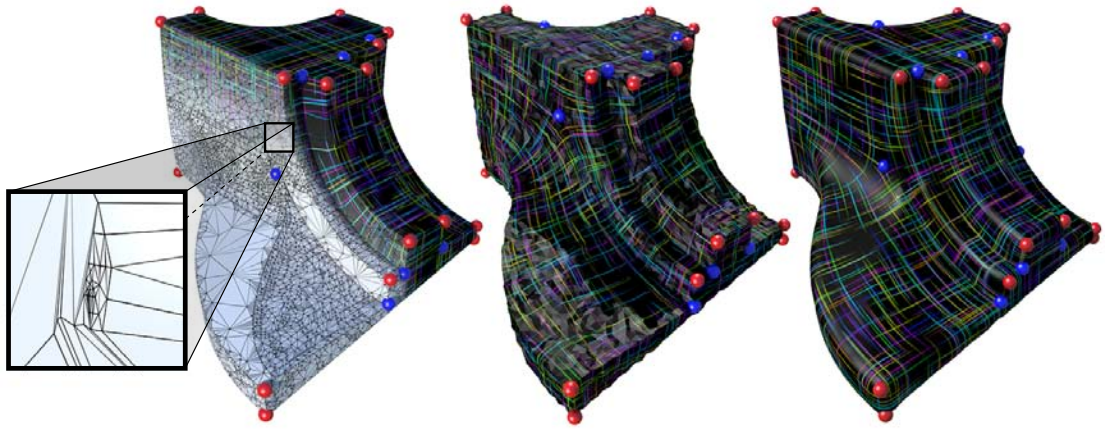

Figure 16. Our method is robust to ill-conditioned, noisy, and over-regularized meshes.

8 Conclusion

We have presented a principled formulation of n-direction fields on surfaces that naturally leads to efficient numerical algorithms. Although we have focused primarily on automatic singularity placement, our formulation provides a useful theory that can be applied in a variety of applications, especially in light of well-established connections with existing work (Section 4). These connections open up a number of interesting avenues for future investigation. For example, in some applications it is still desirable to edit fields manually in order to accommodate other (*e.g.*æsthetic) criteria. Editing could easily be achieved in our framework via a dualized version of the *trivial connections* algorithm of Crane *et al.* [**16**]. One might also use holomorphic and/or anti-holomorphic n-vector fields as the basis for surface parameterization, creating a direct path to quad meshing (*e.g.*, [**29**]). Different weighting schemes for the magnitude of the field ϕ may also permit the incorporation of other interesting criteria via the alignment term E_l.

Acknowledgments. This research was supported by a Google PhD Fellowship, the Hausdorff Research Institute for Mathematics, BMBF Research Project GEOMEC, SFB / Transregio 109 "Discretization in Geometry and Dynamics," and the TU München Institute for Advanced Study, funded by the German Excellence Initiative. Meshes provided by the Stanford Computer Graphics Laboratory and the AIM@SHAPE Shape Repository.

A Minimum of Alignment Energy

In Section 5 we claimed that the normalized solution $\psi = \tilde{\psi}/\|\tilde{\psi}\|$ of

$$(A - \lambda(t)I)\tilde{\psi} = \phi$$

for $\lambda(t) \in (-\infty, \lambda_1)$ yields a global minimum of $E_{s,t}(\psi)$ over all $\|\psi\| = 1$. Here we prove this fact and show the relationship between λ and t. For simplicity we assume throughout that $\|\phi\| = 1$. Furthermore, we will only use the real part of the product $\langle\langle ., . \rangle\rangle$ in this section.

At a constrained minimum the constraint must be fulfilled and the gradients of energy and constraint must be parallel, *i.e.*, there exists a Lagrange multiplier $\lambda/2$ such that

$$A\psi - \tfrac{t}{1-t}\phi = \lambda(t)\psi. \tag{A.1}$$

Substituting $\psi = \tilde{\psi}/\|\tilde{\psi}\|$ we find it to solve this equation for $t = (1 + \|\tilde{\psi}\|)^{-1}$.

Notice now how for $t \to 0$, ψ approaches the eigenvector belonging to the smallest eigenvalue λ_1 (we are working with a positive semidefinite Hermitian operator), since we are looking for a minimum of the energy. To see the behavior for $t \to 1$ consider the inner product of Equation (A.1) with ψ

$$\langle\langle A\psi, \psi \rangle\rangle = \lambda(t) + \tfrac{t}{1-t}\langle\langle \phi, \psi \rangle\rangle.$$

Together with

$$\langle\langle A\psi, \psi\rangle\rangle \leq \langle\langle A\phi, \phi\rangle\rangle,$$

which we will show in a moment, we get

$$\lambda(t) \leq \langle\langle A\phi, \phi\rangle\rangle - \tfrac{t}{1-t}\langle\langle \phi, \psi\rangle\rangle.$$

As $t \to 1$ the scalar product $\langle\langle \phi, \psi\rangle\rangle \to 1$ and so $\lambda(t) \to -\infty$.

To see that

$$\langle\langle A\psi, \psi\rangle\rangle \leq \langle\langle A\phi, \phi\rangle\rangle$$

we consider that for a minimizer ψ

$$\langle\langle A\psi, \psi\rangle\rangle - \tfrac{t}{1-t}\langle\langle \phi, \psi\rangle\rangle \leq \langle\langle A\phi, \phi\rangle\rangle - \tfrac{t}{1-t}\|\phi\|^2.$$

With $\|\phi\| = 1$ we get

$$\langle\langle A\psi, \psi\rangle\rangle \leq \langle\langle A\phi, \phi\rangle\rangle - \tfrac{t}{1-t}(1 - \langle\langle \phi, \psi\rangle\rangle)$$

and since $\langle\langle \phi, \psi\rangle\rangle \leq 1$ the desired bound.

Together we see that as $t \to 0$, $\lambda(t) \to \lambda_1$ and for $t \to 1$, $\lambda(t) \to -\infty$.

Finally we show that there exists a strictly monotone analytic function $t(\lambda) : (-\infty, \lambda_1) \to (0, 1)$.

Since the resolvent

$$R(\lambda) = (A - \lambda I)^{-1}$$

is a bounded self-adjoint operator which is an analytic function of λ so is

$$\psi(\lambda) = \frac{R(\lambda)\phi}{\|R(\lambda)\phi\|}$$

and $t(\lambda) = (1 + \|R(\lambda)\phi\|^{-1})$.

Now assume by contradiction that $t(\lambda)$ is not strictly monotone. Then there exists a $\lambda_0 \in (-\infty, \lambda_1)$ for which $t'(\lambda_0) = 0$. Taking the derivative at λ_0 of Equation (A.1) we find

$$(A - \lambda_0 I)\psi'(\lambda_0) - \psi(\lambda_0) = 0.$$

Taking the inner product of this equation with $\psi'(\lambda_0)$ the second term vanishes (since $\|\psi\| = 1$) leading to $\langle\langle (A - \lambda_0 I)\psi'(\lambda_0), \psi'(\lambda_0)\rangle\rangle = 0$, which contradicts the positive definiteness of $A - \lambda_0 I$.

Together these results show that solving

$$(A - \lambda(t)I)\tilde{\psi} = \phi$$

for a choice of $\lambda \in (-\infty, \lambda_1)$ finds a global minimizer of $E_{s,t}(\psi)$ for $t(\lambda)$, a strictly monotone function with range $(0, 1)$. For $\lambda(t) \to \lambda_1$, $t \to 0$ (*i.e.*, no alignment) while for $\lambda(t) \to -\infty$, $t \to 1$ (*i.e.*, no smoothing).

B Poincaré-Hopf

In this Section we assign to every n-direction field ψ an index $\operatorname{index}_t \psi \in \{-1, 0, 1\}$ on each triangle t. This is how we locate the singularities and label them as positive or negative. Under some smoothness assumption on the surface we will prove a discrete version of the Poincaré-Hopf theorem. A proof in the smooth case can be found in [**61**].

Recall that we describe the parallel transport of n-vectors by transport coefficients r_{ij} defined for each oriented edge e_{ij}. They can be thought of as an angle valued 1-form, *i.e.*, they satisfy $r_{ji} = r_{ij}^{-1}$. For each face t_{ijk} there is a unique real number $\Omega_{ijk} \in (-\pi, \pi)$ such that

$$r_{ij} r_{jk} r_{ki} = e^{\imath \Omega_{ijk}}.$$

We call Ω the curvature 2-form of the transport.

For a face $t = t_{ijk}$ we call the total curvature pushed into the triangle from its three vertices the *geometric curvature*

$$\sigma_t := s_i \alpha_i^{jk} + s_j \alpha_j^{ki} + s_k \alpha_k^{ij} - \pi$$

of t. A triangle mesh is called n-smooth if for each face we have

$$|\sigma_t| < \tfrac{\pi}{n}.$$

On an n-smooth mesh we have $\Omega_t = n\sigma_t$ and therefore a Gauß-Bonnet type theorem:

$$\sum_{t\in T} \Omega_t = 2n\pi\chi.$$

Here χ is the Euler characteristic of the mesh.

Let now ψ be an n-direction field given by complex numbers u_i of norm one for each vertex. Then for each edge e_{ij} we define the *rotation angle* of ψ as the unique number $\omega_{ij} \in (-\pi, \pi)$ such that

$$u_j = e^{\imath\omega_{ij}} r_{ij} u_i.$$

Then for each face $t = t_{ijk}$ we have $\omega_{ij}+\omega_{jk}+\omega_{ki}+\Omega_{ijk} \in (-4\pi, 4\pi)$ and $e^{\imath(\omega_{ij}+\omega_{jk}+\omega_{ki}+\Omega_{ijk})} = 1$. This means that for every face we have an integer

$$\operatorname{index}_t \psi := \tfrac{1}{2\pi}(\omega_{ij} + \omega_{jk} + \omega_{ki} + \Omega_{ijk}) \in \{-1, 0, 1\}.$$

In the language of Discrete Exterior Calculus this could be expressed as

$$\operatorname{index} \psi = \tfrac{1}{2\pi}(d\omega + \Omega)$$

If we sum this equation over all faces, the rotation angles cancel and we are left with the

Discrete Poincaré-Hopf theorem: On an n-smooth closed triangle mesh the index sum of every n-direction field equals $2n\chi$ where χ is the Euler characteristic of the mesh.

C Difference of Anti-Holomorphic and Holomorphic Energy

Here we show that for any smooth n-vector field ψ we have $E_A(\psi) - E_H(\psi) = \frac{1}{2}\int_M nK|\psi|^2 dA - \frac{1}{2}\int_{\partial M} \operatorname{Im}\langle\nabla\psi, \psi\rangle$. First we spell out all the requirements, each of which is fulfilled in our setting, of the theorem and a few facts that we will use during the derivation.

Let M be an oriented surface with Riemannian metric. Denote by J the complex structure on M. Suppose we have a complex line bundle L over M, *i.e.*, L is a 2-dimensional real vector bundle with a complex structure J. Our ψ are sections of such a line bundle. Assume that L comes with a complex connection ∇, meaning ∇ is compatible with J: $\nabla_X(J\psi) = J\nabla_X\psi$ for all sections ψ of L and all vector fields X. Then a section ψ is called holomorphic resp. anti-holomorphic if

$$0 = \bar\partial_X\psi := \tfrac{1}{2}(\nabla_X\psi + J\nabla_{JX}\psi)$$
$$0 = \partial_X\psi := \tfrac{1}{2}(\nabla_X\psi - J\nabla_{JX}\psi).$$

If ψ is a section of L, at each point $p \in M$ the linear map $\nabla\psi$ from the 2-dimensional real vector space T_pM into the 2-dimensional real vector space L_p can be decomposed into a complex linear part $(\partial\psi)_p$ and an anti-linear part $(\bar\partial\psi)_p$ and we can write

$$\nabla = \partial + \bar\partial.$$

Now suppose in addition that each L_P comes with a Hermitian scalar product $\langle .,. \rangle_p$ that is invariant under J and that ∇ is a metric connection, *i.e.*,

$$X\langle\psi, \varphi\rangle = \langle\nabla_X\psi, \varphi\rangle + \langle\psi, \nabla_X\varphi\rangle.$$

Replacing here X by some Lie bracket $[X, Y]$ we obtain

$$\begin{aligned}
&\langle\nabla_{[X,Y]}\psi, \varphi\rangle + \langle\psi, \nabla_{[X,Y]}\varphi\rangle = [X, Y]\langle\psi, \varphi\rangle \\
&\quad = X(\langle\nabla_Y\psi, \varphi\rangle + \langle\phi, \nabla_Y\varphi\rangle) - Y(\langle\nabla_X\psi, \varphi\rangle + \langle\psi, \nabla_X\varphi\rangle) \\
&\quad = \langle\nabla_X\nabla_Y\psi, \varphi\rangle + \langle\psi, \nabla_X\nabla_Y\varphi\rangle - \langle\nabla_Y\nabla_X\psi, \varphi\rangle - \langle\psi, \nabla_Y\nabla_X\varphi\rangle
\end{aligned}$$

Thus the curvature tensor R of L defined by

$$R(X, Y)\psi = \nabla_X\nabla_Y\psi - \nabla_Y\nabla_X\psi - \nabla_{[X,Y]}\psi$$

satisfies

$$\langle R(X, Y)\psi, \varphi\rangle + \langle\psi, R(X, Y)\varphi\rangle = 0.$$

In particular this implies that $\langle R(X,Y)\psi,\psi\rangle$ is always imaginary. Thus there is a real-valued 2-form Ω on M such that

$$R(X,Y)\psi = \Omega(X,Y)J\psi.$$

In case L is the tangent bundle TM, the Gaussian curvature K of M at a point p is defined in terms of a unit vector $X \in T_pM$ as

$$K = \langle R(X,JX)JX, X\rangle = -\Omega(X,JX).$$

Thus, if σ denotes the volume form, the curvature 2-form of the tangent bundle is $\Omega = -K\sigma$. We are mostly interested in the case $L = TM^{\otimes n}$, in which case we have $\Omega = -nK\,dA$ and therefore

$$R(X,Y)\psi = -nK\,dA(X,Y)J\psi.$$

The Dirichlet energy of a section ψ of L is defined as

$$E_D(\psi) = \tfrac{1}{2}\int_M |\nabla\psi|^2 dA,$$

where we view $\nabla\psi$ as a 1-form with values in L and the squared norm of such a 1-form ω at a point $p \in M$ is defined as

$$|\omega_p|^2 = |\omega(X)|^2 + |\omega(Y)|^2$$

where $\{X,Y\}$ form an orthonormal basis for T_pM.

Similarly, we define the holomorphic resp. anti-holomorphic energy of ψ as

$$E_H(\psi) = \tfrac{1}{2}\int_M |\bar\partial\psi|^2 dA \qquad E_A(\psi) = \tfrac{1}{2}\int_M |\partial\psi|^2 dA.$$

Theorem 1. *The holomorphic and anti-holomorphic energies of a section ψ are related as*

$$E_A(\psi) - E_H(\psi) = \tfrac{1}{2}\int_M nK|\psi|^2 dA - \tfrac{1}{2}\int_{\partial M} \mathrm{Im}\langle\nabla\psi,\psi\rangle.$$

Proof. Because of

$$\mathrm{Re}\langle\nabla\psi,\psi\rangle = \tfrac{1}{2}d\langle\psi,\psi\rangle$$

we can define a real-valued 1-form η on M via

$$\langle\nabla\psi,\psi\rangle = \tfrac{1}{2}d\langle\psi,\psi\rangle + \imath\eta(X).$$

Then for a locally defined unit vector field X we have

$$\begin{aligned}\imath d\eta(X,JX) &= X\langle\nabla_{JX}\psi,\psi\rangle - (JX)\langle\nabla_X\psi,\psi\rangle - \langle\nabla_{[X,JX]}\psi,\psi\rangle\\ &= \langle R(X,JX)\psi,\psi\rangle + \imath(\langle J\nabla_{JX}\psi,\nabla_X\psi\rangle + \langle\nabla_X\psi, J\nabla_{JX}\psi\rangle)\\ &= \imath(nK|\psi|^2 + |\bar\partial\psi|^2 - |\partial\psi|^2)dA(X,JX).\end{aligned}$$

Thus

$$d\eta = (nK|\psi|^2 + |\bar\partial\psi|^2 - |\partial\psi|^2)dA$$

and our claim follows by Stokes theorem. □

D Integrals

In this section we give a derivation of all the integrals needed in the finite element discretization of the smooth theory, and begin with the construction of the basis sections.

D.1 PL Basis Sections

The PL basis sections Ψ_i are defined by extending the basis X_i from each vertex into the incident triangles through parallel transport along radii, giving us a unit basis section Φ_i supported on all incident triangles. Φ_i is then linearly attenuated with the standard PL hat function at v_i to give us Ψ_i. On a single incident triangle, say t_{ijk}, this amounts to

$$\Psi_i = b_i \Phi_i,$$

where b_i is the standard barycentric coordinate function for v_i in t_{ijk}.

While this procedure, together with the fact that the curvature is defined everywhere, uniquely defines the PL basis sections Ψ_i we do not have a closed form expression for them. Yet we can work out all required integrals in closed form due of our earlier assumption that the curvature is constant in each triangle (Equation (6.4)). Specifically, the constancy of curvature over t_{ijk} gives us the holonomy angle $\Omega_{t'}$ (Equation (6.3)) along the boundary of any sub-triangle $t' \subset t_{ijk}$ as a fraction of the area:

$$\Omega_{t'} = \int_{t'} nK dA = nK_{ijk}|t'|,$$

where $|t'|$ denotes the area of t' (see Equation (6.4)).

This brings us to the starting position.

D.2 Mass Matrix

Consider the Hermitian product (complex anti-linear resp. linear in the left resp. right factor) of basis sections restricted to the triangle t_{ijk}

$$\langle\!\langle \Psi_j, \Psi_k \rangle\!\rangle_{ijk} := \int_{t_{ijk}} \langle \Psi_j, \Psi_k \rangle dA$$

These determine the L_2 metric on the space of PL sections over a triangle and define the local mass matrix.

Let v be inside t_{ijk} with barycentric coordinates b_i, b_j, b_k

$$v = b_i v_i + b_j v_j + b_k v_k,$$

and consider the integrand of the mass matrix as a function of v

$$\langle \Psi_j(v), \Psi_k(v) \rangle = b_j b_k \langle \Phi_j(v), \Phi_k(v) \rangle. \tag{D.1}$$

We know that Φ_j is parallel along e_{jk}, implying

$$\Phi_j(v_k) = r_{jk} \Phi_k(v_k)$$

(*cf.* Equation (2.1); Figure 17). Moreover Φ_k is parallel along the ray from v_k to v and Φ_j along

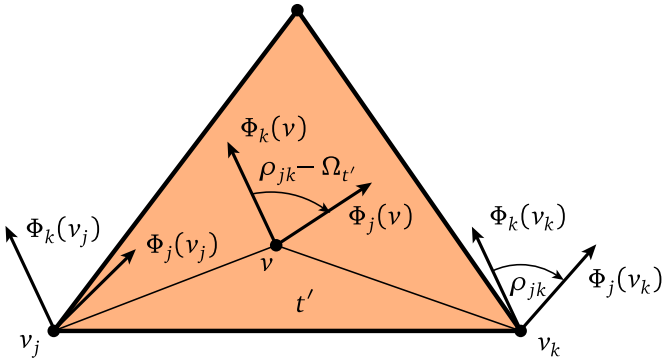

Figure 17. The angle between two basis sections at some point v in t_{ijk} can be deduced from the holonomy $\Omega_{t'}$ of the sub-triangle $t' = \{v_k, v, v_j\}$ and the known transport along e_{jk}.

the ray from v to v_j. Since parallel transport around $t' = \{v_k, v, v_j\}$, *i.e.*, from v_k to v on to v_j and finally back to v_k recovers the holonomy angle of t' we get

$$\langle \Phi_j(v), \Phi_k(v) \rangle r_{jk} = e^{\imath \Omega_{t'}} = e^{\imath \Omega_{ijk} b_i}. \tag{D.2}$$

(using Eqs. (6.3) and (6.4). Putting Eqs. (D.1) and (D.2) together and integrating symbolically we get

$$\langle\!\langle \Psi_j, \Psi_k \rangle\!\rangle_{ijk} = \bar{r}_{jk} |t_{ijk}| \frac{6e^{\imath\Omega_{ijk}} - 6 - 6\imath\Omega_{ijk} + 3\Omega_{ijk}^2 + \imath\Omega_{ijk}^3}{3\Omega_{ijk}^4}.$$

For the product of a basis section with itself the curvature does not play a role and one obtains $\|\Psi_i\|_{ijk}^2 = |t_{ijk}|/6$.

This completely determines the metric on the space of basis n-vector fields.

D.3 Dirichlet Energy

To compute the Dirichlet energy we need the covariant derivatives of Φ_j and Φ_k. To this end we will employ a particular (linear) parameterization $f(x, y)$ of the embedding p_{ijk} of t_{ijk} with $f(0,0) = p_i$, $f(1,0) = p_j$ and $f(0,1) = p_k$, and will treat Φ and Ψ as defined on the image of f in this section. Denote by ∂_x and ∂_y the tangent coordinate frame corresponding to the coordinates x, y (Figure 18).

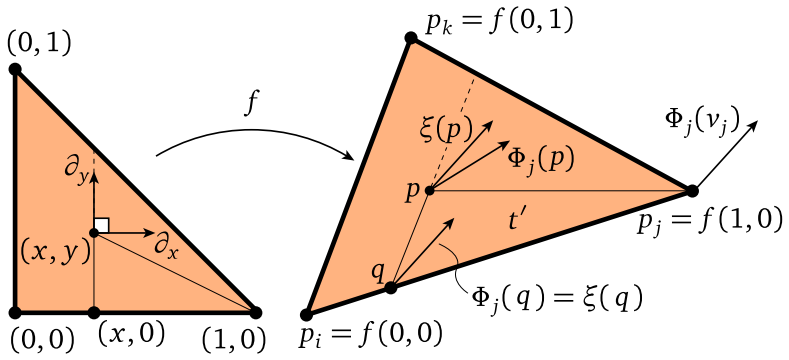

Figure 18. Using the parameterization f we compute the covariant derivative of a basis n-vector field Φ_j using the local section ξ.

To compute $\nabla_{\partial_y} \Phi_j$ at some point p in the interior of the triangle, let ξ be a section on the triangle t_{ijk} that agrees with Φ_j for $y = 0$ and is parallel along the lines $\{x = \text{const}\}$, giving $\nabla_{\partial_y} \xi = 0$. For $p = f(x, y)$ let $q := f(x, 0)$, then the holonomy around t_{pqj} is $\Omega_{pqj} = (1 - x)y\Omega_{ijk}$ and consequently

$$e^{\imath(1-x)y\Omega_{ijk}} \xi(p) = \Phi_j(p).$$

This gives us

$$\nabla_{\partial_y} \Phi_j = (\partial_y e^{\imath(1-x)y\Omega_{ijk}})\xi = \imath(1 - x)\Omega_{ijk}\Phi_j.$$

The derivative $\nabla_{\partial_x} \Phi_k$ follows immediately through interchange of x and y. Taking account of the orientation, we obtain

$$\nabla_{\partial_x} \Phi_k = -\imath(1 - y)\Omega_{ijk}\Phi_k.$$

Since by construction Φ_j is parallel along rays from v_j and Φ_k along rays from v_k we get a linear relationship between the covariant derivative of Φ_j with respect to ∂_x and ∂_y (and similarly for Φ_k)

$$0 = (1 - x)\nabla_{\partial_x} \Phi_j - y\nabla_{\partial_y} \Phi_j \quad 0 = x\nabla_{\partial_x} \Phi_k + (y - 1)\nabla_{\partial_y} \Phi_k$$

which give us all the covariant derivatives for Φ_j and Φ_k. For $\Psi_j = x\Phi_j$ and $\Psi_k = y\Phi_k$ this results in

$$\nabla_{\partial_x} \Psi_j = (1 + \imath\Omega_{ijk} xy)\Phi_j, \qquad \nabla_{\partial_y} \Psi_j = \imath\Omega x(1 - x)\Phi_j,$$
$$\nabla_{\partial_x} \Psi_k = -\imath\Omega_{ijk} y(1 - y)\Phi_k, \qquad \nabla_{\partial_y} \Psi_k = (1 - \imath\Omega_{ijk} xy)\Phi_k.$$

To simplify the computation of the integrals we now switch from the basis $\{\partial_x, \partial_y\}$ to an orthogonal basis $\{E_1, E_2\}$. Letting

$$g_{11} = |p_j - p_i|^2, \quad g_{12} = \langle p_j - p_i, p_k - p_i\rangle = g_{21}, \quad g_{22} = |p_k - p_i|^2,$$

the orthogonal basis follows as

$$E_1 = \tfrac{1}{\sqrt{g_{11}}}\partial_x, \quad E_2 = \tfrac{1}{2|t_{ijk}|\sqrt{g_{11}}}(g_{11}\partial_y - g_{12}\partial_x).$$

With respect to this basis we get

$$\begin{aligned}
\nabla_{E_1}\Psi_j &= \tfrac{1}{\sqrt{g_{11}}}(1 + \imath\Omega_{ijk}xy)\Phi_j,\\
\nabla_{E_2}\Psi_j &= \tfrac{1}{2|t_{ijk}|\sqrt{g_{11}}}(-g_{12} + \imath\Omega_{ijk}(g_{11}x(1-x) - g_{12}xy))\Phi_j,\\
\nabla_{E_1}\Psi_k &= \tfrac{1}{\sqrt{g_{11}}}(-\imath\Omega_{ijk}y(1-y))\Phi_k,\\
\nabla_{E_2}\Psi_k &= \tfrac{1}{2|t_{ijk}|\sqrt{g_{11}}}(g_{11} + \imath\Omega_{ijk}(g_{12}y(1-y) - g_{11}xy))\Phi_k.
\end{aligned}$$

With this we have all the necessary components to compute the Dirichlet products and integrate them to yield

$$\begin{aligned}
\langle\!\langle \nabla\Psi_j, \nabla\Psi_j\rangle\!\rangle_{ijk} =& \tfrac{1}{4|t_{ijk}|}\left[g_{22} + \Omega_{ijk}^2\tfrac{3g_{11}-3g_{12}+g_{22}}{90}\right],\\
\langle\!\langle \nabla\Psi_j, \nabla\Psi_k\rangle\!\rangle_{ijk} =& \tfrac{\bar r_{jk}}{|t_{ijk}|\Omega_{ijk}^4}[(3g_{11} + 4g_{12} + 3g_{22})\\
&+ \imath\Omega_{ijk}(g_{11} + g_{12} + g_{22}) - \imath\Omega_{ijk}^3\tfrac{g_{12}}{6}\\
&+ \Omega_{ijk}^4\tfrac{g_{11}-2g_{12}+g_{22}}{24} - \imath\Omega_{ijk}^5\tfrac{g_{11}-2g_{12}+g_{22}}{60}\\
&+ (-(3g_{11} + 4g_{12} + 3g_{22})\\
&+ \imath\Omega_{ijk}(2g_{11} + 3g_{12} + 2g_{22})\\
&+ \Omega_{ijk}^2\tfrac{g_{11}+2g_{12}+g_{22}}{2})e^{\imath\Omega_{ijk}}].
\end{aligned}$$

D.4 Boundary Terms

The proof in Appendix C applies to a single triangle t and we only need to compute the boundary term, since the $nK\|\Psi\|_{ijk}^2$ term is the curvature weighted mass matrix (Appendix D.2).

We begin by observing that the 1-form $\langle\nabla\Psi_j, \Psi_k\rangle$ is nonzero only on e_{jk}. Moreover, since $\langle\nabla\Psi_i, \Psi_i\rangle = d|\Psi_i|^2$, it is real for matching indices giving us

$$\int_{\partial t}\operatorname{Im}\langle\nabla\Psi, \Psi\rangle = \sum_{e_{ij}\in\partial t}\operatorname{Im}\int_{e_{ij}}\left(\bar\alpha_i\alpha_j\langle\nabla\Psi_i, \Psi_j\rangle + \bar\alpha_j\alpha_i\langle\nabla\Psi_j, \Psi_i\rangle\right).$$

Note also that $d\langle\Psi_i, \Psi_j\rangle = \langle\nabla\Psi_i, \Psi_j\rangle + \overline{\langle\nabla\Psi_j, \Psi_i\rangle}$. Since $\Psi_i(v_j) = 0 = \Psi_j(v_i)$ Stokes' theorem yields

$$0 = \int_{\partial e_{ij}}\langle\Psi_i, \Psi_j\rangle = \int_{e_{ij}}\langle\nabla\Psi_i, \Psi_j\rangle + \int_{e_{ij}}\overline{\langle\nabla\Psi_j, \Psi_i\rangle},$$

which we use to simplify the boundary edge sum

$$\int_{\partial t}\operatorname{Im}\langle\nabla\Psi, \Psi\rangle = \sum_{e_{ij}\in\partial t} 2\operatorname{Im}\left(\bar\alpha_i\alpha_j\int_{e_{ij}}\langle\nabla\Psi_i, \Psi_j\rangle\right).$$

We now turn to an individual edge term. Parameterize e_{jk} with a constant speed $\gamma\colon [0,1]\to M$, *i.e.*, $|\gamma'| = |p_{jk}|$. Due to $d\boldsymbol{b}_j(\gamma') = -1$ for b_j the barycentric coordinate function of v_j, and the parallelity of Φ_j along e_{jk}, we get

$$\nabla_{\gamma'}\Psi_j = d\boldsymbol{b}_j(\gamma')\Phi_j + b_j\nabla_{\gamma'}\Phi_j = -\Phi_j,$$

and hence

$$\int_{e_{jk}}\langle\nabla\Psi_j, \Psi_k\rangle = \int_0^1\langle\nabla_{\gamma'}\Psi_j, \Psi_k\rangle dt = -\int_0^1 b_k\langle\Phi_j, \Phi_k\rangle dt = -\tfrac{\bar r_{jk}}{2}.$$

This finally yields

$$\int_{\partial t}\operatorname{Im}\langle\nabla\Psi, \Psi\rangle = -\sum_{e_{jk}\in\partial t}\operatorname{Im}\left(\bar r_{jk}\bar\alpha_j\alpha_k\right)$$

Notice that for an edge e_{jk} with two incident triangles (not on the boundary) the corresponding terms from each triangle cancel out, leaving us with only a sum over the boundary of the triangle mesh.

In the flat case ($\Omega_{ijk} = 0$) these boundary terms are simply the areas of the range of ψ seen as a complex function, *i.e.*, $2E_H(\psi) = E_D(\psi) - A(\psi(M))$ (compare with [**48**, Eq. 2], resp. [**56**, Eq. 2.1]).

D.5 Curvature Alignment

To perform alignment with principal curvature directions we need to find a 2-vector corresponding to these directions (2-vector since principal curvature directions are indistinguishable under rotations by π). In the smooth setting, the principal curvature directions are the eigendirections of the shape operator S, given by the derivative of the Gauss map

$$dN = df \circ S.$$

The *Hopf differential* Q is the trace-free part of the shape operator

$$S = H \cdot I + Q,$$

where $H = (\kappa_1 + \kappa_2)/2$ is the mean curvature and I the identity. Q has the principal curvature directions as eigendirections with eigenvalues $(\kappa_1 - \kappa_2)/2$ and $(\kappa_2 - \kappa_1)/2$. Thus the information provided by Q contains the (unoriented) direction of maximal curvature together with a positive "intensity." Thus Q can be viewed as a 2-vector field.

Using a local complex coordinate z we can write any tangent vector as $a\frac{\partial}{\partial x}$ where x is the real part of z and $a \in \mathbb{C}$. Q is anti-linear as an endomorphism of the tangent space (it anti-commutes with J) and therefore we have

$$Q(a\tfrac{\partial}{\partial x}) = q\bar{a}\tfrac{\partial}{\partial x}.$$

So with respect to a complex coordinate Q is described by a complex function q.

In the discrete setting the shape operator survives in a distributional sense concentrated along edges [**14**]. Consider an edge e and its two incident triangles and map them isometrically to the plane with e along the x-axis, then

$$S_e = \delta_y \beta_e \tfrac{\partial}{\partial y} dy = \delta_y \tfrac{\beta_e}{2} \tfrac{\partial}{\partial x} dz - \delta_y \tfrac{\beta_e}{2} \tfrac{\partial}{\partial x} d\bar{z}.$$

Here δ_y is the delta distribution along across e, β_e the dihedral angle at e, $dz = dx + \imath dy$, and $d\bar{z} = dx - \imath dy$. Hence

$$q_e = -\delta_y \tfrac{\beta_e}{2}.$$

Since q exists only as a distribution, we treat it as a functional on the space of smooth sections. In particular we can pair it with each of our PL 2-vector basis sections

$$\tilde{q}_i := q(\Psi_i) = \sum_{e \ni i} q_e(\Psi_i) = -\tfrac{1}{4} \sum_{e \ni i} r_{ie} \beta_e |p_e|.$$

where the transport coefficient $r_{ie} = e^{\imath 2\theta_i(X_i, e)}$ depends on the rescaled Euclidean angle between X_i and e. Hence the coefficients q of the PL Hopf differential solve the matrix problem

$$Mq = \tilde{q}.$$

2 | Stripe Patterns on Surfaces

Felix Knöppel, Keenan Crane, Ulrich Pinkall, Peter Schröder
ACM Trans. Graph. 34, 4 (July), 2015 (to appear).

Abstract. *Stripe patterns* are ubiquitous in nature, describing macroscopic phenomena such as stripes on plants and animals, down to material impurities on the atomic scale. We propose a method for synthesizing stripe patterns on triangulated surfaces, where singularities are automatically inserted in order to achieve user-specified orientation and line spacing. Patterns are characterized as global minimizers of a convex-quadratic energy which is well-defined in the smooth setting. Computation amounts to finding the principal eigenvector of a symmetric positive-definite matrix with the same sparsity as the standard graph Laplacian. The resulting patterns are globally continuous, and can be applied to a variety of tasks in design and texture synthesis.

1 Introduction

A diverse collection of natural phenomena exhibit the same characteristic pattern: unoriented stripes of uniform width that bifurcate at isolated points called *edge dislocations* [**31**, pp. 44–46]. Why does this bifurcation occur? Surfaces of revolution (Fig. 3) provide an instructive example: for any given stripe width there is a maximum integer number of stripes that can fit around the circumference at each height h. As the radius r grows or shrinks this number can suddenly jump—stripes must therefore branch or coalesce in order to maintain a continuous transition between neighboring regions.

(Photo courtesy Joshua Hill)

Figure 1. Natural phenomena like cacti *(left)* exhibit characteristic branching patterns in an effort to maintain evenly-spaced features. *Right:* Our algorithm allows one to quickly and automatically synthesize similar patterns that seamlessly cover arbitrary surfaces.

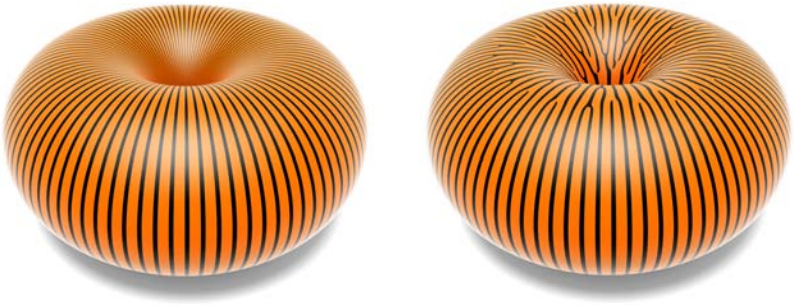

Figure 2. In general, patterns without singularities exhibit uneven spacing, and patterns with even spacing must have singularities.

Our approach to stripe pattern synthesis is closely related to existing methods for field-aligned parameterization (Sec. 1.1). The main point of departure is that existing algorithms try to *prevent* singularities that were not specified as part of the input, whereas we intentionally *allow* new singularities, which are essential to modeling natural phenomena. In particular, we use a formulation akin to Ray *et al.* **[59]**, but omit both the curl correction step and the unit-norm constraint. As a result, we can formulate our problem via a simple convex quadratic energy; the practical payoff is an algorithm about an order of magnitude faster than those based on nonconvex constraints or mixed integer programming (Sec. 5.1). The final algorithm is simple to implement (App. C) requiring no global cutting or integer jumps, is robust to errors in the input (Fig. 16), and requires little work beyond factoring a single sparse matrix. The output is an angle α at each triangle corner which can be used to drive a periodic texture or shader. We also introduce a novel interpolation scheme (Sec. 4.3) that ensures the pattern is globally continuous (C^0) away from a finite collection of isolated singular points.

1.1 Related Work

Synthesis of stripe patterns has a long history due to the prevalence of these patterns in nature (Figs. 1, 5, 7, 8, and 20). Placement of evenly-spaced stripes is also essential in vector field visualization **[28, 46, 66]** and nonphotorealistic rendering **[27]**. Methods based on reaction-diffusion equations can generate stripe patterns aligned with nonorientable features **[68, 76]**, but must solve the full time evolution of a nonlinear parabolic PDE whose parameters can be difficult to control; moreover, these methods solve directly for the color function rather than a parameterization, precluding more complex shaders like Fig. 1, right. Image-based texture synthesis can also produce patterns aligned with orientable **[58, 40]** and nonorientable vector fields **[78, 70, 69]**, but features in such patterns result from, *e.g.*, statistical similarity of image

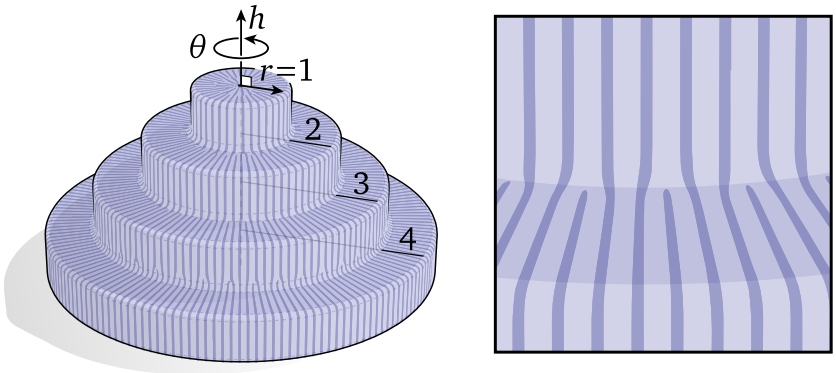

Figure 3. To maintain constant stripe width, each "level" of this surface of revolution must exhibit a different number of stripes. Judiciously placed branch points are therefore essential to keeping the pattern continuous. *Inset:* each pair of stripes from the second level perfectly branches into three.

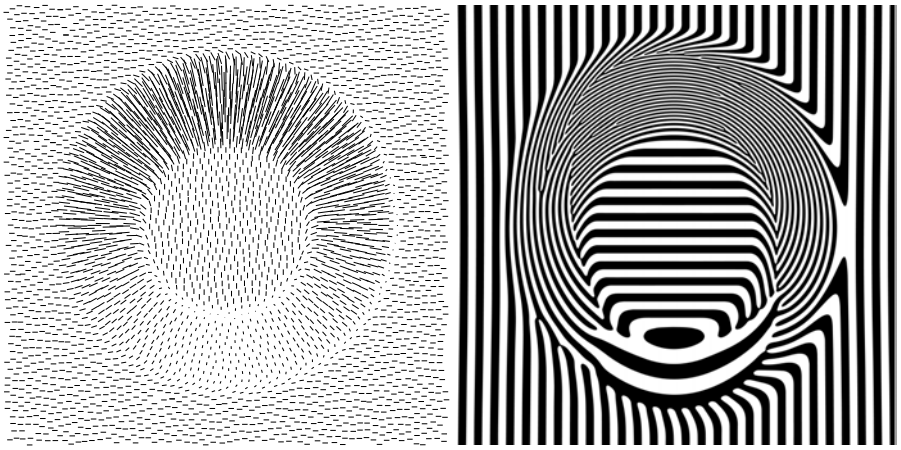

Figure 4. An arbitrary vector field *(left)* drives stripe orientation and spacing *(right)*; there are no special conditions on the input.

features rather than geometric phenomena. Input to these methods can be generated using a variety of techniques **[27, 54, 22, 62, 5, 16, 33, 20]**.

Field-aligned parameterization is also a key component of global remeshing algorithms, where there are two popular representations: either angle-valued coordinate functions with discrete *period jumps* **[5]**, or vector-valued functions whose angular components provide the final coordinates **[59, 80]**. Both choices traditionally lead to difficult nonconvex problems: the former due to integer variables; the latter due to a quartic unit-norm constraint at each point. The unit constraint implicit in both representations also leads to energies that do not converge under refinement and are therefore unstable with respect to tessellation **[33**, Fig. 4**]**. Nonorientable features are handled using either *matchings* **[59, 67, 5]**, or a *branched covering* **[29, 30]**, though in retrospect these two approaches are essentially equivalent: matchings determine a covering space, and optimization on a covering is no more costly than on the original domain, as long as one respects *conjugate symmetry* (Sec. A). The advantage of the covering space perspective is that it allows for a meaningful formulation in the smooth setting (Sec. 2.3), helps to simplify implementation (Sec. 4.1), and clarifies interpolation—especially near branch points (Sec. 4.3). Finally, several quadrangulation methods automatically insert singularities, but cannot be used for stripe patterns since they do not align to a given field **[44]** or do not generalize to arbitrary n-direction fields **[50, 51]**; moreover, these methods do not guarantee anything about global optimality.

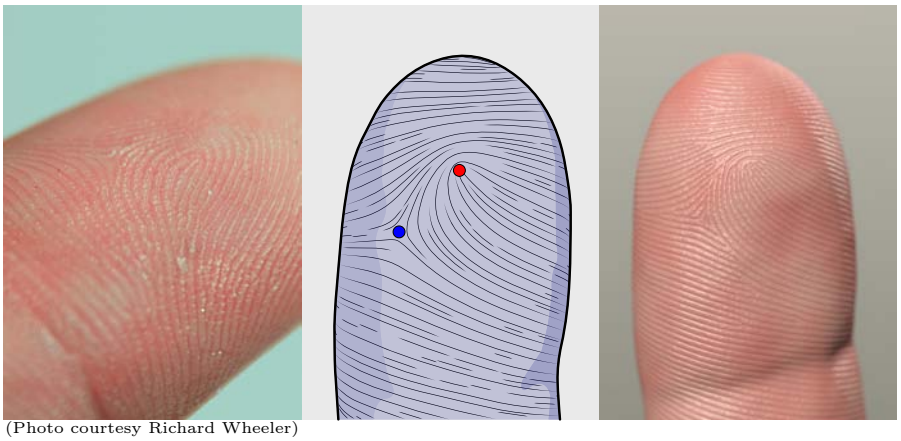

(Photo courtesy Richard Wheeler)

Figure 5. We can mimic real fingerprints *(left)* by using a field with prescribed singularities *(center)* to guide our stripe pattern *(right)*.

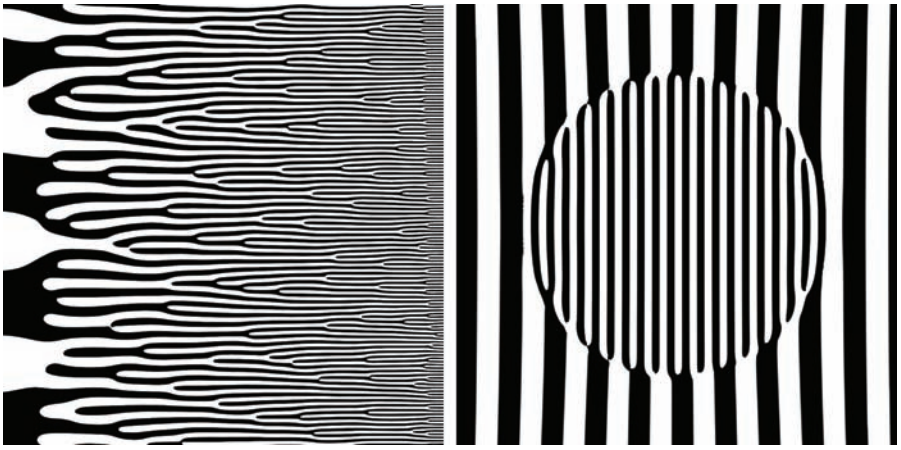

Figure 6. Our method handles line spacing that varies spatially *(left)*, even if the target spacing function is discontinuous *(right)*.

Our method adopts the same energy as Ray *et al.* [**59**], but differs in several important ways: (i) dropping quartic (unit-norm) constraints improves performance by about 10x (Sec. 5.1); (ii) permitting new singularities during parameterization allows us to directly handle nonintegrable input (Sec. 2.2); (iii) a new sub-triangle interpolation scheme (Sec. 4.3) means that pattern frequency is not limited by mesh resolution. We also provide an interpretation of the energy in the smooth setting (Sec. 2) and show new applications to texture synthesis and design (Sec. 5.3).

2 Smooth Formulation

Given a surface M, we seek a scalar function α that varies along a given unit vector field X at a given rate ν. There are two principal challenges having to do with (i) the *integrability* of the vector field $Z := \nu X$ and (ii) the representation of α itself. Roughly speaking, Z is integrable if it is locally expressible as the gradient of a scalar potential, but if we simply adopt a *real*-valued representation then we may not be able to find a potential that is globally continuous. Later, we also address *nonorientable* patterns by formulating the same problem on a *double cover* of M

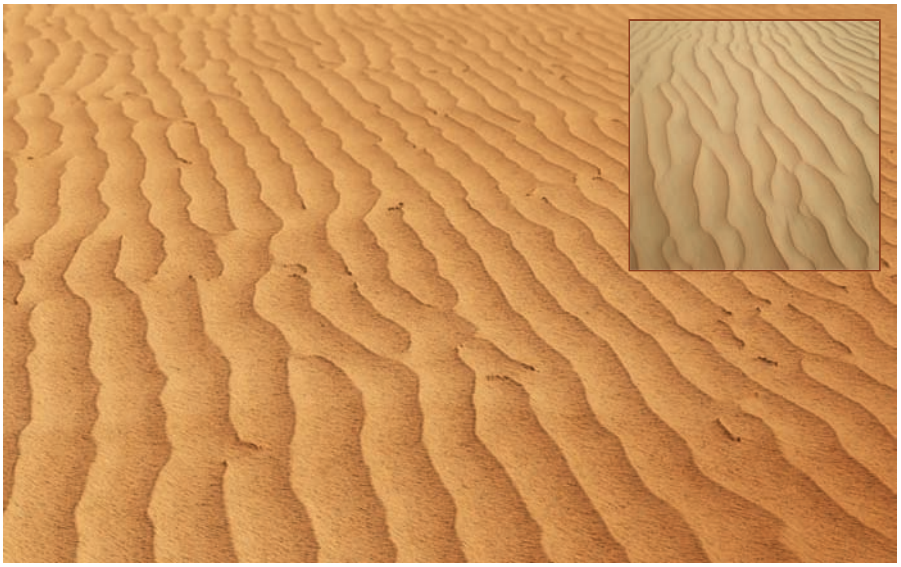

(Inset photo courtesy Lauren Edgar)

Figure 7. Randomly perturbing a constant vector field yields a pattern reminiscent of real *aeolian wind ripples (inset)*, here rendered as a displacement map over a flat plane.

(Sec. 2.3). Throughout we use $\langle \cdot, \cdot \rangle$ and $|\cdot|$ to denote the real inner product and norm on vectors, $||\cdot||$ for the L^2 norm on functions, dA for the usual area measure, and $\imath$ for the imaginary unit. We also make occasional use of (discrete) differential forms—see Crane *et al.* [**15**, Chapter 3] for a brief introduction.

2.1 Periodic Functions

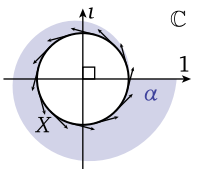

A key challenge in computing globally continuous parameterizations is choosing a suitable representation for the coordinate functions. As a motivating example, consider a smooth unit tangent field X on the circle $\alpha \mapsto (\cos\alpha, \sin\alpha)$. Locally we can think of X as the derivative of the angle α, but globally it cannot be expressed as the derivative of a smooth function since any such function would have to increase indefinitely as we walk multiple times around the circle. We therefore adopt an alternative representation: rather than a real function α, we work with a complex function ψ whose argument $\arg\psi := \text{atan2}(\text{Im}\,\psi, \text{Re}\,\psi) \in [-\pi, \pi)$ serves as a proxy for α. Although $\arg\psi$ is not globally continuous when viewed as a real-valued function, it still has a well-defined rate of change when viewed as a map to the unit circle, namely $d\arg\psi := \langle d\psi, i\psi\rangle / |\psi|^2$. Hence, we obtain a representation of angle that is globally continuous and differentiable away from zeros (in the case of the circle, we have just $\psi = \cos\alpha + \imath\sin\alpha$). For surfaces, nothing changes except that the domain of ψ is now a surface rather than a curve. Importantly, we do not need to cut the surface into a topological disk since coordinates are globally continuous *by construction.*

2.2 Integrability

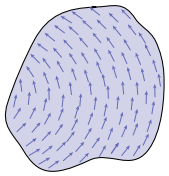

Vector fields with nonzero curl are *nonintegrable*, *i.e.*, even on small, simply-connected patches they cannot be expressed as the gradient of any scalar function. A common technique is to reduce or eliminate curl prior to parameterization via *curl correction* [**59**] or *Helmholtz-Hodge decomposition* [**29**]. We instead work directly with nonintegrable input by allowing new singularities during parameterization—empirically, these singularities tend to appear in regions where there is significant curl, or where stripe spacing changes quickly (see Figs. 2, 3, 4, and 6). More explicitly, consider the complex function ψ (Sec. 2.1)—since we care only about its argument $\alpha := \arg\psi$ (and not its magnitude), we initially consider optimization over functions $\varphi := e^{\imath\alpha}$ with unit norm at each point. Conceptually, we want the angle α to change at rate ν along the direction X, and remain constant along the orthogonal direction. In other words, we want $d\alpha(Y) = \nu\langle X, Y\rangle$ for all vector fields Y, or equivalently, $d\alpha = \omega$ for $\omega := \langle \nu X, \cdot\rangle = \langle Z, \cdot\rangle$. Differentiating φ we get $d\varphi = \imath d\alpha\varphi$, which means we can achieve the desired angular velocity $d\alpha$ by solving

$$d\varphi = \imath\omega\varphi. \tag{2.1}$$

Unless Z is exactly integrable, this equation admits only the trivial solution $\varphi \equiv 0$. We therefore consider its L^2 residual

$$\mathcal{E}(\varphi) := \int_M |(d - \imath\omega)\varphi|^2. \tag{2.2}$$

Since the operator $\widetilde{\nabla} := d - \imath\omega$ defines a *connection* on M, $\mathcal{E}$ can be viewed as a *Dirichlet energy* on complex functions. However, this energy is not well-defined for unit functions φ that encode the type of features we observe in natural stripe patterns (*e.g.*, Fig. 1, left), since the angular velocity $d\alpha$ approaches infinity as we approach an edge dislocation. Following the approach advocated in Knöppel *et al.* [**33**, Sec. 3], we therefore define the energy of a *unit* function φ as the minimum Dirichlet energy among all possible rescalings by a real-valued function $a \geq 0$:

$$\hat{\mathcal{E}}(\varphi) := \min_{a\geq 0, ||a||=1} \int_M |\widetilde{\nabla}(a\varphi)|^2 \, dA.$$

The constraint $||a|| = 1$ prevents the solution $a \equiv 0$. Minimizing this energy over unit functions φ is then equivalent to minimizing Dirichlet energy over *all* complex functions ψ:

$$\min_{|\varphi|=1} \hat{\mathcal{E}} \quad = \quad \min_{|\varphi|=1, a\geq 0, ||a||=1} \mathcal{E}(a\varphi) \quad = \quad \min_{||\psi||=1} \mathcal{E}(\psi). \tag{2.3}$$

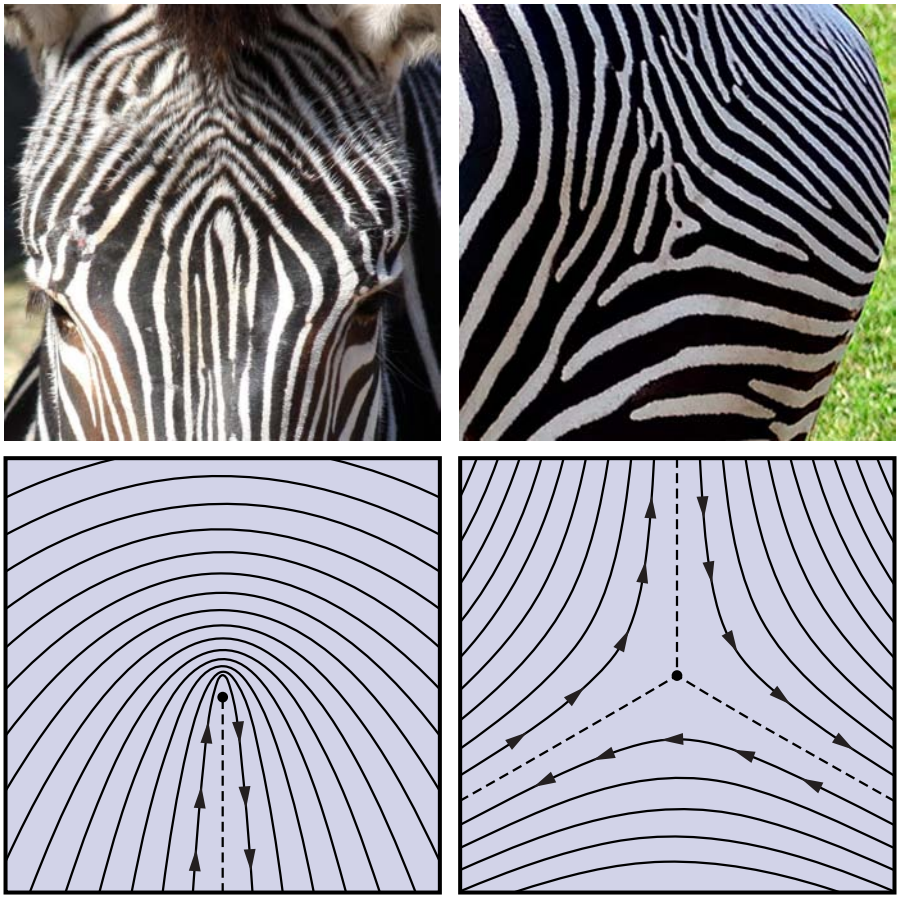

(Photos courtesy Trisha Shears and André Karwath)

Figure 8. Naturally-occurring stripe patterns *(top)* often include *nonorientable singularities*, *i.e.*, features that do not follow any continuously varying vector field *(bottom)*.

In the case of parameterization, it is precisely the ability to scale ψ down to zero that allows us to form new singularities, thereby improving alignment in nonintegrable regions. Just as with standard Dirichlet energy, minimizers of this energy among functions ψ with unit L^2-norm are eigenfunctions associated with the smallest eigenvalue of the corresponding Laplace operator **[48]**—discretization of this operator is discussed in Sec. 3.

2.3 Nonorientable Patterns

The formulation above is suitable in the orientable case, but in general we want to produce stripe patterns with *nonorientable* features, which often appear in nature (Fig 8). Another common geometric example is a *principal direction field*—see for instance Fig. 17, left. In general, such features can be encoded as a *line field*, *i.e.*, a choice of unoriented line in each tangent plane. A useful way to work with such fields, pioneered in geometry processing by Kälberer *et al.* **[29, 30]**, is as a vector field on a *branched double cover* of the original surface. Consider for instance Fig. 9 (top left)—here it is impossible to give the innermost stripe an orientation that is consistent across the dashed line. Yet by gluing together two oppositely-oriented copies of the surface, we obtain a consistently-oriented vector field $\widetilde{Z}$ on a new surface $\widetilde{M}$ called the *double cover*. Away from singularities, a double cover looks like two identical copies of the original surface M (Fig. 10). In practice we never explicitly construct the cover, but will use this mental image to turn algorithms which are naturally expressed on $\widetilde{M}$ into algorithms that can be implemented using standard data structures on the original surface M.

The key idea behind this construction is that features in a given line field determine the topology of a corresponding covering space. In particular, suppose that lines vary continuously except at a

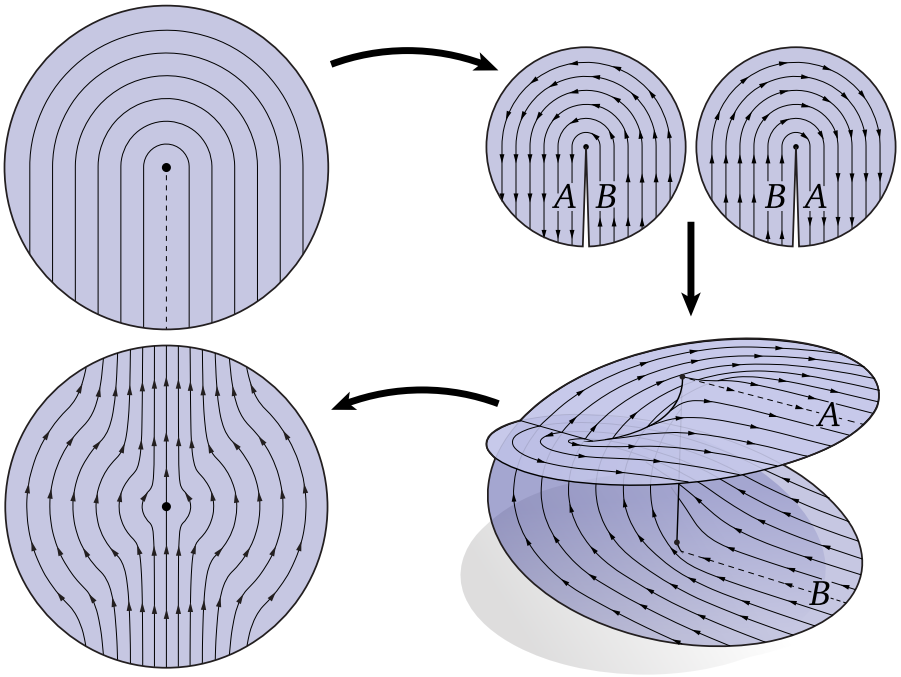

Figure 9. To deal with a nonorientable singularity *(top left)*, one can make two oppositely oriented copies *(top right)*, and cut and glue along opposite edges to obtain a branched double cover *(bottom right)*. Untwisting this figure into the plane reveals that the new vector field has a globally consistent orientation *(bottom left)*.

collection of isolated *branch points* $\{p_1, \ldots, p_k\}$ around which the field is nonorientable (Fig. 8). At each such point the "trick" demonstrated in Fig. 9 can be used to produce an orientable vector field $\widetilde{Z}$ on the double cover $\widetilde{M}$ of $\widehat{M} := M \setminus \{p_1, \ldots, p_k\}$. Away from these points the double cover has two points $x_1, x_2 \in \widetilde{M}$ "above" each point $x \in \widehat{M}$ which are mapped back "down" to the base by the *projection map* $\pi(x_1) = \pi(x_2) = x$; the *sheet interchange* function $\tau(x_1) = x_2$ swaps these two points (Fig. 10).

Conjugate Symmetry. We now have an oriented vector field $\widetilde{Z}$ on a surface $\widetilde{M}$ which—at least conceptually—could be parameterized precisely as described in Sec. 2.3, since we can simply forget that this data comes from a double cover. However, we ultimately need texture coordinates on the base surface M. Suppose for a moment that the minimizer ψ on $\widetilde{M}$ is *conjugate symmetric*,

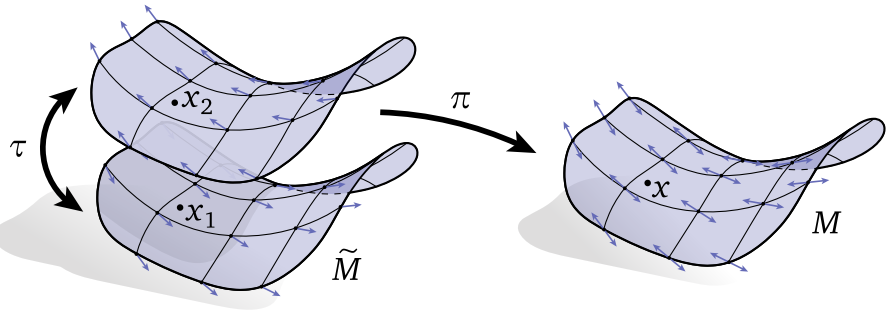

Figure 10. Away from singularities, an unoriented *line field* on a patch of M can be expressed as a pair of oppositely and consistently-oriented vector fields on two identical copies of this patch in $\widetilde{M}$. The map τ switches between the two copies, and the map π takes us back down to the original surface.

meaning that it gets conjugated under sheet interchange:

$$\psi \circ \tau = \overline{\psi}. \tag{2.4}$$

The coordinate function $\alpha = \arg \psi$ is then *antisymmetric* with respect to sheet interchange, *i.e.*, $\alpha \circ \tau = -\alpha$, and plugging it into any *even* 2π-periodic function (like cosine) hence yields the same color value whether we use $\arg(\psi_{x_1})$ or $\arg(\psi_{x_2})$. An arbitrary minimizer ψ will not in general satisfy Eq. 2.4, but as shown in App. A, *one of the energy minimizers will always be conjugate-symmetric.* We can therefore always obtain a stripe pattern that is globally continuous away from isolated branch points. Moreover, the algorithm can be implemented on the base mesh without doubling the memory or computation cost, since we need only store a single value to represent both ψ and $\overline{\psi}$.

3 Discretization

As in the smooth setting we begin with oriented vector fields, then use a double cover to reduce the nonorientable case to the orientable one. The ultimate goal is to compute a value of $\alpha \in \mathbb{R}$ at each triangle corner—plugging interpolated values of α into any even periodic function will then produce a pattern of oscillating stripes orthogonal to X. The input to our algorithm is a manifold simplicial 2-complex K and positive edge lengths ℓ_{ij} satisfying the triangle inequality on each face, along with a unit tangent vector X_i and target line frequency $\nu_i \in \mathbb{R}^+$ which together define a (not necessarily unit) vector $Z_i := \nu_i X_i$ at each vertex $i \in V$. Overall, the algorithm consists of four basic steps:

(1) Project Z onto each edge to get angular displacements ω.

(2) Build a Laplace-like matrix A with entries determined by ω.

(3) Find an eigenvector ψ of A with smallest eigenvalue.

(4) In each face, compute texture coordinates α from ψ and ω.

Notation. For any simplicial complex K, we use V, E, F to denote its vertices, edges, and faces, respectively. Individual simplices are expressed as a list of incident vertices, *e.g.*, $ijk \in F$ denotes the face with vertices $i, j, k \in V$. The orientation of a simplex is determined by the order of its indices, and an edge $ij \in E$ is *canonically oriented* if $i < j$. An expression of the form $x_i = \sum_{ij \in E} y_{ij}$ indicates that a value x_i associated with vertex i is obtained by summing the quantity y_{ij} over all edges incident on i. Likewise, $x_i = \sum_{ijk \in F} y_{ijk}$ denotes a sum over all incident faces.

3.1 Discrete Tangent Spaces

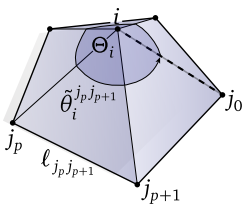

Rather than thinking of tangent vectors X_i as elements of $\mathbb{R}^3$, we view them as angles $\phi_i \in [-\pi, \pi)$ relative to a reference edge $ij_0 \in E$ chosen at each vertex $i \in V$, effectively adopting polar coordinates. This intrinsic setup simplifies the rest of our algorithm, and does not require us to define normals at vertices. Moreover, let $\hat{\theta}_i^{jk}$ denote the angle at corner i of triangle ijk, and let $\Theta_i := \sum_{ijk \in F} \hat{\theta}_i^{jk}$ be the sum of all angles incident on vertex i. For each oriented edge ij_a incident on i, we store the angle

$$\theta_{ij_a} := \tfrac{2\pi}{\Theta_i} \sum_{p=0}^{a-1} \hat{\theta}_i^{j_p j_{p+1}},$$

where the j_p are indexed counter-clockwise around vertex i, and $p+1$ is taken modulo the degree of the vertex. These rescaled angles are effectively the polar coordinates of the edges, starting at $\theta_{ij_0} = 0$ (see Knöppel *et al.* [**33**, Sec. 6] for further interpretation). Hence, given any unit vector X_i encoded by an angle ϕ_i, the parallel vector at a neighboring vertex j can be computed via

$$\phi_i \mapsto \phi_i \underbrace{-\theta_{ij} + (\theta_{ji} + \pi)}_{=:\rho_{ij}}, \tag{3.1}$$

i.e., by using the shared edge ij as a common frame of reference (see inset below). For convenience, we store the values ρ_{ij}, noting that $\rho_{ji} = -\rho_{ij}$.

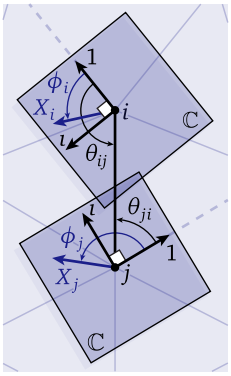

Complex Representation. In many cases, we will instead encode tangent vectors as complex numbers to avoid the difficulty of comparing angles modulo 2π. An expression like Eq. 3.1 can then be written as $X_i \mapsto e^{\imath\rho_{ij}} X_i$, which in complex arithmetic expresses a counter-clockwise rotation of X_i by the angle ρ_{ij}. It is important to note, however, that these complex values still only carry meaning relative to the reference edge chosen at each vertex, which we identify with the real axis $1 + 0\imath \in \mathbb{C}$.

3.2 Discrete Energy

As discussed in Sec. 2.2, we ideally want an angle-valued function α such that $d\alpha = \omega$, *i.e.*, such that the angular velocity $d\alpha$ matches our input field. To discretize this relationship, we linearly interpolate Z and integrate along each edge $ij \in E$, yielding

$$\alpha_j - \alpha_i = \int_{ij} d\alpha = \int_{ij} \omega = \tfrac{1}{2}(\langle e_{ij}, Z_i\rangle + \langle e_{ij}, Z_j\rangle) =: \omega_{ij}. \tag{3.2}$$

The resulting quantity $\omega_{ij} = -\omega_{ji} \in \mathbb{R}$ determines a *discrete 1-form* [**15**, Sec. 3.6] that describes the target angular displacement as we move from i to j; we will use

$$P_{ij}(X) := e^{\imath\omega_{ij}} X$$

to denote the transport of a vector X from i to j via ω. If extrinsic data is available (*i.e.*, vectors and edges in $\mathbb{R}^3$) then this quantity can be computed as above. If only intrinsic data is available (*i.e.*, lengths and angles), one can just as easily write ω as

$$\omega_{ij} = \tfrac{1}{2}\ell_{ij}(\nu_i \cos(\phi_i - \theta_{ij}) + \nu_j \cos(\phi_j - \theta_{ji})). \tag{3.3}$$

Recalling that $\arg\psi$ is used as a proxy for α (Sec. 2.1), we therefore ask that the rotated value of ψ_i agrees exactly with ψ_j, providing a discrete version of Equation (2.1):

$$e^{\imath\omega_{ij}}\psi_i = \psi_j \quad \forall ij \in E.$$

As in the smooth setting we cannot expect to satisfy this relationship exactly, since not every discrete vector field is integrable. Hence, we instead minimize the L^2 residual

$$\mathcal{E}(\psi) := \sum_{ij\in E} w_{ij}|\psi_j - e^{\imath\omega_{ij}}\psi_i|^2, \tag{3.4}$$

where $w_{ij} := (\cot\beta_{ij} + \cot\beta_{ji})/2$ and β_{ij}, β_{ji} are the two angles opposite edge ij; these *cotangent weights* account for the shape of mesh elements [**45**, Sec. 3.2]. For boundary edges we set the unknown cotan to zero, corresponding to zero-Neumann boundary conditions. Likewise, we effectively omit branch triangles (Sec. 3.3) by simply setting their weights to zero.

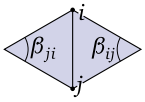

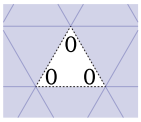

3.3 Discrete Double Cover

To handle nonorientable input, we now define the double cover $\widetilde{K} = (\widetilde{V}, \widetilde{E}, \widetilde{F})$ induced by vectors Z_i at the vertices of our original mesh K, as well as an *orientable* vector field $\widetilde{Z}$ on $\widetilde{K}$ (Fig. 11). Note that at no point do we actually *build* $\widetilde{K}$—this description merely serves to clarify implementation.

Vertices: For each vertex $i \in V$, $\widetilde{V}$ contains a pair of vertices i_1, i_2, which we associate with vectors Z_i and $-Z_i$, respectively.

Edges: For each edge $ij \in E$, let the quantity

$$s_{ij} := \begin{cases} +1 & \text{if } \phi_j - (\phi_i + \rho_{ij}) \in [-\pi/2, \pi/2), \\ -1 & \text{otherwise} \end{cases}$$

encode the relative orientation of Z_i and Z_j. If $s_{ij} = +1$ then $\widetilde{E}$ contains edges i_1j_1 and i_2j_2; otherwise, it contains i_1j_2 and i_2j_1.

Faces: For each face $ijk \in F$, let $s_{ijk} := s_{ij}s_{jk}s_{ki}$. If $s_{ijk} = +1$ then ijk is *regular*, and the six corresponding edges in $\widetilde{E}$ form two distinct triangles which we include in $\widetilde{F}$. Otherwise, ijk is a *branch triangle* and these edges form a hexagon which we omit from $\widetilde{F}$. We will use $B \subset F$ to denote the set of branch triangles.

We also define a 1-form $\widetilde{\omega}$ on $\widetilde{K}$ corresponding to $\widetilde{Z}$, given by

$$\begin{aligned}\widetilde{\omega}_{i_1j_2} = -\widetilde{\omega}_{i_2j_1} &= \tfrac{1}{2}(\langle e_{ij}, Z_i\rangle - \langle e_{ij}, Z_j\rangle) \quad \text{if } s_{ij} = -1, \text{ and}\\ \widetilde{\omega}_{i_1j_1} = -\widetilde{\omega}_{i_2j_2} &= \tfrac{1}{2}(\langle e_{ij}, Z_i\rangle + \langle e_{ij}, Z_j\rangle) \quad \text{if } s_{ij} = +1,\end{aligned}$$

along with a corresponding parallel transport map $\widetilde{P}_{pq}(Z) := e^{\imath\widetilde{\omega}_{pq}}Z$. Notice that $\widetilde{\omega} \circ \tau = -\widetilde{\omega}$, *i.e.*, switching sheets flips the sign of $\widetilde{\omega}$. For convenience, we define

$$\widehat{\omega}_{ij} = \begin{cases}\widetilde{\omega}_{i_1j_2}, & \text{if } s_{ij} = -1,\\ \widetilde{\omega}_{i_1j_1}, & \text{if } s_{ij} = +1,\end{cases}$$

and use $\hat{P}_{ij}(Z) := e^{\imath\widehat{\omega}_{ij}}Z$ to denote parallel transport by $\widehat{\omega}$.

3.4 Conjugate Symmetric Energy

We now formulate an expression for the discrete energy (Eq. 3.4) on $\widetilde{K}$, restricted to conjugate-symmetric functions $\widetilde{\psi}$ (Sec. 2.3). To do so, we store a single value $\psi_i \in \mathbb{C}$ at each vertex $i \in V$ of our original mesh, and define (but do not explicitly store)

$$\widetilde{\psi}_{i_1} := \psi_i, \qquad \widetilde{\psi}_{i_2} := \overline{\psi}_i.$$

For a given edge $ij \in E$, one then observes that the two corresponding edges on the double cover make identical contributions to the energy. For example, when $s_{ij} = +1$, we have terms

$$w_{ij}|\psi_j - e^{\imath\widetilde{\omega}_{i_1j_1}}\psi_i|^2 \qquad \text{and} \qquad w_{ij}|\overline{\psi}_j - e^{\imath\widetilde{\omega}_{i_2j_2}}\overline{\psi}_i|^2,$$

which are identical in magnitude since $e^{\imath\widetilde{\omega}_{i_2j_2}} = e^{-\imath\widetilde{\omega}_{i_1j_1}} = \overline{e^{\imath\widetilde{\omega}_{i_1j_1}}}$, and conjugation does not affect the norm (likewise for $s_{ij} = -1$). The energy associated with a given edge $ij \in E$ is therefore

$$\tilde{\mathcal{E}}_{ij} := \begin{cases}2w_{ij}|\overline{\psi}_j - e^{\imath\widehat{\omega}_{ij}}\psi_i|^2, & s_{ij} = -1,\\ 2w_{ij}|\psi_j - e^{\imath\widehat{\omega}_{ij}}\psi_i|^2, & s_{ij} = +1,\end{cases} \tag{3.5}$$

yielding a total energy $\tilde{\mathcal{E}} := \sum_{ij\in E}\tilde{\mathcal{E}}_{ij}$. Notice that this energy is defined entirely in terms of a single value ψ per vertex and a single value $\widehat{\omega}$ per edge of our original mesh K: we do not need to double the degrees of freedom.

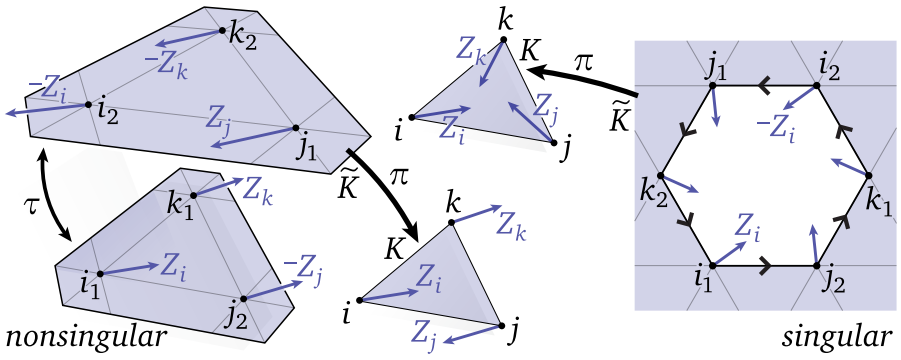

Figure 11. The double cover of a triangle is determined by the relative orientation of vectors Z along each edge. *Left:* If the vector field can be locally oriented, we simply make two copies of the original triangle. *Right:* Otherwise we have a *branch triangle,* and use the six corresponding edges to form a hexagonal boundary cycle.

4 Implementation

We now describe essential aspects of implementation; detailed pseudocode can be found in App. C.

4.1 Matrix Representation

Although $\tilde{\mathcal{E}}$ is expressed using complex variables, it must be encoded as a real matrix in order to represent conjugation of the stored values of ψ, which is not a complex-linear operation. In particular, for complex numbers $u = a + b\imath$, $v = c + d\imath$ and $z = x + y\imath$ the following relations are easily verified:

$$\langle zu, v\rangle = (a\ b)\big[z\big](c\ d)^{\mathsf{T}}, \quad \langle zu, \overline{v}\rangle = (a\ b)\overline{\big[z\big]}(c\ d)^{\mathsf{T}},$$

where

$$\big[z\big] := \begin{pmatrix} x & y \\ -y & x \end{pmatrix} \qquad \text{and} \qquad \overline{\big[z\big]} := \begin{pmatrix} x & -y \\ -y & -x \end{pmatrix}.$$

Expanding Eq. 3.5 and applying these expressions then allow us to express $\tilde{\mathcal{E}}$ as a matrix $\mathsf{A} \in \mathbb{R}^{2|V|\times 2|V|}$ with diagonal blocks

$$\mathsf{A}_{ii} = \sum_{ij\in E} \big[w_{ij}\big]$$

for each vertex $i \in V$, and off-diagonal blocks

$$\mathsf{A}_{ij} = \begin{cases} -w_{ij}\overline{\big[e^{\widehat{\imath\omega_{ij}}}\big]}, & s_{ij} = -1 \\ -w_{ij}\big[e^{\widehat{\imath\omega_{ij}}}\big], & s_{ij} = +1 \end{cases}, \qquad \mathsf{A}_{ji} = \mathsf{A}_{ij}^{\mathsf{T}}$$

for each canonically oriented edge $ij \in E$. The resulting matrix is symmetric positive-definite with the same structure as the usual cotan-Laplace matrix. Finally, letting $\mathcal{A}_{ijk}$ denote the area of triangle ijk, we encode the usual L^2 norm on functions via the block diagonal lumped mass matrix $\mathsf{B} \in \mathbb{R}^{2|V|\times 2|V|}$ with entries

$$\mathsf{B}_{ii} = \tfrac{1}{3} \sum_{ijk\in F} \mathcal{A}_{ijk},$$

i.e., one-third the total area of triangles incident on vertex i.

4.2 Global Optimization

We now seek values $\psi_i = a_i + b_i\imath$ that minimize $\tilde{\mathcal{E}}$, which we encode as a vector $\mathsf{x} \in \mathbb{R}^{2|V|}$ of interleaved values $a_i, b_i \in \mathbb{R}$. To avoid the trivial solution $\psi \equiv 0$, we ask that ψ have unit L^2 norm—using the matrices above, this problem can be stated as

$$\min_{\mathsf{x}\in\mathbb{R}^{2|V|}} \mathsf{x}^{\mathsf{T}}\mathsf{A}\mathsf{x} \quad \text{s.t.} \quad \mathsf{x}^{\mathsf{T}}\mathsf{B}\mathsf{x} = 1. \tag{4.1}$$

Differentiating the Lagrangian of this problem yields the generalized eigenvalue problem $\mathsf{A}\mathsf{x} = \lambda\mathsf{B}\mathsf{x}$ for an eigenvector x associated with the smallest eigenvalue λ, which we solve by factoring A and applying the inverse power method (Alg. 9). Hence, the total cost of our algorithm is dominated by a single matrix factorization.

4.3 Texture Coordinates

After computing ψ, we still need to extract our final texture coordinates α. However, simply using the values $\arg(\psi_i)$ is less than ideal since (i) we will experience aliasing if the target frequency ν is greater than the mesh spacing (Fig. 12) and (ii) linear interpolation will produce discontinuities along the boundaries of triangles containing zeros or branch points (Fig. 13, 14). We therefore adjust α as described below. On $\widetilde{K}$ these coordinates are antisymmetric with respect to sheet interchange and encode a globally continuous map to the unit circle, as shown in App. B. Hence, for any given base triangle we can use coordinates from either of the two covering triangles to obtain a globally continuous stripe pattern. Note that for sufficiently fine meshes one can often ignore singular triangles and simply use linear interpolation everywhere, as done in Figs. 1, 5, 19, 20, 21, and 7.

The Spinning Form. Along each edge $pq \in \widetilde{E}$ we had a target angular displacement $\widetilde{\omega}_{pq}$, which in general we cannot achieve due to nonintegrability. The *spinning form* $\widetilde{\sigma}$ is the angular displacement closest to $\widetilde{\omega}$ that agrees with the obtained minimizer of $\tilde{\mathcal{E}}$, *i.e.*, for which $\arg(e^{\imath\widetilde{\sigma}_{pq}}\widetilde{\psi}_p) = \arg(\widetilde{\psi}_q)$. These values will be used to adjust our texture coordinates. In particular, let

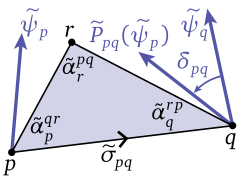

$$\delta_{pq} := \arg(\widetilde{P}_{pq}(\widetilde{\psi}_p)/\widetilde{\psi}_q)$$

be the smallest angular distance from the obtained vector $\widetilde{\psi}_q$ at vertex q to the desired vector $\widetilde{P}_{pq}(\widetilde{\psi}_p)$. Then

$$\widetilde{\sigma}_{pq} := \widetilde{\omega}_{pq} - \delta_{pq}.$$

Note that, like $\widetilde{\omega}$, $\widetilde{\sigma}$ is a discrete 1-form on $\widetilde{K}$, *i.e.*, $\widetilde{\sigma}_{pq} = -\widetilde{\sigma}_{qp}$.

Frequency Adjustment. We first adjust the texture coordinates locally in each triangle to properly account for the target frequency ν. More explicitly, let $\tilde{\alpha}_p^{qr}$ denote the coordinate at corner p of a triangle $pqr \in \widetilde{F}$. We let $\tilde{\alpha}_p^{qr} := \arg(\widetilde{\psi}_p)$ at the first corner, and use the spinning form to define coordinates at the other two corners: $\tilde{\alpha}_q^{rp} := \tilde{\alpha}_p^{qr} + \widetilde{\sigma}_{pq}$ and $\tilde{\alpha}_r^{pq} := \tilde{\alpha}_p^{qr} + \widetilde{\sigma}_{pr}$. To render a given base triangle $ijk \in F$ we need only extract the coordinates from one of its two covering triangles, which we can do using the values ψ and $\widehat{\omega}$ stored on our original mesh K. In particular, if we let

$$\hat{\delta}_{ij} := \begin{cases} \arg(\hat{P}_{ij}(\psi_i)/\overline{\psi}_j), & s_{ij} = -1, \\ \arg(\hat{P}_{ij}(\psi_i)/\psi_j), & s_{ij} = +1, \end{cases}$$

and $\hat{\sigma} := \hat{\omega} + \hat{\delta}$, then our local coordinates are just

$$\begin{aligned} \alpha_i^{jk} &:= \arg(\psi_i), \\ \alpha_j^{ki} &:= \arg(\psi_i) + \hat{\sigma}_{ij}, \\ \alpha_k^{ij} &:= \arg(\psi_i) + \hat{\sigma}_{ik}. \end{aligned}$$

In nonsingular regions, α can then be interpolated linearly over each triangle. In general, however, we must account for singularities due to either (i) zeros of the function ψ or (ii) branch points arising from the field Z; each of these cases is treated below.

Zeros. As we walk around the boundary of any regular triangle $ijk \in F \setminus B$ covered by a triangle $i_1qr \in \widetilde{F}$, the spinning form describes a net change in angle of

$$\widetilde{\sigma}_{i_1q} + \widetilde{\sigma}_{qr} + \widetilde{\sigma}_{ri_1} =: 2\pi n_{ijk}$$

for some integer $n_{ijk} \in \mathbb{Z}$ which we refer to as the *index* of ψ. If $n_{ijk} \neq 0$, this means that the angle α on the base domain must "jump" somewhere and hence cannot be interpolated by any continuous function (including a linear one). We instead use an interpolant that mimics the

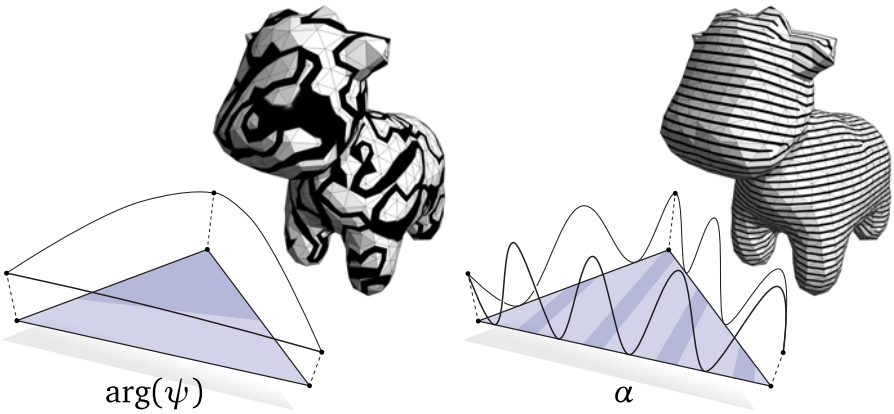

Figure 12. *Left:* for coarse meshes or high frequency stripes, naïve interpolation of angles leads to severe aliasing. *Right:* by augmenting the angles according to the target frequency, we can render stripes far above the resolution of the mesh.

Figure 13. In triangles containing zeros, piecewise linear interpolation *(left)* can be substantially improved *(middle)* by using a closed-form, nonlinear interpolant *(right)*.

behavior of the complex function $z \mapsto \arg(z)$, but remains linear along the boundary so that it agrees with interpolants in adjacent triangles. Explicitly, let

$$\mathrm{lArg}_n(t_i, t_j, t_k) := \begin{cases} \frac{\pi\, n}{3}\left(1 + \frac{t_j - t_i}{1-3t_k}\right), & t_k \le t_i \text{ and } t_k \le t_j, \\ \frac{\pi\, n}{3}\left(3 + \frac{t_k - t_j}{1-3t_i}\right), & t_i \le t_j \text{ and } t_i \le t_k, \\ \frac{\pi\, n}{3}\left(5 + \frac{t_i - t_k}{1-3t_j}\right), & t_j \le t_k \text{ and } t_j \le t_i, \end{cases} \tag{4.2}$$

where $t_i, t_j, t_k \in \mathbb{R}$ are barycentric coordinates (Fig. 13, right). One can easily check that the function lArg is piecewise linear along the boundary, and describes a continuous map to the unit circle away from an isolated singularity at the barycenter m. To get our final interpolant, we first subtract the value of $\mathrm{lArg}_{n_{ijk}}$ from our texture coordinates α_i^{jk} at all three triangle corners (Alg. 10, lines 22–24). Texture coordinates are then computed by evaluating Eq. 4.2 in a fragment shader and adding the result to the usual linearly interpolated values of α.

Branch Points. Finally, we extend our texture coordinates to any branch triangle $ijk \in B$, which on the double cover $\widetilde{K}$ corresponds to a hexagonal boundary cycle (p, q, r, s, t, u) starting at $p = i_1$ (Fig. 11, right). Conceptually, we imagine triangulating the hexagon by inserting a new vertex $\widetilde{m}$ at the middle and extending coordinates linearly over each face. In practice, we simply draw the barycentric subdivision of the base triangle ijk (Fig. 14, right) using texture coordinates

$$\begin{aligned} \beta_i &:= \arg(\psi_i) & \beta_j &:= \beta_i + \widetilde{\sigma}_{pq} \\ \beta_m &:= \beta_i + (\widetilde{\sigma}_{pq} + \widetilde{\sigma}_{qr} + \widetilde{\sigma}_{rs})/2 & \beta_k &:= \beta_j + \widetilde{\sigma}_{qr} \\ & & \beta_l &:= \beta_k + \widetilde{\sigma}_{rs} \end{aligned}$$

with linear interpolation in each sub-triangle. (See App. B for further discussion; an explicit implementation is given in Alg. 10.)

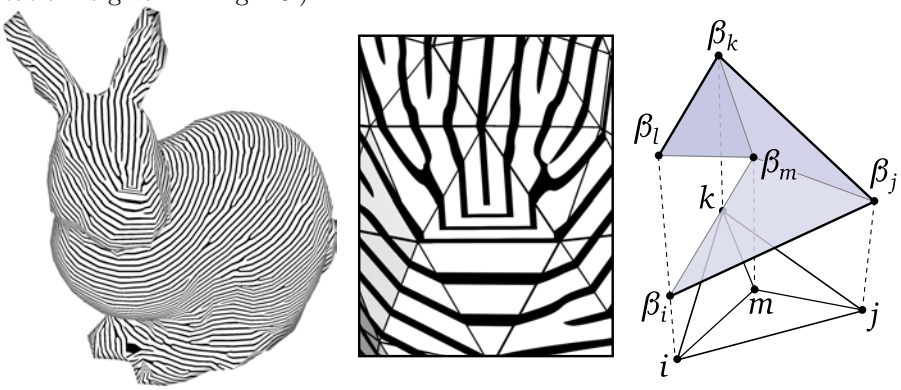

Figure 14. *Left:* stripe pattern on a coarse mesh of 708 vertices. *Center:* tip of the nose with a branch point in the middle and nearby triangles containing zeros of index 1 and 3. *Right:* piecewise linear texture coordinates used to draw branch points.

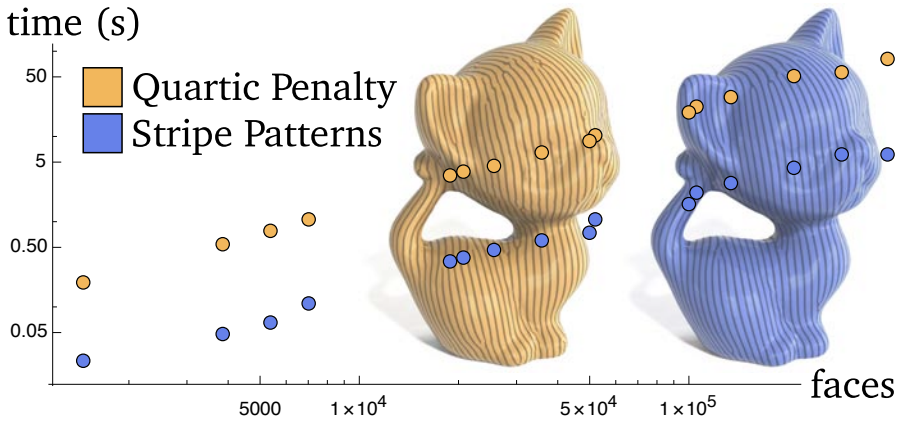

Figure 15. Our method *(blue)* produces results of comparable quality to expensive, nonlinear optimization *(yellow)*, yet is on average 10.54 times faster across a broad collection of examples.

5 Results

5.1 Performance

We implemented our method in C++ using SuiteSparse [**9**]; performance was measured on a single core of a 2.6 GHz Intel Core i7. Typical run times were less than a second; for example, we required 591ms, 14ms, 804ms, and 386ms to generate patterns in Figs. 1, 12, 18, and 21 on meshes with 28k, 1k, 44k, and 39k faces (resp.). This level of performance makes it possible to interactively edit or even animate patterns. We also compared the performance of our eigenvector scheme to the quartic penalty scheme suggested by Ray *et al.* [**59**, Sec. 2.5] using the same matrix for the quadratic term and their suggested solver parameters. Across a broad range of examples our eigenvector method performed over an order of magnitude faster (Fig. 15); moreover, the penalty scheme sometimes gets trapped in local minima, yielding unwanted distortion.

5.2 Robustness

In Fig. 16 we severely distort the input mesh, yet still obtain a well-behaved stripe pattern; this type of behavior is a hallmark of elliptic variational problems based on Laplace-like operators. The method also has no trouble handling sharp discontinuities in the input vector field Z (Fig. 6, right); likewise, outliers and noise in Z are handled gracefully (Fig. 7.)

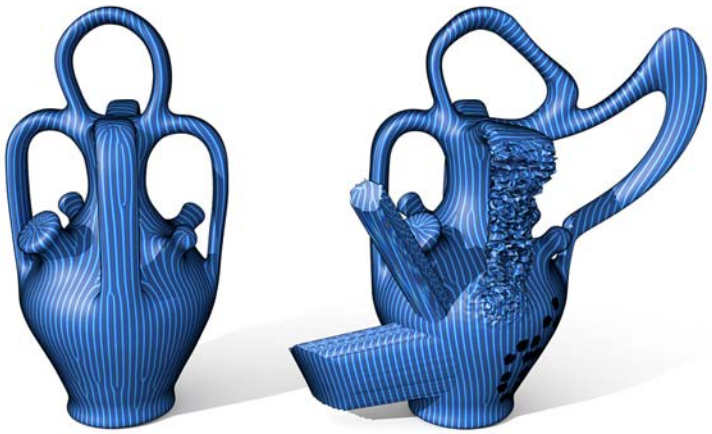

Figure 16. Due to the elliptic nature of our problem, the method is robust to noise, holes, and other severe errors in the input.

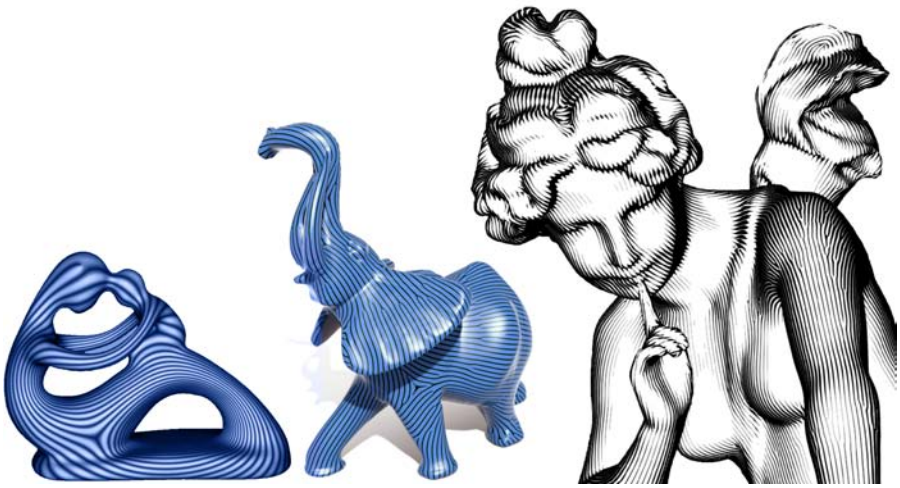

Figure 17. Minimum principal curvature directions with either nonuniform *(left)* or uniform spacing *(center)* provide a natural pattern orientation. *Right:* this kind of pattern can be used to automatically drive a real-time hatching shader, requiring no user input apart from a single global scale parameter.

5.3 Applications

Figures throughout show applications of stripe patterns to design and texture synthesis. For real-time hatching **[63]**, a precomputed stripe pattern pattern circumvents the need to trace integral curves, *etc.*(Fig. 17). In Fig. 21 we used marching triangles to generate geometric stripes. In Figs. 1 and 18, we combined two orthogonal stripe patterns to drive a 2D displacement map with bilateral symmetry across each axis. In Fig. 19 we mimic the artistic style of a ceramic mug. All of the preceding examples used optimally smooth or curvature-aligned fields as input **[33]**, but in principle any technique can be applied—Fig. 4 shows an example where Z is described by artist-driven input; Fig. 5 uses the smoothest field with prescribed singularities **[16]** to emulate fingerprints.

6 Conclusion

Remarkably, many disparate natural phenomena are well-captured by one simple energy (Eq. 2.2), making stripe patterns a versatile tool for design. More broadly, stripe patterns contribute to a growing collection of tasks in geometry processing and simulation where singular features and sparse approximations are efficiently obtained by minimizing Dirichlet energy on a particular *Hermitian line bundle*—other recent examples include the design of smooth vector fields **[33]**, extraction of smoke rings **[72]**, and conformal volume deformation **[11]**. A better understanding of this phenomenon will no doubt lead to valuable future developments.

Acknowledgments. This work was supported by the DFG Collaborative Research Center TRR 109, "Discretization in Geometry and Dynamics," an NSF Mathematical Sciences Postdoctoral Research Fellowship (Award #1304254), and Intel. Meshes in Figs. 1, 14, 18, 22 are courtesy the Stanford Computer Graphics Lab; in Figs. 5, 15, 16, 17 courtesy AIM@Shape; and in Fig. 21 courtesy Martin Newell.

A Conjugate Symmetry

Here we show that the energy $\mathcal{E}$ on $\widetilde{M}$ always has a global minimizer ψ that is *conjugate symmetric*, *i.e.*, such that

$$\psi = T\psi := \overline{\psi} \circ \tau.$$

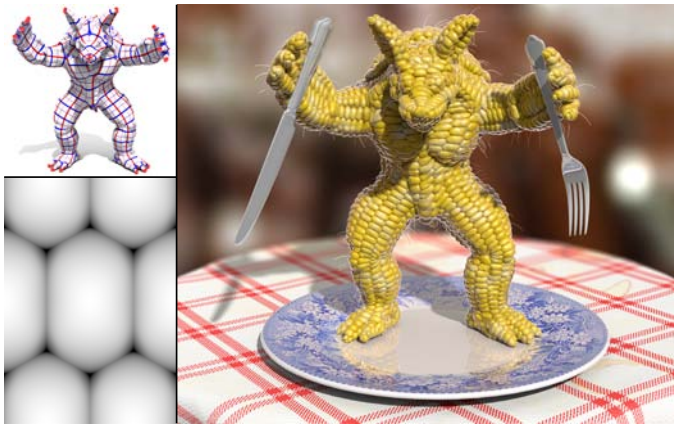

Figure 18. Two orthogonal stripe patterns (red and blue) are computed separately and combined to drive a displacement map with bilateral symmetry across both axes.

(Photo courtesy DJ Crane)

Figure 19. *Left:* Photograph of a real mug with uniform, branching stripes. *Right:* Virtual mug synthesized using our method.

(Photo courtesy Carl Clifford)

Figure 20. Stripe patterns are ubiquitous in nature—here we caricature a real angelfish *(left)* using a stripe pattern *(right)*.

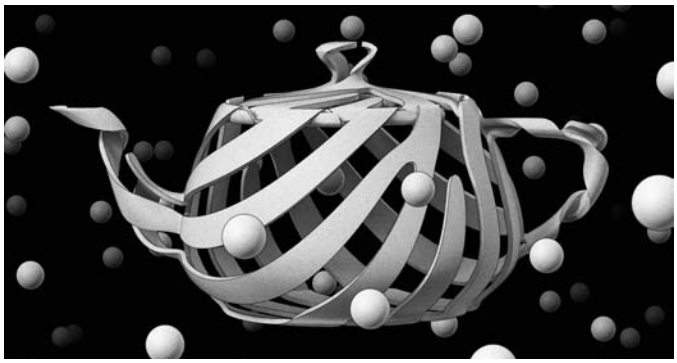

Figure 21. Stripe patterns can also be used to achieve a variety of artistic effects, here inspired by the work of M.C. Escher.

Notice that $T(\iota\psi) = -\iota T\psi$, in other words, T is *complex anti-linear*. Hence, we can decompose ψ as

$$\psi = \underbrace{\tfrac{1}{2}(\mathrm{Id} + T)\psi}_{=:\psi_1} + \tfrac{1}{2}(\mathrm{Id} - T)\psi = \psi_1 + \iota\,\underbrace{(\tfrac{1}{2}(\mathrm{Id} + T)(-\iota\psi))}_{=:\psi_2},$$

i.e., $\psi = \psi_1 + \iota\psi_2$. Since $T^2 = \mathrm{Id}$, one easily checks that both ψ_1 and ψ_2 are conjugate symmetric and $T\psi = \psi_1 - \iota\psi_2$. Moreover, since ψ_1 and $\iota\psi_2$ are orthogonal ($(\mathrm{Id} - T)(\mathrm{Id} + T) = 0$) we get

$$\|\psi\|^2 = \|\psi_1\|^2 + \|\psi_2\|^2.$$

Further, $\mathcal{E}(\psi) = \mathcal{E}(T\psi)$, *i.e.*, conjugation does not change the energy. Applying the polarization identity then yields

$$\begin{aligned}\mathcal{E}(\psi) &= \tfrac{1}{2}(\mathcal{E}(\psi) + \mathcal{E}(T\psi)) \\ &= \tfrac{1}{2}(\mathcal{E}(\psi_1 + \iota\psi_2) + \mathcal{E}(\psi_1 - \iota\psi_2)) \\ &= \mathcal{E}(\psi_1) + \mathcal{E}(\psi_2).\end{aligned}$$

Suppose now that ψ with $\|\psi\|^2 = 1$ solves (Equation (2.3)). Then for any other nonzero function ϕ we have $\mathcal{E}(\phi)/\|\phi\|^2 \geq \mathcal{E}(\psi)$, which when applied to ψ_1 (resp. ψ_2) implies

$$\mathcal{E}(\psi_1) \geq \mathcal{E}(\psi)\|\psi_1\|^2 \qquad \mathcal{E}(\psi_2) \geq \mathcal{E}(\psi)\|\psi_2\|^2.$$

But these inequalities must actually be *equalities*, otherwise summing them would yield a contradiction. Hence, there will always be some unit, conjugate symmetric function $\psi_1/\|\psi_1\|$ or $\psi_2/\|\psi_2\|$ *with the same energy as* ψ. In other words, restricting optimization to conjugate symmetric functions will still yield a global minimizer.

B Texture Coordinates

We now show that the coordinate function $\tilde{\alpha} : \widetilde{M} \to \mathbb{R}/2\pi\mathbb{Z}$ on the double cover is (i) antisymmetric with respect to sheet interchange, *i.e.*,

$$\tilde{\alpha} \circ \tau = -\tilde{\alpha}, \tag{B.1}$$

and (ii) encodes a globally continuous map to the unit circle, away from isolated singularities. Together, these properties guarantee that $\tilde{\alpha}$ can be used to draw a continuous stripe pattern. More formally, they ensure that for any even 2π-periodic function $u : \mathbb{R} \to \mathbb{R}$, there exists a function $c : K \setminus B \to \mathbb{R}$ such that away from branch triangles B,

$$u \circ \tilde{\alpha} = c \circ \pi,$$

i.e., $u(\tilde{\alpha})$ descends to some function c on the base. We will first define $\tilde{\alpha}$ on edges, and then extend it to triangle interiors.

To begin, let $\widetilde{\psi}$ be a minimizer of our discrete energy (Eq. 3.5), which is conjugate-symmetric by construction, and let $\varphi := \widetilde{\psi}/|\widetilde{\psi}|$. If $\widetilde{\sigma}$ is the corresponding *spinning form* (Sec. 4.3), one easily checks that $e^{\iota\widetilde{\sigma}_{pq}}\varphi_p = \varphi_q$ along any edge $pq \in \widetilde{E}$. Hence, for any path $\tilde{\gamma}$ from p to q along edges of $\widetilde{K}$, we have

$$e^{\iota\int_{\tilde{\gamma}}\widetilde{\sigma}}\varphi_p = \varphi_q, \tag{B.2}$$

where the integral is just the sum of $\widetilde{\sigma}$ over all the oriented edges of $\tilde{\gamma}$. From Equation (B.2) it follows that the integral of $\widetilde{\sigma}$ over any closed path ($p = q$) has a value in $2\pi\mathbb{Z}$. Thus, if we integrate $\widetilde{\sigma}$ over the whole domain (starting at an arbitrary vertex), we obtain a piecewise-linear function $\tilde{\alpha}$ on edges such that

$$e^{\iota\tilde{\alpha}_p} = \varphi_p. \tag{B.3}$$

at each vertex p. By conjugate symmetry of $\widetilde{\psi}$, and since $\widetilde{\omega}$ changes sign under τ, we see that $\widetilde{\sigma}$ also changes sign under τ. Together with Equation (B.3) this yields Equation (B.1).

We now extend $\tilde{\alpha}$ to triangle interiors. For triangles $pqr \in \widetilde{F}$ where the index $\tilde{n}_{pqr} := (\widetilde{\sigma}_{pq} + \widetilde{\sigma}_{qr} + \widetilde{\sigma}_{rp})/2\pi$ equals zero (*i.e.*, where $\widetilde{\psi}$ is regular), we can simply extend $\tilde{\alpha}$ linearly—clearly Equation (B.1) stays valid in this case. When $\tilde{n}_{pqr}$ is nonzero, we interpolate boundary values using the nonlinear function lArg described in Sec. 4.3. Here Eq. B.1 follows from the behavior of $\widetilde{\sigma}$ under τ: changing sheets flips the sign of $\tilde{n}_{pqr}$ and thus of $\mathrm{lArg}_{\tilde{n}_{pqr}}$.

Finally, consider any branch triangle $ijk \in B$, which corresponds to a hexagonal boundary cycle on the double cover $\widetilde{K}$ (Sec. 3.3). We will extend $\tilde{\alpha}$ over a triangulation connecting each boundary vertex to a "middle" vertex m. In particular, consider a path $\tilde{\gamma} = (p, q, r, s)$ on $\widetilde{K}$ between vertices $\tau(p) = s$ that projects to the path $\gamma = (i, j, k, i)$ on K. Let $\xi := \int_{\tilde{\gamma}} \widetilde{\sigma}$, and assume $\tilde{\gamma}$ is oriented such that $\tilde{\alpha}_s = \tilde{\alpha}_p + \xi$. From Equation (B.2) we have

$$e^{\imath \int_{\tilde{\gamma}} \widetilde{\sigma}} \varphi_p = \varphi_s = \overline{\varphi_p},$$

which implies that $\xi + \tilde{\alpha}_s = -\tilde{\alpha}_s + 2a\pi$ for some $a \in \mathbb{Z}$, or equivalently that the texture coordinate $\tilde{\alpha}_m := \tilde{\alpha}_s + \xi/2$ equals $a\pi$. Hence, $u(\tilde{\alpha}_p) = u(a\pi - \xi/2) = u(a\pi + \xi/2) = u(\tilde{\alpha}_s)$, which means that the piecewise linear extension of $\tilde{\alpha}$ satisfies Eq. B.1.

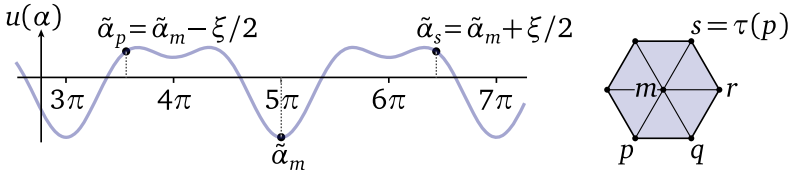

C Pseudocode

The input to the main algorithm StripePattern is a simplicial surface $K = (V, E, F)$, a collection of edge lengths $\ell \in \mathbb{R}^{|E|}$ satisfying the triangle inequality in each face, unit vectors $X \in \mathbb{C}^{|V|}$ describing the desired pattern orientation, and positive values $\nu \in \mathbb{R}^{|V|}$ giving the target line frequency. The output is a collection of coordinates $\alpha \in \mathbb{R}^{3|F|+2|B|}$ (one for each triangle corner, and two more for the midpoint m and duplicate vertex l of each branch triangle), along with indices $n_{ijk}, S_{ijk} \in \mathbb{Z}^{|F|}$ for each triangle indicating whether one should apply nonlinear interpolation or draw the barycentric subdivision (resp.).

Algorithm 4 The main stripe pattern algorithm (Sec. 2).

```
1: procedure StripePattern(K, ℓ, X, ν)                    ▷ K = (V, E, F)
2:     θ ← VertexAngles(K, ℓ)
3:     ω, s ← EdgeData(K, ℓ, θ, X, ν)
4:     A ← EnergyMatrix(K, ℓ, ω, s)
5:     B ← MassMatrix(K, ℓ)
6:     ψ ← PrincipalEigenvector(A, B)
7:     α, n, S ← TextureCoordinates(K, ψ, ω, s)
8:     return α, n, S
9: end procedure
```

Algorithm 5 Computes polar coordinates of outgoing halfedges around each vertex (Sec. 3.1).

```
1: procedure VertexAngles(K, ℓ)
2:     for each i ∈ V do
3:         Θ_i ← 0                                        ▷ cumulative angle
4:         for p ← {0, ..., Degree(K, i) − 1} do
5:             θ_{ij_p} ← Θ_i
6:             Θ_i ← Θ_i + TipAngle(K, ℓ, i, j_p, j_{p+1})   ▷ returns θ̂_i^{jk}
7:         end for
8:         for p ← {0, ..., Degree(K, i) − 1} do
9:             θ_{ij_p} ← 2πθ_{ij_p} / Θ_i
10:        end for
11:    end for
12:    return θ
13: end procedure
```

Algorithm 6 Initializes basic edge data (Sec. 3.1).

procedure EdgeData(K, ℓ, θ, X, ν)
  for each $ij \in E$ **do** ▷ canonically oriented ($i < j$)
    $\rho_{ij} \leftarrow -\theta_{ij} + \theta_{ji} + \pi$
    $s_{ij} \leftarrow \operatorname{sgn}(\langle e^{\imath \rho_{ij}} X_i, X_j \rangle)$
    $\phi_i \leftarrow \arg(X_i)$
    $\phi_j \leftarrow \arg(s_{ij} X_j)$
    $\omega_{ij} \leftarrow \frac{\ell_{ij}}{2}(\nu_i \cos(\phi_i - \theta_{ij}) + \nu_j \cos(\phi_j - \theta_{ji}))$
  end for
  return ω, s
end procedure

Algorithm 7 Builds the matrix defining the energy (Sec. 4.1).

procedure EnergyMatrix(K, ℓ, ω, s)
  $\mathsf{A} \leftarrow 0 \in \mathbb{R}^{2|V| \times 2|V|}$ ▷ start with all zeros
  for each $ij \in E$ **do** ▷ canonically oriented ($i < j$)
    $\beta_i, \beta_j \leftarrow$ OppositeAngles(K, ℓ, ij)
    $w_{ij} \leftarrow \frac{1}{2}(\cot \beta_{ij} + \cot \beta_{ji})$
    $\mathsf{A}_{ii} \leftarrow \mathsf{A}_{ii} + [w_{ij}]$
    $\mathsf{A}_{jj} \leftarrow \mathsf{A}_{jj} + [w_{ij}]$
    if $s_{ij} \geq 0$ **then**
      $\mathsf{A}_{ij} \leftarrow -w_{ij}[e^{\imath \omega_{ij}}]$
    else
      $\mathsf{A}_{ij} \leftarrow -w_{ij}\overline{[e^{\imath \omega_{ij}}]}$
    end if
    $\mathsf{A}_{ji} \leftarrow \mathsf{A}_{ij}^T$
  end for
  return A
end procedure

Algorithm 8 Builds the mass matrix associated with vertices (Sec. 4.1).

procedure MassMatrix(K, ℓ)
  $\mathsf{B} \leftarrow 0 \in \mathbb{R}^{2|V| \times 2|V|}$ ▷ start with all zeros
  for each $ijk \in F$ **do**
    $\mathcal{A}_{ijk} \leftarrow$ TriangleArea(K, ℓ, ijk)
    $\mathsf{B}_{ii} \leftarrow \mathsf{B}_{ii} + \mathcal{A}_{ijk}/3$
    $\mathsf{B}_{jj} \leftarrow \mathsf{B}_{jj} + \mathcal{A}_{ijk}/3$
    $\mathsf{B}_{kk} \leftarrow \mathsf{B}_{kk} + \mathcal{A}_{ijk}/3$
  end for
  return A
end procedure

Algorithm 9 Computes an eigenvector corresponding to the smallest eigenvalue via the inverse power method (Sec. 4.2).

procedure PrincipalEigenvector(A, B)
  $\mathsf{L} \leftarrow$ CholeskyFactor(A)
  $\mathsf{x} \leftarrow$ UniformRand(Size(A))
  for $i = 1, \ldots, N$ **do**
    $\mathsf{x} \leftarrow$ BackSolve(L, Bx)
    $\mathsf{x} \leftarrow \mathsf{x}/\sqrt{\mathsf{x}^\mathsf{T}\overline{\mathsf{Bx}}}$
  end for
  return x
end procedure

The remaining routines are standard, but are defined here for completeness:

- Degree(K, i) - returns the degree of vertex $i \in V$.
- TipAngle(K, ℓ, i, j, k) - returns the angle at vertex $i \in V$ in triangle $ijk \in F$.
- OppositeAngles(K, ℓ, ij) - returns the two angles opposite edge $ij \in E$ (Sec. 3.2).
- TriangleArea(K, ℓ, ijk) - returns area of triangle $ijk \in F$.
- CholeskyFactor(A) - returns Cholesky factor of matrix A.
- UniformRand(n) - returns a vector of n numbers in the interval $[-1, 1]$, picked uniformly at random.
- Size(A) - returns the dimension of a square matrix A.
- BackSolve(L, b) - if L is a Cholesky factor of A, solves the linear system $\mathsf{Ax} = \mathsf{b}$.

Algorithm 10 Computes final texture coordinates. (Sec. 4.3).

procedure TEXTURECOORDINATES(K, ψ, ω, s)
  for each $ijk \in F$ **do**
    $c_{ij} \leftarrow i < j \ ? \ 1 : -1$ ▷ is each edge canonical?
    $c_{jk} \leftarrow j < k \ ? \ 1 : -1$
    $c_{ki} \leftarrow k < i \ ? \ 1 : -1$
    $z_i, z_j, z_k \leftarrow \psi_i, \psi_j, \psi_k$ ▷ get local copies of edge data
    $\upsilon_{ij}, \upsilon_{jk}, \upsilon_{ki} \leftarrow c_{ij}\omega_{ij}, c_{jk}\omega_{jk}, c_{ki}\omega_{ki}$
    $S_{ijk} \leftarrow s_{ij}s_{jk}s_{ki}$ ▷ compute branch index
    if $S_{ijk} < 0$ **then** ▷ branch triangle
      $\upsilon_{ki} \leftarrow -\upsilon_{ki}$ ▷ want transport to $\tau(i_1)$, not i_1
    end if
    if $s_{ij} < 0$ **then** ▷ make values at j consistent w/ i
      $z_j \leftarrow \overline{z}_j$
      $\upsilon_{ij} \leftarrow \ \ c_{ij}\upsilon_{ij}$
      $\upsilon_{jk} \leftarrow -c_{jk}\upsilon_{jk}$
    end if
    if $S_{ijk}s_{ki} < 0$ **then** ▷ make values at k consistent w/ i
      $z_k \leftarrow \overline{z}_k$
      $\upsilon_{ki} \leftarrow -c_{ki}\upsilon_{ki}$
      $\upsilon_{jk} \leftarrow \ \ c_{jk}\upsilon_{jk}$
    end if
    $\alpha_i^{jk} \leftarrow \arg(z_i)$ ▷ compute angles at triangle corners
    $\alpha_j^{ki} \leftarrow \alpha_i^{jk} + \upsilon_{ij} - \arg(e^{\imath\upsilon_{ij}} z_i/z_j)$
    $\alpha_k^{ij} \leftarrow \alpha_j^{ki} + \upsilon_{jk} - \arg(e^{\imath\upsilon_{jk}} z_j/z_k)$
    $\alpha_l^{ijk} \leftarrow \alpha_k^{ij} + \upsilon_{ki} - \arg(e^{\imath\upsilon_{ki}} z_k/z_i)$
    $\alpha_m^{ijk} \leftarrow \alpha_i^{jk} + (\alpha_l^{ijk} - \alpha_i^{jk})/2$ ▷ midpoint coordinate
    $n_{ijk} \leftarrow \frac{1}{2\pi}(\alpha_l^{ijk} - \alpha_i)$ ▷ compute zero index
    $\alpha_j^{ki} \leftarrow \alpha_j^{ki} - 2\pi n/3$ ▷ adjust zeros
    $\alpha_k^{ij} \leftarrow \alpha_k^{ij} - 4\pi n/3$
  end for
  return α, n, S
end procedure

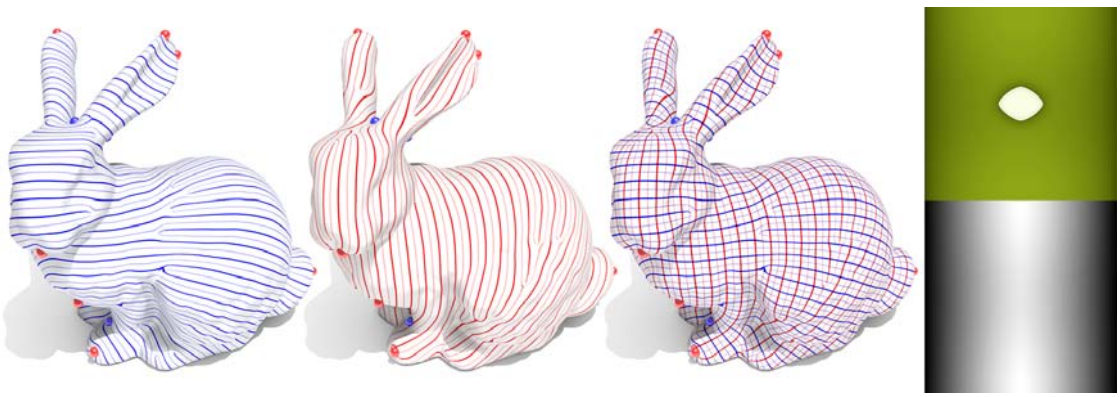

Figure 22. Two independent stripe patterns *(left, center left)* oriented along orthogonal fields are combined *(center right)* to drive the periodic texture and displacement maps *(right)* used in Fig. 1.

3 | Complex Line Bundles over Simplicial Complexes and their Applications

Felix Knöppel, Ulrich Pinkall
arXiv:1506.07853[math.DG]

Abstract. Discrete vector bundles are important in Physics and recently found remarkable applications in Computer Graphics. This article approaches discrete bundles from the viewpoint of Discrete Differential Geometry, including a complete classification of discrete vector bundles over finite simplicial complexes. In particular, we obtain a discrete analogue of a theorem of André Weil on the classification of hermitian line bundles. Moreover, we associate to each discrete hermitian line bundle with curvature a unique piecewise-smooth hermitian line bundle of piecewise constant curvature. This is then used to define a discrete Dirichlet energy which generalizes the well-known cotangent Laplace operator to discrete hermitian line bundles over Euclidean simplicial manifolds of arbitrary dimension.

1 Introduction

Vector bundles are fundamental objects in Differential Geometry and play an important role in Physics [**7**]. The Physics literature is also the main place where discrete versions of vector bundles were studied: First, there is a whole field called Lattice Gauge Theory where numerical experiments concerning connections in bundles over discrete spaces (lattices or simplicial complexes) are the main focus. Some of the work that has been done in this context is quite close to the kind of problems we are going to investigate here [**12, 13, 25**].

Vector bundles make their most fundamental appearance in Physics in the form of the complex line bundle whose sections are the wave functions of a charged particle in a magnetic field. Here the bundle comes with a connection whose curvature is given by the magnetic field [**7**]. There are situations where the problem itself suggests a natural discretization: The charged particle (electron) may be bound to a certain arrangement of atoms. Modelling this situation in such a way that the electron can only occupy a discrete set of locations then leads to the "tight binding approximation" [**37, 3, 64**].

Recently vector bundles over discrete spaces also have found striking applications in Geometry Processing and Computer Graphics. We will describe these in detail in Section 2.

In order to motivate the basic definitions concerning vector bundles over simplicial complexes let us consider a smooth manifold $\tilde{\mathrm{M}}$ that comes with smooth triangulation (Figure 1).

Let $\tilde{\mathrm{E}}$ be a smooth vector bundle over $\tilde{\mathrm{M}}$ of rank $\mathfrak{K}$. Then we can define a discrete version E of $\tilde{\mathrm{E}}$ by restricting $\tilde{\mathrm{E}}$ to the vertex set $\mathcal{V}$ of the triangulation. Thus E assigns to each vertex $i \in \mathcal{V}$ the $\mathfrak{K}$-dimensional real vector space $\mathrm{E}_i := \tilde{\mathrm{E}}_i$. This is the way vector bundles over simplicial complexes are defined in general: Such a bundle E assigns to each vertex i a $\mathfrak{K}$-dimensional real vector space E_i in such a way that $\mathrm{E}_i \cap \mathrm{E}_j = \emptyset$ for $i \neq j$.

So far the notion of a discrete vector bundle is completely uninteresting mathematically: The obvious definition of an isomorphism between two such bundles E and $\hat{\mathrm{E}}$ just would require vector space isomorphism $f_i \colon \mathrm{E}_i \to \hat{\mathrm{E}}_i$ for each vertex i. Thus, unless we put more structure on our bundles, any two vector bundles of the same rank over a simplicial complex are isomorphic.

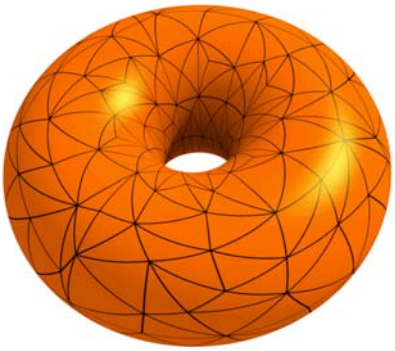

Figure 1. A smooth triangulation of a manifold.

Suppose now that $\tilde{\mathrm{E}}$ comes with a connection ∇. Then we can use the parallel transport along edges ij of the triangulation to define vector space isomorphisms

$$\eta_{ij} \colon \tilde{\mathrm{E}}_i \to \tilde{\mathrm{E}}_j$$

This leads to the standard definition of a connection on a vector bundle over a simplicial complex: Such a connection is given by a collection of isomorphisms $\eta_{ij} \colon \mathrm{E}_i \to \mathrm{E}_j$ defined for each edge ij such that

$$\eta_{ji} = \eta_{ij}^{-1}.$$

Now the classification problem becomes non-trivial because for an isomorphism f between two bundles E and $\hat{\mathrm{E}}$ with connection we have to require compatibility with the transport maps η_{ij}:

$$f_j \circ \eta_{ij} = \hat{\eta}_{ij} \circ f_i.$$

Given a connection η and a closed edge path $\gamma = e_\ell \cdots e_1$ (compare Section 4) of the simplicial complex we can define the monodromy $P_\gamma \in \mathrm{Aut}(\mathrm{E}_i)$ around γ as

$$P_\gamma = \eta_{e_\ell} \circ \ldots \circ \eta_{e_1}.$$

In particular the monodromies around triangular faces of the simplicial complex provide an analog for the smooth curvature in the discrete setting. In Section 4 we will classify vector bundles with connection in terms of their monodromies.

Let us look at the special case of a rank 2 bundle E that is oriented and comes with a Euclidean scalar product. Then the 90°-rotation in each fiber makes it into 1-dimensional complex vector space, so we effectively are dealing with a hermitian complex line bundle. If ijk is an oriented face of our simplicial complex, the monodromy $P_{\partial\, ijk} \colon \mathrm{E}_i \to \mathrm{E}_i$ around the triangle ijk is multiplication by a complex number h_{ijk} of norm one. Writing $h_{ijk} = e^{\imath\alpha_{ijk}}$ with $-\pi < \alpha_{ijk} \leq \pi$ we see that this monodromy can also be interpreted as a real curvature $\alpha_{ijk} \in (-\pi, \pi]$. It thus becomes apparent that the information provided by the connection η cannot encode any curvature that integrated over a single face is larger than $\pm\pi$. This can be a serious restriction for applications: We effectively see a cutoff for the curvature that can be contained in a single face.

Remember however our starting point: We asked for structure that can be naturally transferred from the smooth setting to the discrete one. If we think again about a triangulated smooth manifold it is clear that we can associate to each two-dimensional face ijk the integral Ω_{ijk} of the curvature 2-form over this face. This is just a discrete 2-form in the sense of discrete exterior calculus [**17**]. Including this discrete curvature 2-form with the parallel transport η brings discrete complex line bundles much closer to their smooth counterparts:

Definition. *A hermitian line bundle with curvature over a simplicial complex $\mathcal{X}$ is a triple $(\mathrm{E}, \eta, \Omega)$. Here E is complex hermitian line bundle over $\mathcal{X}$, for each edge ij the maps $\eta_{ij} \colon \mathrm{E}_i \to \mathrm{E}_j$ are unitary and the closed real-valued 2-form Ω on each face ijk satisfies*

$$\eta_{ki} \circ \eta_{jk} \circ \eta_{ij} = e^{\imath\Omega_{ijk}}\, \mathrm{id}_{\mathrm{E}_i}.$$

In Section 7 we will prove for hermitian line bundles with curvature the discrete analog of a well-known theorem by André Weil on the classification of hermitian line bundles.

In Section 8 we will define for hermitian line bundles with curvature a degree (which can be an arbitrary integer) and we will prove a discrete version of the Poincaré-Hopf index theorem concerning the number of zeros of a section (counted with sign and multiplicity).

Finally we will construct in Section 10 for each hermitian line bundle with curvature a piecewise-smooth bundle with a curvature 2-form that is constant on each face. Sections of the discrete bundle can be canonically extended to sections of the piecewise-smooth bundle. This construction will provide us with finite elements for bundle sections and thus will allow us to compute the Dirichlet energy on the space of sections.

2 Applications of Vector Bundles in Geometry Processing

Several important tasks in Geometry Processing (see the examples below) lead to the problem of coming up with an optimal normalized section ϕ of some Euclidean vector bundle E over a compact manifold with boundary M. Here "normalized section" means that ϕ is defined away from a certain singular set and where defined it satisfies $|\phi| = 1$.

In all the mentioned situations E comes with a natural metric connection ∇ and it turns out that the following method for finding ϕ yields surprisingly good results:

Among all sections ψ of E *find one which minimizes* $\int_{\mathrm{M}} |\nabla\psi|^2$ *under the constraint* $\int_{\mathrm{M}} |\psi|^2 = 1$. *Then away from the zero set of ψ use $\phi = \psi/|\psi|$.*

The term "optimal" suggests that there is a variational functional which is minimized by ϕ and this is in fact the case. Moreover, in each of the applications there are heuristic arguments indicating that ϕ is indeed a good choice for the problem at hand. For the details we refer to the original papers. Here we are only concerned with the Discrete Differential Geometry involved in the discretization of the above variational problem.

2.1 Direction Fields on Surfaces

Here M is a surface with a Riemannian metric, E = TM is the tangent bundle and ∇ is the Levi-Civita connection. Figure 2 shows the resulting unit vector field ϕ. If we consider TM as a

Figure 2. An optimal direction field on a surface.

complex line bundle, normalized sections of the tensor square $\mathrm{L} = \mathrm{TM} \otimes \mathrm{TM}$ describe unoriented direction fields on M. Similarly, "higher order direction fields" like cross fields are related to higher tensor powers of TM. Higher order direction fields also have important applications in Computer Graphics.

2.2 Stripe Patterns on Surfaces

A *stripe pattern* on a surface M is a map which away from a certain singular set assigns to each point $p \in \mathrm{M}$ an element $\phi(p) \in \mathbb{S} = \{z \in \mathbb{C} || z| = 1\}$. Such a map ϕ can be used to color M in a periodic fashion according to a color map that assigns a color to each point on the unit circle $\mathbb{S}$. Suppose we are given a 1-form ω on M that specifies a desired direction and spacing of the stripes, which means that ideally we would wish for something like $\phi = e^{i\alpha}$ with $d\alpha = \omega$. Then the algorithm in **[34]** says that we should use a ϕ that comes from taking E as the trivial bundle $\mathrm{E} = \mathrm{M} \times \mathbb{C}$ and $\nabla\psi = d\psi - i\omega\psi$. Sometimes the original data come from an unoriented direction field and (in order to obtain the 1-form ω) we first have to move from M to a double branched cover $\tilde{\mathrm{M}}$ of M. This is for example the case in Figure 3.

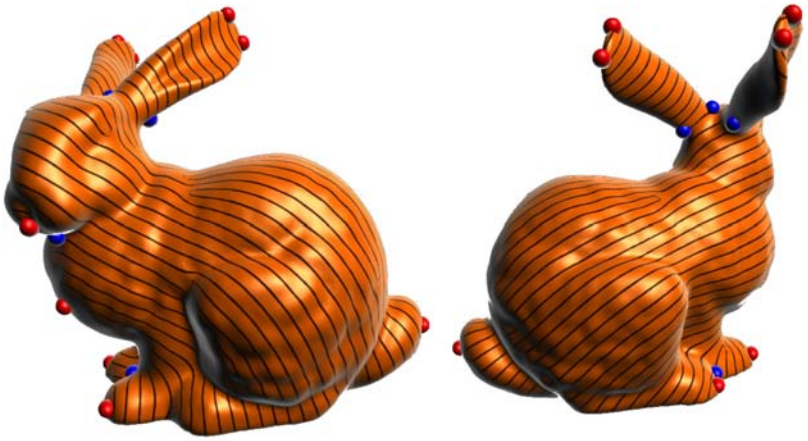

Figure 3. An optimal stripe pattern aligned to an unoriented direction field.

2.3 Decomposing Velocity Fields into Fields Generated by Vortex Filaments

The velocity fields that arise in fluid simulations quite often can be understood as a superposition of interacting vortex rings. It is therefore desirable to have an algorithm that reconstructs the underlying vortex filaments from a given velocity field. Let the velocity field $\mathbf{v}$ on a domain $\mathrm{M} \subset \mathbb{R}^3$ be given as a 1-form $\omega = \langle \mathbf{v}, \cdot \rangle$. Then the algorithm proposed in **[72]** uses the function $\phi\colon \mathrm{M} \to \mathbb{C}$ that results from taking the trivial bundle $\mathrm{E} = \mathrm{M} \times \mathbb{C}$ endowed with the connection $\nabla\psi = d\psi - i\omega\psi$. Note that so far this is just a three-dimensional version of the situation in Section 2.2. This time however we even forget ϕ in the end and only retain the zero set of ψ as the filament configuration we are looking for.

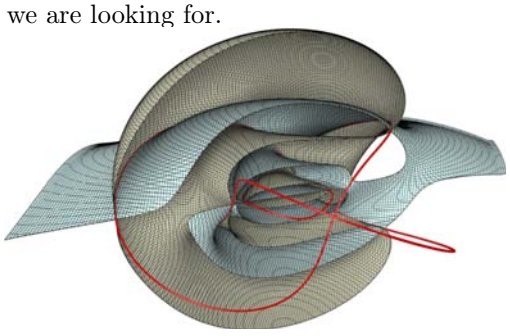

Figure 4. A knotted vortex filament defined as the zero set of a complex valued function ψ. It is shown as the intersection of the zero set of $\operatorname{Re}\psi$ with the zero set of $\operatorname{Im}\psi$.

2.4 Close-To-Conformal Deformations of Volumes

Here the data are a domain $\mathrm{M} \subset \mathbb{R}^3$ and a function $u\colon \mathrm{M} \to \mathbb{R}$. The task is to find a map $f\colon \mathrm{M} \to \mathbb{R}^3$ which is approximately conformal with conformal factor e^u, i.e. for all tangent vectors $X \in \mathrm{TM}$ we want

$$|df(X)| \approx e^u |X|.$$

The only exact solutions of this equations are the Möbius transformations. For these we find

$$df(X) = e^u \overline{\psi} X \psi$$

for some map $\psi\colon \mathrm{M} \to \mathbb{H}$ with $|\psi| = 1$ which in addition satisfies

$$d\psi(X) = -\tfrac{1}{2}(\operatorname{grad} u \times X)\, \psi.$$

Note that here we have identified $\mathbb{R}^3$ with the space of purely imaginary quaternions. Let us define a connection ∇ on the trivial rank 4 vector bundle $\mathrm{M} \times \mathbb{H}$ by

$$\nabla_X \psi := d\psi(X) + \tfrac{1}{2}(\operatorname{grad} u \times X)\psi.$$

Then we can apply the usual method and find a section $\phi\colon \mathrm{M} \to \mathbb{H}$ with $|\phi| = 1$. In general there will not be any $f\colon \mathrm{M} \to \mathbb{R}^3$ that satisfies

$$df(X) = e^u \overline{\phi} X \phi \tag{2.1}$$

exactly but we can always look for an f that satisfies (2.1) in the least squares sense. See Figure 5 for an example.

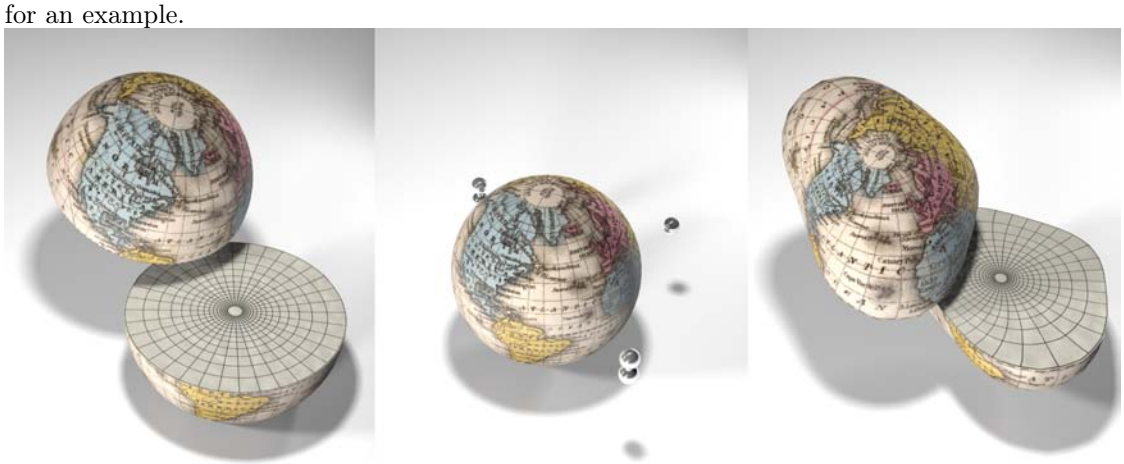

Figure 5. Close-to-conformal deformation of a sphere based on a desired conformal factor specified as the potential of a collection of point charges.

3 Discrete Vector Bundles with Connection

An *(abstract) simplicial complex* is a collection $\mathcal{X}$ of finite non-empty sets such that if σ is an element of $\mathcal{X}$ so is every non-empty subset of σ [**49**].

An element of a simplicial complex $\mathcal{X}$ is called a *simplex* and each non-empty subset of a simplex σ is called a *face of* σ. The elements of a simplex are called *vertices* and the *dimension of a simplex* is defined to be one less than the number of its vertices: $\dim \sigma := |\sigma| - 1$. A simplex of dimension k is also called a *k-simplex*. The *dimension of a simplicial complex* is defined as the maximal dimension of its simplices.

To avoid technical difficulties, we will restrict our attention to *finite* simplicial complexes only. The main concepts are already present in the finite case. Though, the definitions carry over verbatim to infinite simplicial complexes and several statements remain true in this case.

Definition 1. *Let $\mathbb{F}$ be a field and let $\mathcal{X}$ be a simplicial complex with vertex set $\mathcal{V}$. A discrete $\mathbb{F}$-vector bundle E of rank $\mathfrak{K} \in \mathbb{N}$ over $\mathcal{X}$ is a map $\pi\colon \mathrm{E} \to \mathcal{V}$ such that for each vertex $i \in \mathcal{V}$ the fiber over i*

$$\mathrm{E}_i := \pi^{-1}(\{i\})$$

has the structure of a $\mathfrak{K}$-dimensional $\mathbb{F}$-vector space. We slightly abuse notation and denote a discrete vector bundle over a simplicial complex schematically by $\mathrm{E} \to \mathcal{X}$.

Clearly, the fibers can be equipped with additional structures. In particular, a real vector bundle whose fibers are Euclidean vector spaces is called a *discrete Euclidean vector bundle*. Similarly, a complex vector bundle whose fibers are hermitian vector spaces is called a *discrete hermitian vector bundle*.

Now, let $\sigma = \{i_0, \dots, i_k\}$ be a k-simplex. We define two orderings of its vertices to be equivalent if they differ by an even permutation. Such an equivalence class is then called an *orientation* of σ and a simplex together with an orientation is called an *oriented simplex*. We will denote the oriented k-simplex just by the word $i_0 \cdots i_k$. Further, an oriented 1-simplex is simply called an *edge*.

Definition 2. *Let $\mathrm{E} \to \mathcal{X}$ be a discrete vector bundle over a simplicial complex. A discrete connection on E is a map η which assigns to each edge ij an isomorphism $\eta_{ij}\colon \mathrm{E}_i \to \mathrm{E}_j$ of vector spaces such that*

$$\eta_{ji} = \eta_{ij}^{-1}.$$

Here and in the following a morphism of vector spaces is a linear map that also preserves all additional structures - if any present. E.g., if we are dealing with hermitian vector spaces, then a morphism is a complex-linear map that preserves the hermitian metric, i.e. it is a complex linear isometric immersion. Now let us define morphisms of discrete vector bundles with connection.

Definition 3. *A morphism of discrete vector bundles with connection is a map $f\colon \mathrm{E} \to \mathrm{F}$ between discrete vector bundles $\mathrm{E} \to \mathcal{X}$ and $\mathrm{F} \to \mathcal{X}$ with connections η and θ (resp.) such that*

i) for each vertex i we have that $f(\mathrm{E}_i) \subset \mathrm{F}_i$ and the map $f_i = f|_{\mathrm{E}_i} : \mathrm{E}_i \to \mathrm{F}_i$ is a morphism of vector spaces,

ii) for each edge ij the following diagram commutes:

$$\begin{array}{ccc} \mathrm{F}_i & \xrightarrow{\theta_{ij}} & \mathrm{F}_j \\ {\scriptstyle f_i}\uparrow & = & \uparrow{\scriptstyle f_j} \\ \mathrm{E}_i & \xrightarrow[\eta_{ij}]{} & \mathrm{E}_j \end{array},$$

i.e. $\theta_{ij} \circ f_i = f_j \circ \eta_{ij}$.

Clearly, an *isomorphism* is a morphism which has an inverse map, which is also a morphism. Two discrete vector bundles with connection are called *isomorphic*, if there exists an isomorphism between them. Again let $\mathcal{V}$ denote the vertex set of $\mathcal{X}$. A discrete vector bundle $\mathrm{E} \to \mathcal{X}$ with connection η is called *trivial*, if it is isomorphic to the *product bundle*

$$\underline{\mathbb{F}^{\mathfrak{K}}} := \mathcal{V} \times \mathbb{F}^{\mathfrak{K}}$$

over $\mathcal{X}$ equipped with the connection which assigns to each edge the identity $\mathrm{id}_{\mathbb{F}^{\mathfrak{K}}}$.

Let $\mathrm{E} \to \mathcal{X}$ be a discrete vector bundle with connection and let $\mathcal{V}$ denote the vertex set of $\mathcal{X}$. A *section* of a discrete vector bundle $\mathrm{E} \to \mathcal{X}$ is a map $\psi\colon \mathcal{V} \to \mathrm{E}$ such that the following diagram commutes

$$\begin{array}{ccc} & & \mathrm{E} \\ & {\scriptstyle \psi}\nearrow & = \downarrow{\scriptstyle \pi} \\ \mathcal{V} & \xrightarrow[id]{} & \mathcal{V} \end{array},$$

i.e. $\pi \circ \psi = id$. As usual, the space of sections of E will be denoted by $\Gamma(\mathrm{E})$.

Definition 4. *Let $\mathrm{E} \to \mathcal{X}$ be a discrete vector bundle with connection η. A section $\Phi \in \Gamma(\mathrm{E})$ is called parallel, if $\eta_{ij}(\phi_i) = \phi_j$ for each edge ij of $\mathcal{X}$.*

Proposition 1. *A discrete vector bundle* $\mathrm{E} \to \mathcal{X}$ *with connection of rank* $\mathfrak{K}$ *is trivial if and only if it has* $\mathfrak{K}$ *linearly independent parallel sections.*

Proof. Let E be trivial. Then there is an isomorphism $f\colon \mathrm{E} \to \underline{\mathbb{F}^{\mathfrak{K}}}$. Parallel sections of the trivial bundle are just constant maps $\mathcal{V} \to \mathbb{F}^{\mathfrak{K}}$. For $j = 1, \dots, \mathfrak{K}$, we define sections ϕ^j by $\phi^j_i := f^{-1}((i, \epsilon_j))$, where ϵ_j denotes the j-th canonical basis vector of $\mathbb{F}^{\mathfrak{K}}$. Since f is an isomorphism the ϕ^j is parallel. Clearly, these sections are linearly independent. Conversely, given $\mathfrak{K}$ linearly independent parallel sections, these form at each vertex i a basis of the fiber E_i. The corresponding coordinates establish an isomorphism with the trivial bundle. □

Clearly, each vector space operation gives rise to an operation on discrete vector bundles with connection. E.g. if $\mathrm{E} \to \mathcal{X}$ and $\mathrm{F} \to \mathcal{X}$ are discrete vector bundles with connection, then the tensor product $\mathrm{E} \otimes \mathrm{F} \to \mathcal{X}$ is the discrete vector bundle with fiber $(\mathrm{E} \otimes \mathrm{F})_i = \mathrm{E}_i \otimes \mathrm{F}_i$ over the vertex i. If η and θ denote the connections of E and F (resp.), then the connection $\eta \otimes \theta$ on $\mathrm{E} \otimes \mathrm{F}$ is simply given by $(\eta \otimes \theta)_{ij} = \eta_{ij} \otimes \theta_{ij}$. Thus we can build direct sums, tensor products and duals of discrete vector bundles.

Let E and F be discrete vector bundles with connections η and θ, respectively. If $f\colon \mathrm{E} \to \mathrm{F}$ is an isomorphism then, by the commutative edge diagrams, we obtain for each edge ij the following relation:

$$\theta_{ij} \circ f_i \circ \eta_{ij}^{-1} = f_j$$

If we regard f as a section of the tensor product $\mathrm{F} \otimes \mathrm{E}^*$, then the above equation states that f is parallel. Conversely, if $\operatorname{rank} \mathrm{E} = \operatorname{rank} \mathrm{F}$, every non-vanishing parallel section of $\mathrm{F} \otimes \mathrm{E}^*$ yields an isomorphism between E and F.

Proposition 2. *Two vector bundles* E *and* F *of equal rank are isomorphic if and only if* $\mathrm{F} \otimes \mathrm{E}^*$ *has a non-vanishing parallel section. In particular,* $\mathrm{E} \otimes \mathrm{E}^*$ *is trivial.*

It is a natural question to ask how many non-isomorphic discrete vector bundles with connection exist on a given simplicial complex. This question is related to the topology of the simplicial complex and can be studied by monodromy.

4 Monodromy - A Discrete Analogue of Kobayashi's Theorem

Let $\mathcal{X}$ be a simplicial complex. Each edge of $\mathcal{X}$ has a start and a target vertex. We denote the map that sends an edge to its start vertex by s and the map that sends the edge to its target vertex by t:

$$s(ij) := i, \quad t(ij) := j.$$

A *(discrete) path* γ is then simply a sequence of successive edges $(e_1, \dots, e_\ell)$, i.e. $s(e_{k+1}) = t(e_k)$ for all $k = 1, \dots, \ell - 1$, and will be denoted by the word:

$$\gamma = e_\ell \cdots e_1.$$

A path from i to j is a path $\gamma = e_\ell \cdots e_1$ such that $i = s(e_1)$ and $j = t(e_\ell)$. We also say that γ starts at i and ends at j. If $\gamma = e_m \cdots e_1$ is a path from i to j and $\tilde{\gamma} = e_\ell \cdots e_{m+1}$ is a path from j to k, then we can define a new path $\tilde{\gamma}\gamma$ from i to k as follows:

$$\tilde{\gamma}\gamma = e_\ell \cdots e_1.$$

The path $\tilde{\gamma}\gamma$ is called the *concatenation* of γ and $\tilde{\gamma}$. In this sense we can regard an edge e as an *elementary path* from its start to its target vertex. With this identification the inverse e^{-1} of an elementary path $e = ij$ is then given by its opposite edge, i.e. $e^{-1} = ji$. The *inverse* of a path $\gamma = e_\ell \cdots e_1$ is then defined by

$$\gamma^{-1} := e_1^{-1} \cdots e_\ell^{-1}.$$

Let $\mathrm{E} \to \mathcal{X}$ be a discrete vector bundle with connection η. Now, given a discrete path $\gamma = e_\ell \cdots e_1$ from i to j, we define the *parallel transport along* γ as the map $P_\gamma\colon \mathrm{E}_i \to \mathrm{E}_j$ given by

$$P_\gamma := \eta_{e_\ell} \circ \cdots \circ \eta_{e_1}.$$

Proposition 3. *Let $\mathrm{E} \to \mathcal{X}$ be a discrete vector bundle with connection η and let γ and $\tilde{\gamma}$ be discrete paths in $\mathcal{X}$ such that $\tilde{\gamma}$ starts where γ ends. Then:*

$$P_{\tilde{\gamma}\gamma} = P_{\tilde{\gamma}} \circ P_{\gamma}, \quad P_{\gamma^{-1}} = P_{\gamma}^{-1}.$$

Proof. The proposition obviously follows from the definitions. □

Proposition 4. *Let $f\colon \mathrm{E} \to \tilde{\mathrm{E}}$ be an isomorphism of discrete vector bundles. Let P and $\tilde{P}$ denote the parallel transport on E and $\tilde{\mathrm{E}}$, respectively. Then, for each path γ from a vertex i to a vertex j,*

$$\tilde{P}_{\gamma} = f_j \circ P_{\gamma} \circ f_i^{-1}.$$

Proof. Denote the connections of E and $\tilde{\mathrm{E}}$ by η and $\tilde{\eta}$, respectively. Since f is an isomorphism, the f_i are invertible and we can express $\tilde{\eta}$ for each edge e as follows

$$\tilde{\eta}_e = f_{t(e)} \circ \eta_e \circ f_{s(e)}^{-1}$$

Now, let $\gamma = e_1 \cdots e_\ell$ be a path from the vertex i to the vertex j. Since $s(e_1) = i$, $t(e_\ell) = j$ and $s(e_{k+1}) = t(e_k)$ for $0 \leq k < \ell$, we obtain

$$\tilde{P}_{\gamma} = \tilde{\eta}_{e_\ell} \circ \cdots \circ \tilde{\eta}_{e_1} = f_{t(e_\ell)} \circ \eta_{e_\ell} \circ \cdots \circ \eta_{e_1} \circ f_{s(e_1)}^{-1} = f_j \circ P_{\gamma} \circ f_i^{-1},$$

as was claimed. □

A *loop based at a vertex i* is a path that starts and ends at i. The *loop space based at i* is then the set $\mathcal{LS}(\mathcal{X}, i)$ of all loops based at i. To extract the essential information out of parallel transport we will consider certain loops as equivalent.

A *spike* is a path of the form $e^{-1}e$. Clearly, if a loop contains a spike, we can delete the spike and obtain a new loop based at the same vertex:

$$e_\ell \cdots e_{k+1} e^{-1} e \, e_k \cdots e_1 \longrightarrow e_\ell \cdots e_{k+1} e_k \cdots e_1.$$

Similarly certain spikes can be inserted into loops. These operations, deleting or inserting spikes, will be referred to as *elementary moves*. We define an equivalence relation on the loop space $\mathcal{LS}(\mathcal{X}, i)$ as follows:

$$\gamma \sim \tilde{\gamma} :\Longleftrightarrow \tilde{\gamma} \text{ can be obtained from } \gamma \text{ by a sequence of elementary moves.}$$

The concatenation of discrete paths induces a group structure on the quotient space $\mathcal{LG}(\mathcal{X}, i) := \mathcal{LS}(\mathcal{X}, i)/_{\sim}$:

$$[\tilde{\gamma}][\gamma] = [\tilde{\gamma}\gamma], \quad [\gamma]^{-1} = [\gamma^{-1}].$$

The group $\mathcal{LG}(\mathcal{X}, i)$ is called the *discrete path group in $\mathcal{X}$ with base point i*. In the smooth case, the path group appears e.g. in **[35]** and more recently in **[47]**.

Remark 1: *The k-skeleton of a simplicial complex $\mathcal{X}$ is the simplicial complex formed by all simplices in $\mathcal{X}$ of dimension $\leq k$. Clearly, $\mathcal{LG}(\mathcal{X}, i)$ is nothing else than the first fundamental group of the 1-skeleton of $\mathcal{X}$.*

If $\mathcal{X}$ is *connected*, i.e. any two vertices i and j of $\mathcal{X}$ can be joined by a path, then the groups $\mathcal{LG}(\mathcal{X}, i)$ and $\mathcal{LG}(\mathcal{X}, j)$ are isomorphic. An isomorphism is established by conjugation with any path γ from i to j. By Proposition 1, it is clear that all discrete vector bundles over a connected simplicial complexes with vanishing path group must be trivial. If the path group does not vanish, there are obvious obstructions. These are encoded by the monodromy of the bundle.

Proposition 5. *Let $\mathrm{E} \to \mathcal{X}$ be a discrete vector bundle with connection over a connected simplicial complex. The parallel transport pushes forward to a representation of the loop group with base point i:*

$$\mathfrak{M}\colon \mathcal{LG}(\mathcal{X}, i) \to \mathrm{Aut}(\mathrm{E}_i), \quad [\gamma] \mapsto P_{\gamma}.$$

The representation $\mathfrak{M}$ will be called the monodromy of discrete vector bundle E.

Proof. Obviously, the parallel transport is invariant under elementary moves. Hence $\mathfrak{M}$ is well-defined. That $\mathfrak{M}$ is a group homomorphism is just Proposition 3. □

Isomorphy of discrete vector bundles carries over to their monodromy as follows.

Proposition 6. *Isomorphic discrete vector bundles with connection have isomorphic monodromies, i.e. the monodromies lie in the same conjugacy class.*

Proof. Let $f\colon \mathrm{E} \to \tilde{\mathrm{E}}$ be an isomorphism of discrete vector bundles with connection over a simplicial complex $\mathcal{X}$. Then, by Proposition 4, the monodromies $\mathfrak{M}\colon \mathcal{LG}(\mathcal{X}, i) \to \mathrm{Aut}(\mathrm{E}_i)$ and $\tilde{\mathfrak{M}}\colon \mathcal{LG}(\mathcal{X}, i) \to \mathrm{Aut}(\tilde{\mathrm{E}}_i)$ are related as follows:

$$\tilde{\mathfrak{M}}([\gamma]) = f_i \circ \mathfrak{M}([\gamma]) \circ f_i^{-1}, \text{ for each } [\gamma] \in \mathcal{LG}(\mathcal{X}, i).$$

But this means that $\mathfrak{M}$ and $\tilde{\mathfrak{M}}$ are isomorphic representations. □

In fact, as we will see, the monodromy completely determines a discrete vector bundle with connection up to isomorphism. This provides a complete classification of discrete vector bundles with connection.

Let $\mathcal{X}$ be a connected simplicial complex. Let $\mathrm{E} \to \mathcal{X}$ be a discrete $\mathbb{F}$-vector bundle of rank $\mathfrak{K}$ with connection and let $\mathfrak{M}\colon \mathcal{LG}(\mathcal{X}, i) \to \mathrm{Aut}(\mathrm{E}_i)$ denote its monodromy. Any choice of a basis of the fiber E_i determines a group homomorphism $\rho \in \mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathrm{GL}(\mathfrak{K}, \mathbb{F})\big)$. Any different choice of basis determines a group homomorphism $\tilde{\rho}$ which is related to ρ by conjugation, i.e. there is $S \in \mathrm{GL}(\mathfrak{K}, \mathbb{F})$ such that

$$\tilde{\rho}([\gamma]) = S \cdot \rho([\gamma]) \cdot S^{-1} \text{ for all } [\gamma] \in \mathcal{LG}(\mathcal{X}, i).$$

Hence the monodromy $\mathfrak{M}$ determines a well-defined conjugacy class of group homomorphisms from $\mathcal{LG}(\mathcal{X}, i)$ to $\mathrm{GL}(\mathfrak{K}, \mathbb{F})$, which we will simply denote by $[\mathfrak{M}]$. The group $\mathrm{GL}(\mathfrak{K}, \mathbb{F})$ will be referred to as the *structure group of* E.

Let $\mathfrak{V}_{\mathbb{F}}^{\mathfrak{K}}(\mathcal{X})$ denote the *set of isomorphism classes $\mathbb{F}$-vector bundles of rank $\mathfrak{K}$ with connection over* $\mathcal{X}$ and let $\mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathrm{GL}(\mathfrak{K}, \mathbb{F})\big)/_{\sim}$ denote the *set of conjugacy classes of group homomorphisms from the path group $\mathcal{LG}(\mathcal{X}, i)$ into the structure group* $\mathrm{GL}(\mathfrak{K}, \mathbb{F})$. The following theorem is a discrete analogue of Kobayashi's theorem on smooth bundles (compare **[35]**).

Theorem 2. $F\colon \mathfrak{V}_{\mathbb{F}}^{\mathfrak{K}}(\mathcal{X}) \to \mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathrm{GL}(\mathfrak{K}, \mathbb{F})\big)/_{\sim}$, $[\mathrm{E}] \mapsto [\mathfrak{M}]$ *is bijective.*

Proof. By Proposition 6, F is well-defined. First we show injectivity. Consider two discrete vector bundles E and $\tilde{\mathrm{E}}$ over $\mathcal{X}$ with connections η and $\tilde{\eta}$, respectively, and let $\mathfrak{M}$ and $\tilde{\mathfrak{M}}$ denote their monodromies. Suppose that $[\mathfrak{M}] = [\tilde{\mathfrak{M}}]$. Hence, if we choose bases $\{V_1, \dots, V_{\mathfrak{K}}\}$ of E_i and $\{\tilde{V}_1, \dots, \tilde{V}_{\mathfrak{K}}\}$ of $\tilde{\mathrm{E}}_i$, then $\mathfrak{M}$ and $\tilde{\mathfrak{M}}$ are represented by group homomorphisms $\rho, \tilde{\rho} \in \mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathrm{GL}(\mathfrak{K}, \mathbb{F})\big)$ (resp.) both of which are related by conjugation and, without loss of generality, we can assume that $\rho = \tilde{\rho}$. Now, let $\mathcal{T}$ be a spanning tree of $\mathcal{X}$ with root i. Then, for each vertex j of $\mathcal{X}$ there is a path $\gamma_{i,j}$ from the root i to the vertex j entirely contained in $\mathcal{T}$. Since the $\mathcal{T}$ contains no loops the path $\gamma_{i,j}$ is essentially unique, i.e. any two such paths differ by a sequence of elementary moves. Thus, we can extend the bases parallelly along $\mathcal{T}$ to each vertex of $\mathcal{X}$ and obtain sections $\{X^1, \dots, X^{\mathfrak{K}}\} \subset \Gamma(\mathrm{E})$ and $\{\tilde{X}^1, \dots, \tilde{X}^{\mathfrak{K}}\} \subset \Gamma(\tilde{\mathrm{E}})$ providing bases at each fiber. With respect to these bases the connections η and $\tilde{\eta}$ are represented by elements of $\mathrm{GL}(\mathfrak{K}, \mathbb{F})$. Clearly, by construction, for each edge e in $\mathcal{T}$ the connection is represented by just the identity matrix. Moreover, to each edge $e = jk$ not contained in $\mathcal{T}$ there corresponds a unique loop $[\gamma_e] \in \mathcal{LG}(\mathcal{X}, i)$. With the notation above, it is given by $\gamma_e = \gamma_{i,k}^{-1}\, e\, \gamma_{i,j}$. In particular, on the edge e both connections are represented by the same matrix $\rho([\gamma_e]) = \tilde{\rho}([\gamma_e])$. Thus if we define $f\colon \mathrm{E} \to \tilde{\mathrm{E}}$ such that $f(X^m) := \tilde{X}^m$ for $m = 1, \dots, \mathfrak{K}$ we obtain an isomorphism, i.e. $\mathrm{E} \cong \tilde{\mathrm{E}}$. Hence F is injective. Now, let $\rho \in \mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathrm{GL}(\mathfrak{K}, \mathbb{F})\big)$. To see that F is surjective we use $\mathcal{T}$ to equip the product bundle $\mathrm{E} := \mathcal{V} \times \mathbb{F}^{\mathfrak{K}}$ with a particular connection η. Namely, if e lies in $\mathcal{T}$ we set $\eta_e = \mathrm{id}$ else we set $\eta_e := \rho([\gamma_e])$. Clearly, by construction, $F([\mathrm{E}]) = [\rho]$. Thus F is surjective. □

5 Discrete Line Bundles - The Abelian Case

Let $\mathcal{X}$ be a connected simplicial complex. A *discrete line bundle* is a discrete vector bundle $\mathrm{L} \to \mathcal{X}$ of rank $\mathfrak{K} = 1$. In this case the structure group is the multiplicative group of the underlying field $\mathbb{F}_* := \mathbb{F} \setminus \{0\}$. Since $\mathbb{F}_*$ is abelian, we obtain

$$\mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathbb{F}_*)\big)/_\sim = \mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathbb{F}_*\big).$$

Clearly, $\mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathbb{F}_*\big)$ carries a natural group structure. Moreover, the isomorphism classes of discrete line bundles over $\mathcal{X}$ itself build an abelian group. The group structure is just given by the tensor product: Let $[\mathrm{L}], [\tilde{\mathrm{L}}] \in \mathfrak{V}^1_{\mathbb{F}}(\mathcal{X})$, then

$$[\mathrm{L}][\tilde{\mathrm{L}}] = [\mathrm{L} \otimes \tilde{\mathrm{L}}], \quad [\mathrm{L}]^{-1} = [\mathrm{L}^*].$$

The identity element is given by the trivial bundle. In the following we will denote the *group of isomorphism classes of* $\mathbb{F}$*-line bundles over* $\mathcal{X}$ by $\mathcal{L}^{\mathbb{F}}_{\mathcal{X}}$.

It is easily checked that the map $F\colon \mathcal{L}^{\mathbb{F}}_{\mathcal{X}} \to \mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathbb{F}_*\big)$, $[\mathrm{L}] \mapsto [\mathfrak{M}]$ is a group homomorphism. By Theorem 2, F is an isomorphism.

Now, since $\mathbb{F}_*$ is abelian, each homomorphism $\rho \in \mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathbb{F}_*\big)$ factors through the *abelianization*

$$\mathcal{LG}(\mathcal{X}, i)_{ab} = \mathcal{LG}(\mathcal{X}, i)/[\mathcal{LG}(\mathcal{X}, i), \mathcal{LG}(\mathcal{X}, i)],$$

i.e. for each $\rho \in \mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathbb{F}_*\big)$ there is a unique $\rho_{ab} \in \mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i)_{ab}, \mathbb{F}_*\big)$ such that

$$\rho = \rho_{ab} \circ \pi_{ab}.$$

Here $\pi_{ab}\colon \mathcal{LG}(\mathcal{X}, i) \to \mathcal{LG}(\mathcal{X}, i)_{ab}$ denotes the canonical projection. This yields an isomorphism between $\mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i), \mathbb{F}_*\big)$ and $\mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i)_{ab}, \mathbb{F}_*\big)$. In particular,

$$\mathcal{L}^{\mathbb{F}}_{\mathcal{X}} \cong \mathrm{Hom}\big(\mathcal{LG}(\mathcal{X}, i)_{ab}, \mathbb{F}_*\big).$$

Actually, as we will see, the abelianization $\mathcal{LG}(\mathcal{X}, i)_{ab}$ is naturally isomorphic to the group of closed 1-chains.

The *group of* k*-chains* $\mathrm{C}_k(\mathcal{X}, \mathbb{Z})$ is defined as the free abelian group which is generated by the k-simplices of $\mathcal{X}$. More precisely, let $\mathcal{X}^{or}_k$ denote the *set of oriented* k*-simplices of* $\mathcal{X}$. Clearly, for $k > 0$, each k-simplex has two orientations. Interchanging these orientations yields a fixed-point-free involution $\rho_k\colon \mathcal{X}^{or}_k \to \mathcal{X}^{or}_k$. The group of k-chains is then explicitly given as follows:

$$\mathrm{C}_k(\mathcal{X}, \mathbb{Z}) := \big\{c\colon \mathcal{X}^{or}_k \to \mathbb{Z} \mid c \circ \rho_k = -c\big\}.$$

Since simplices of dimension zero have only one orientation, $\mathcal{X}^{or}_0 = \mathcal{X}_0$. Thus,

$$\mathrm{C}_0(\mathcal{X}, \mathbb{Z}) := \big\{c\colon \mathcal{X}^{or}_k \to \mathbb{Z}\big\}.$$

It is common to identify an oriented k-simplex σ with its *elementary* k*-chain*, i.e. the chain which is 1 for σ, -1 for the oppositely oriented simplex and zero else. With this identification a k-chain c can be written as a formal sum of oriented k-simplices with integer coefficients:

$$c = \sum_{i=1}^{m} n_i \sigma_i, \quad n_i \in \mathbb{Z},\ \sigma_i \in \mathcal{X}^{or}_k.$$

The *boundary operator* $\partial_k\colon \mathrm{C}_k(\mathcal{X}, \mathbb{Z}) \to \mathrm{C}_{k-1}(\mathcal{X}, \mathbb{Z})$ is then the homomorphism which is uniquely determined by

$$\partial_k\, i_0 \cdots i_k = \sum_{j=0}^{k} (-1)^j\, i_0 \cdots \widehat{i_j} \cdots i_k.$$

It well-known and easily checked that $\partial_k \circ \partial_{k+1} \equiv 0$. Thus we get a chain complex

$$0 \xleftarrow{\partial_0} \mathrm{C}_0(\mathcal{X}, \mathbb{Z}) \xleftarrow{\partial_1} \mathrm{C}_1(\mathcal{X}, \mathbb{Z}) \xleftarrow{\partial_2} \cdots \xleftarrow{\partial_k} \mathrm{C}_k(\mathcal{X}, \mathbb{Z}) \xleftarrow{\partial_{k+1}} \cdots .$$

The *simplicial Homology groups* $\mathrm{H}_k(\mathcal{X}, \mathbb{Z})$ measure how exact this sequence is:

$$\mathrm{H}_k(\mathcal{X}, \mathbb{Z}) := \ker \partial_k / \mathrm{im}\, \partial_{k+1}.$$

The elements of $\ker \partial_k$ are called k*-cycles*, those of $\mathrm{im}\, \partial_{k+1}$ are called k*-boundaries*.

It is a well-known fact that the abelianization of the first fundamental group is the first homology group (see **[26]**). Now, if we combine this with the fact that $\mathcal{LG}(\mathcal{X}, i)$ is a nothing but the first fundamental group of the 1-skeleton of $\mathcal{X}$ and the first homology of the 1-skeleton consists exactly of all closed chains of $\mathcal{X}$, we see that

$$\mathcal{LG}(\mathcal{X}, i)_{ab} \cong \ker \partial_1.$$

The isomorphism is induced by the map $\mathcal{LG}(\mathcal{X}, i) \to \ker \partial_1$ given by $[\gamma] \mapsto \sum_j e_j$, where $\gamma = e_\ell \cdots e_1$. We summarize the above discussion in the following theorem.

Theorem 3. *The group of isomorphism classes of line bundles $\mathcal{L}^{\mathbb{F}}_{\mathcal{X}}$ is naturally isomorphic to the group* $\mathrm{Hom}(\ker \partial_1, \mathbb{F}_*)$*:*

$$\mathcal{L}^{\mathbb{F}}_{\mathcal{X}} \cong \mathrm{Hom}(\ker \partial_1, \mathbb{F}_*).$$

The isomorphism of Theorem 3 can be made explicit using discrete $\mathbb{F}_*$-valued 1-forms associated to the connection of a discrete line bundle.

6 Discrete Connection Forms

Let $\mathcal{X}$ denote a connected simplicial complex. A discrete k-form is nothing else than a k-cochain with coefficients in an abelian group. The exterior derivative survives as the coboundary operator.

Definition 5. *Let $\mathfrak{G}$ be an abelian group. The group of $\mathfrak{G}$-valued discrete k-forms is defined as follows:*

$$\Omega^k(\mathcal{X}, \mathfrak{G}) := \big\{\omega \colon \mathrm{C}_k(\mathcal{X}) \to \mathfrak{G} \mid \omega \text{ group homomorphism}\big\}.$$

The discrete exterior derivative d_k is then defined to be the adjoint of ∂_{k+1}, i.e.

$$d_k \colon \Omega^k(\mathcal{X}, \mathfrak{G}) \to \Omega^{k+1}(\mathcal{X}, \mathfrak{G}), \quad d_k\omega := \omega \circ \partial_{k+1}.$$

By construction, we immediately get that $d_{k+1} \circ d_k \equiv 0$. The corresponding cochain complex is called the *discrete de Rahm complex with coefficients in* $\mathfrak{G}$:

$$0 \to \Omega^0(\mathcal{X}, \mathfrak{G}) \xrightarrow{d_0} \Omega^1(\mathcal{X}, \mathfrak{G}) \xrightarrow{d_1} \cdots \xrightarrow{d_{k-1}} \Omega^k(\mathcal{X}, \mathfrak{G}) \xrightarrow{d_k} \cdots .$$

Analogous to the construction of the homology groups, the *k-th de Rahm Cohomology group* $\mathrm{H}^k(\mathcal{X}, \mathfrak{G})$ *with coefficients in* $\mathfrak{G}$ is defined as the quotient group

$$\mathrm{H}^k(\mathcal{X}, \mathfrak{G}) := \ker d_k / \mathrm{im}\, d_{k-1}.$$

The discrete k-forms in $\ker d_k$ are called *closed*, those in $\mathrm{im}\, d_{k-1}$ are called *exact*.

Now, let $\mathfrak{C}_{\mathrm{L}}$ denote the *space of connections* on the discrete $\mathbb{F}$-line bundle $\mathrm{L} \to \mathcal{X}$:

$$\mathfrak{C}_{\mathrm{L}} := \big\{\eta \mid \eta \text{ connection on } \mathrm{L}\big\}.$$

Clearly, any two connections $\eta, \theta \in \mathfrak{C}_{\mathrm{L}}$ differ by a discrete 1-form $\omega \in \Omega^1(\mathcal{X}, \mathbb{F}_*)$:

$$\theta = \omega\eta.$$

Hence the group $\Omega^1(\mathcal{K}, \mathbb{F}_*)$ acts simply transitively on the space of connections $\mathfrak{C}_{\mathrm{L}}$. In particular, each choice of a *base connection* $\beta \in \mathfrak{C}_{\mathrm{L}}$ establishes an identification

$$\mathfrak{C}_{\mathrm{L}} \ni \eta = \omega\beta \longleftrightarrow \omega \in \Omega^1(\mathcal{K}, \mathbb{F}_*).$$

Remark 2: *Note that each discrete vector bundle admits a trivial connection. To see this just choose for each vertex a basis of the corresponding fiber. The corresponding coordinates establish an identification with the product bundle. Then there is a unique connection that makes the diagrams over all edges commute.*

Definition 6. *Let $\eta \in \mathfrak{C}_{\mathrm{L}}$. A connection form representing the connection η is a 1-form $\omega \in \Omega^1(\mathcal{X}, \mathbb{F}_*)$ such that $\eta = \omega\beta$ for some trivial base connection β.*

Clearly, there are many connection forms representing a connection. We want to see how two such forms are related.

More generally, two connections η and θ in $\mathfrak{C}_{\mathrm{L}}$ lead to isomorphic discrete line bundles if and only if for each fiber there is a vector space isomorphism $f_i \colon \mathrm{L}_i \to \mathrm{L}_i$, such that for each edge ij:

$$\theta_{ij} \circ f_i = f_j \circ \eta_{ij}.$$

Since η_e and θ_e are linear, this boils down to discrete $\mathbb{F}_*$-valued functions and the relation characterizing an isomorphism becomes

$$\theta_{ij} = \big(g_j g_i^{-1}\big)\eta_{ij} = (dg)_{ij}\eta_{ij},$$

i.e. η and θ differ by an exact discrete $\mathbb{F}_*$-valued 1-form. In particular, the difference of two connection forms representing the same connection η is exact.

Thus we obtain a well-defined map sending a discrete line bundle L with connection to the corresponding equivalence class of connection forms

$$[\omega] \in \Omega^1(\mathcal{X}, \mathbb{F}_*)/d\Omega^0(\mathcal{X}, \mathbb{F}_*).$$

Theorem 4. *The map $F\colon \mathcal{L}_{\mathcal{X}}^{\mathbb{F}} \to \Omega^1(\mathcal{X}, \mathbb{F}_*)/d\Omega^0(\mathcal{X}, \mathbb{F}_*)$, $[\mathrm{L}] \mapsto [\omega]$, where ω is a connection form of* L*, is an isomorphism of groups.*

Proof. Clearly, F is well-defined. Let L and $\tilde{\mathrm{L}}$ be two discrete complex line bundle with connections η and θ, respectively. If $\beta \in \mathfrak{C}_{\mathrm{L}}$ and $\tilde{\beta} \in \mathfrak{C}_{\tilde{\mathrm{L}}}$ are trivial, so is $\beta \otimes \tilde{\beta} \in \mathfrak{C}_{\mathrm{L}\otimes\tilde{\mathrm{L}}}$. Hence, with $\eta = \omega\beta$ and $\tilde{\eta} = \tilde{\omega}\tilde{\beta}$, we get

$$F([\mathrm{L} \otimes \tilde{\mathrm{L}}]) = [\omega\tilde{\omega}] = [\omega][\tilde{\omega}] = F([\mathrm{L}])F([\tilde{\mathrm{L}}]).$$

By the preceding discussion, F is injective. Surjectivity is also easily checked. □

Next we will prove that $\Omega^1(\mathcal{X}, \mathbb{F}_*)/d\Omega^0(\mathcal{X}, \mathbb{F}_*)$ is isomorphic to $\mathrm{Hom}(\ker \partial_1, \mathbb{F}_*)$. The isomorphism is given by the identification

$$\Omega^1(\mathcal{X}, \mathbb{F}_*)/d\Omega^0(\mathcal{X}, \mathbb{F}_*) \ni [\omega] \mapsto \omega|_{\ker \partial_1} \in \mathrm{Hom}(\ker \partial_1, \mathbb{F}_*).$$

Clearly, this is a well-defined group homomorphism. We show its bijectivity in two steps. First, the surjectivity is provided by the following general lemma.

Lemma 1. *Let $\mathcal{X}$ be a simplicial complex and $\mathfrak{G}$ be an abelian group. Then the restriction map $\Phi\colon \Omega^k(\mathcal{X}, \mathfrak{G}) \to \mathrm{Hom}(\ker \partial_k, \mathfrak{G})$, $\omega \mapsto \omega|_{\ker \partial_k}$ is surjective.*

Proof. If we choose an orientation for each simplex in $\mathcal{X}$, then ∂_k is given by an integer matrix. Now, there is a unimodular matrix U such that $\partial_k U = (0|H)$ has Hermite normal form. Write $U = (A|B)$, where $\partial_k A = 0$ and $\partial_k B = H$ and let a_i denote the columns of A, i.e. $A = (a_1, \dots, a_\ell)$. Clearly, $a_i \in \ker \partial_k$. Moreover, if $c \in \ker \partial_k$, then $0 = \partial_k c = (0|H)U^{-1}c$. Hence $U^{-1}c = (q, 0)^\top$, $q \in \mathbb{Z}^\ell$, and thus $c = Aq$. Therefore $\{a_i \mid i = 1, \dots, \ell\}$ is a basis of $\ker \partial_k$. Now, let $\mu \in \mathrm{Hom}(\ker \partial_k, \mathbb{Z})$. A homomorphism is completely determined by its values on a basis. We define $\omega = (\mu(a_1), \dots, \mu(a_\ell), 0 \dots, 0)U^{-1}$. Then $\omega \in \Omega^k(\mathcal{X}, \mathbb{Z})$ and $\omega A = (\mu(a_1), \dots, \mu(a_\ell))$. Hence $\Phi(\omega) = \mu$ and Φ is surjective for forms with coefficients in $\mathbb{Z}$. Now, let $\mathfrak{G}$ be an arbitrary abelian group. And $\mu \in \mathrm{Hom}(\ker \partial_k, \mathfrak{G})$. Now, if $a_1, .., a_\ell$ is an arbitrary basis of $\ker \partial_k$, then there are forms $\omega_1, \dots, \omega_\ell \in \Omega^k(\mathcal{X}, \mathbb{Z})$ such that $\omega_i(a_j) = \delta_{ij}$. Since $\mathbb{Z}$ acts on $\mathfrak{G}$, we can multiply ω_i with elements $g \in \mathfrak{G}$ to obtain forms with coefficients in $\mathfrak{G}$. Now, set $\omega = \sum_{i=1}^{\ell} \omega_i \cdot \mu(a_i)$. Then $\omega \in \Omega^k(\mathcal{X}, \mathfrak{G})$ and $\omega(a_i) = \mu(a_i)$ for $i = 1, \dots, \ell$. Thus $\Phi(\omega) = \mu$. Hence Φ is surjective for forms with coefficients in arbitrary abelian groups. □

For $k = 1$ the injectivity is easy to see. If $\omega|_{\ker \partial_1} = 0$, then we define an $\mathbb{F}_*$-valued function f by *integration along paths*: Fix some vertex i. Then

$$f(j) := \int_\gamma \omega := \sum_{e \in \gamma} \omega(e),$$

where γ is some path joining i to j. Since $\omega|_{\ker \partial_1} = 0$, the value $f(j)$ does not depend on the choice of the path γ. One easily checks that $df = \omega$. Together with Lemma 1, this yields the following theorem.

Theorem 5. *The map $F\colon \Omega^1(\mathcal{X}, \mathbb{F}_*)/d\Omega^0(\mathcal{X}, \mathbb{F}_*) \to \mathrm{Hom}(\ker \partial_1, \mathbb{F}_*)$, $[\omega] \mapsto \omega|_{\ker \partial_1}$ is an isomorphism of groups.*

Let us make the relation to Theorem 3 more explicit. Let $\mathrm{L} \to \mathcal{X}$ be a line bundle with connection η, and let ω be a connection form representing η, i.e. $\eta = \omega\beta$ for some trivial base connection β. Now, let $[\gamma] \in \mathcal{LG}(\mathcal{X}, i)$, where $\gamma = e_\ell \cdots e_1$. By linearity and since trivial connections have vanishing monodromy, we obtain

$$\mathfrak{M}([\gamma]) = \eta_{e_\ell} \circ \cdots \circ \eta_{e_1} = \omega_{e_\ell} \cdots \omega_{e_1} \cdot \beta_{e_\ell} \circ \cdots \circ \beta_{e_1} = \omega(\pi_{ab}([\gamma])) \cdot \mathrm{id}|_{\mathrm{L}_i}\,.$$

Hence, by the uniqueness of $[\mathfrak{M}]_{ab}$, we obtain the following theorem that brings everything nicely together.

Theorem 6. *Let $\mathrm{L} \to \mathcal{X}$ be a line bundle with connection η. Let $\mathfrak{M}$ denote its monodromy and let ω be some connection form representing η. Then, with the identifications above,*

$$[\mathfrak{M}]_{ab} = [\omega].$$

7 Curvature - A Discrete Analogue of Weil's Theorem

Let $\mathcal{X}$ be a connected simplicial complex and let $\mathfrak{G}$ denote an abelian group. Since $d^2 = 0$, the exterior derivative descends to a well-defined map defined on $\Omega^k(\mathcal{X}, \mathfrak{G})/d\Omega^{k-1}(\mathcal{X}, \mathfrak{G})$, which again will be denoted by d. Explicitly,

$$d\colon \Omega^k(\mathcal{X}, \mathfrak{G})/d\Omega^{k-1}(\mathcal{X}, \mathfrak{G}) \to \Omega^{k+1}(\mathcal{X}, \mathfrak{G}), \quad [\omega] \mapsto d\omega.$$

Definition 7. *The $\mathbb{F}_*$-curvature of a discrete $\mathbb{F}$-line bundle $\mathrm{L} \to \mathcal{X}$ is the discrete 2-form $\Omega \in \Omega^2(\mathcal{X}, \mathbb{F}_*)$ given by*

$$\Omega = d[\omega],$$

where $[\omega] \in \Omega^1(\mathcal{X}, \mathbb{F}_*)/d\Omega^0(\mathcal{X}, \mathbb{F}_*)$ represents the isomorphism class [L].

Remark 3: *Note that Ω just encodes the parallel transport along the boundary of the oriented 2-simplices of $\mathcal{X}$ - the "local monodromy".*

From the definition it is obvious that the $\mathbb{F}_*$-curvature is invariant under isomorphisms. Thus, given a prescribed 2-form $\Omega \in \Omega^2(\mathcal{X}, \mathbb{F}_*)$, it is a natural question to ask how many non-isomorphic line bundles with curvature Ω exist.

Actually, this questions is answered easily: Suppose $d[\omega] = \Omega = d[\tilde{\omega}]$, then the difference of ω and $\tilde{\omega}$ is closed. Factoring out the exact 1-forms we see that the space of non-isomorphic line bundles with curvature Ω can be parameterized by the first cohomology group $\mathrm{H}^1(\mathcal{X}, \mathbb{F}_*)$. Further, the existence of a line bundle with curvature $\Omega \in \Omega^2(\mathcal{X}, \mathbb{F}_*)$ is clearly equivalent to the exactness of Ω.

But when is a k-form Ω exact? Clearly, it must be closed. Even more, it must vanish on every closed k-chain: If $\Omega = \mathrm{im}\, d$ and S is a closed k-chain, then

$$\Omega(S) = d\omega(S) = \omega(\partial S) = 0.$$

For $k = 1$, as we have seen, this criterion is sufficient to conclude exactness. For $k > 1$ this is not true with coefficients in arbitrary groups.

Example: *Consider a triangulation $\mathcal{X}$ of the real projective plane $\mathbb{R}\mathrm{P}^2$. The zero-chain is the only closed 2-chain and hence each $\mathbb{Z}_2$-valued 2-form vanishes on every closed 2-chain. But $\mathrm{H}^2(\mathcal{X}, \mathbb{Z}_2) = \mathbb{Z}_2$ and hence there exists a non-exact 2-form.*

In the following we will see that this cannot happen for fields of characteristic zero or, more generally, groups that arise as the image of such fields.

Clearly, there is a natural pairing of $\mathbb{Z}$-modules between $\Omega^k(\mathcal{X}, \mathfrak{G})$ and $\mathrm{C}_k(\mathcal{X}, \mathbb{Z})$:

$$\langle ., . \rangle\colon \Omega^k(\mathcal{X}, \mathfrak{G}) \times \mathrm{C}_k(\mathcal{X}, \mathbb{Z}) \to \mathfrak{G}, \quad (\omega, c) \mapsto \omega(c).$$

This pairing is degenerate if and only if $\mathfrak{G}$ is periodic with bounded exponent. In particular, if $\mathfrak{G}$ is a field $\mathbb{F}$ of characteristic zero, $\langle ., . \rangle$ yields a group homomorphism

$$F_k\colon \mathrm{C}_k(\mathcal{X}, \mathbb{Z}) \to \mathrm{Hom}_{\mathbb{F}}(\Omega^k(\mathcal{X}, \mathbb{F}), \mathbb{F}) = (\Omega^k(\mathcal{X}, \mathbb{F}))^*.$$

A basis of $\mathrm{C}_k(\mathcal{X}, \mathbb{Z})$ is mapped under F_k to a basis of $(\Omega^k(\mathcal{X}, \mathbb{F}))^*$ and hence $\mathrm{C}_k(\mathcal{X}, \mathbb{Z})$ appears as n_k-dimensional lattice in $(\Omega^k(\mathcal{X}, \mathbb{F}))^*$.

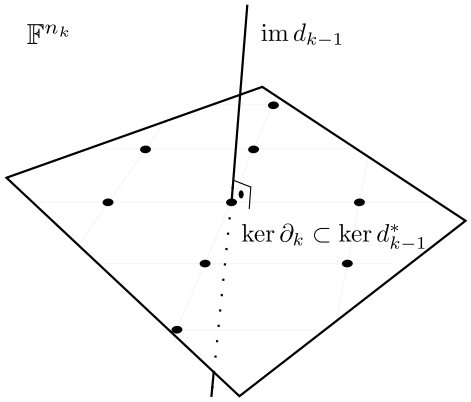

Figure 6. With the identifications 7.1, the space of k-forms becomes a direct sum of the image of d_{k-1} and the kernel of its adjoint d^*_{k-1}, the latter of which contains the closed k-chains as a lattice.

Let d^*_k denote the adjoint of the discrete exterior derivative d_k with respect to the natural pairing between $\Omega^k(\mathcal{X}, \mathbb{F})$ and $(\Omega^k(\mathcal{X}, \mathbb{F}))^*$. Clearly,

$$d^*_k \circ F_k = F_k \circ \partial_{k+1}.$$

Now, since the simplicial complex is finite, we can choose bases of $\mathrm{C}_k(\mathcal{X}, \mathbb{Z})$ for all k. This in turn yields bases of $(\Omega^k(\mathcal{X}, \mathbb{F}))^*$ and hence, by duality, bases of $\Omega^k(\mathcal{X}, \mathbb{F})$. With respect to these bases we have

$$\mathrm{C}_k(\mathcal{X}, \mathbb{Z}) = \mathbb{Z}^{n_k} \subset \mathbb{F}^{n_k} = (\Omega^k(\mathcal{X}, \mathbb{F}))^* = \Omega^k(\mathcal{X}, \mathbb{F}), \tag{7.1}$$

where n_k denotes the number of k-simplices. Moreover, the pairing is represented by the standard product. The operator $d^*_{k-1} = \partial_k$ is then just an integer matrix and

$$\partial_k = d^\top_{k-1}.$$

Clearly, we have $\operatorname{im} d_{k-1} \perp \ker d^*_{k-1}$. And, by the rank-nullity theorem,

$$n_k = \dim \operatorname{im} d^*_{k-1} + \dim \ker d^*_{k-1} = \dim \operatorname{im} d_{k-1} + \dim \ker d^*_{k-1}.$$

Hence, under the identifications above, we have that $\mathbb{F}^{n_k} = \operatorname{im} d_{k-1} \oplus \ker d^*_{k-1}$ (see Figure 6). Moreover, $\ker \partial_k$ contains a basis of $\ker d^*_{k-1}$. From this we conclude immediately the following lemma.

Lemma 2. *Let $\omega \in \Omega^k(\mathcal{X}, \mathbb{F})$, where $\mathbb{F}$ is a field of characteristic zero. Then*

$$\omega \in \operatorname{im} d_{k-1} \iff \langle \omega, c \rangle = 0 \text{ for all } c \in \ker \partial_k.$$

Remark 4: *Note, that for boundary cycles the condition is nothing but the closedness of the form ω. Thus Lemma 2 states that a closed form $\omega \in \Omega^k(\mathcal{X}, \mathbb{F})$ is exact if and only if the integral over all homology classes $[c] \in \mathrm{H}_k(\mathcal{X}, \mathbb{Z})$ vanishes.*

Let $\mathfrak{G}$ be an abelian group. The sequence below will be referred to as the *k-th fundamental sequence of forms with coefficients in $\mathfrak{G}$*:

$$\Omega^{k-1}(\mathcal{X}, \mathfrak{G}) \xrightarrow{d_{k-1}} \Omega^k(\mathcal{X}, \mathfrak{G}) \xrightarrow{\Phi_k} \operatorname{Hom}(\ker \partial_k, \mathfrak{G}) \to 0,$$

where Φ_k denotes the restriction to the kernel of ∂_k, i.e. $\Phi_k(\omega) := \omega|_{\ker \partial_k}$.

Combining Lemma 1 and Lemma 2, we obtain that the fundamental sequence with coefficients in a field $\mathbb{F}$ of characteristic zero is exact for all $k > 1$. This serves as an anchor point. The exactness propagates under surjective group homomorphisms.

Lemma 3. *Let $\mathfrak{A} \xrightarrow{f} \mathfrak{B} \to 0$ be a an exact sequence. Then, if the k-th fundamental sequence of forms is exact with coefficients in $\mathfrak{A}$, so it is with coefficients in $\mathfrak{B}$.*

Proof. By Lemma 1 the restriction map Φ_k is surjective for every abelian group. It is left to check that $\ker \Phi_k = \operatorname{im} d_{k-1}$ with coefficients in $\mathfrak{B}$. Let $\Omega \in \Omega^k(\mathcal{X}, \mathfrak{B})$ such that $\Phi_k(\Omega) = 0$.

Since $f\colon \mathfrak{A} \to \mathfrak{B}$ is surjective, there is a form $\Xi \in \Omega^k(\mathcal{X}, \mathfrak{A})$ such that $\Omega = f \circ \Xi$. Since $0 = \Phi_k(\Omega) = f \circ \Phi_k(\Xi)$, we obtain that $\Phi_k(\Xi)$ takes its values in $\ker f$. Since Φ_k is surjective for arbitrary groups, there is $\Theta \in \Omega^k(\mathcal{X}, \ker f)$ such that $\Phi_k(\Xi) = \Phi_k(\Theta)$. Hence $\Phi_k(\Xi - \Theta) = 0$. Thus there is a form $\xi \in \Omega^{k-1}(\mathcal{X}, \mathfrak{A})$ such that $d_{k-1}\xi = \Xi - \Theta$. Now, let $\omega := f \circ \xi \in \Omega^{k-1}(\mathcal{X}, \mathfrak{B})$. Then

$$d_{k-1}\omega = d_{k-1} f \circ \xi = f \circ d_{k-1}\xi = f \circ (\Xi - \Theta) = f \circ \Xi = \Omega.$$

Hence $\ker \Phi_k = \operatorname{im} d_{k-1}$ and the sequence (with coefficients in $\mathfrak{B}$) is exact. □

Remark 5: *The map $f\colon \mathbb{C} \to \mathbb{C}$, $z \mapsto \exp(2\pi i\, z)$ provides a surjective group homomorphism from $\mathbb{C}$ onto $\mathbb{C}_*$, and similarly from $\mathbb{R}$ onto $\mathbb{S}$. Hence the k-th fundamental sequence of forms is exact for coefficients in $\mathbb{C}_*$ and in the unit circle $\mathbb{S}$.*

Remark 6: *The k-th fundamental sequence with coefficients in an abelian group $\mathfrak{G}$ is exact if and only if $\Omega^k(\mathcal{X}, \mathfrak{G})/d\Omega^{k-1}(\mathcal{X}, \mathfrak{G}) \cong \operatorname{Hom}(\ker \partial_k, \mathfrak{G})$. The isomorphism is just induced by the restriction map Φ_k.*

The following corollary is just an easy consequence of the Remark 5. It nicely displays the fibration of the complex line bundles by their $\mathbb{C}_*$-curvature.

Corollary 1. *For $\mathfrak{G} = \mathbb{S}, \mathbb{C}_*$ the following sequence is exact:*

$$1 \to \mathrm{H}^1(\mathcal{X}, \mathfrak{G}) \hookrightarrow \Omega^1(\mathcal{X}, \mathfrak{G})/d\Omega^0(\mathcal{X}, \mathfrak{G}) \xrightarrow{d} \Omega^2(\mathcal{X}, \mathfrak{G}) \to \operatorname{Hom}(\ker \partial_2, \mathfrak{G}) \to 1.$$

Definition 8. *Let $\Omega^* \in \Omega^k(\mathcal{X}, \mathbb{S})$. A real-valued form $\Omega \in \Omega^2(\mathcal{X}, \mathbb{R})$ is called compatible with Ω^* if $\Omega^* = \exp\big(\imath\Omega\big)$. A discrete hermitian line bundle with curvature is a discrete hermitian line bundle* L *with connection equipped with a closed 2-form compatible with the $\mathbb{S}$-curvature of* L.

For real-valued forms it is common to denote the natural pairing with the k-chains by an integral sign, i.e. if $\omega \in \Omega^k(\mathcal{X}, \mathbb{R})$ and $c \in \mathrm{C}_k(\mathcal{X}, \mathbb{Z})$, then

$$\int_c \omega := \langle \omega, c \rangle = \omega(c).$$

Theorem 7. *Let* L *be a discrete hermitian line bundle with curvature Ω. Then Ω is integral, i.e.*

$$\int_C \Omega \in 2\pi\,\mathbb{Z}, \qquad \text{for all } C \in \ker \partial_2.$$

Proof. By definition the curvature form Ω satisfies $\exp\big(i\Omega\big) = d\omega$ for some connection form $\omega \in \Omega^1(\mathcal{X}, \mathbb{S})$. Thus, if $C \in \ker \partial_2$,

$$\exp\big(\imath \int_{\mathcal{X}} \Omega\big) = \langle \exp(i\Omega), \mathcal{X} \rangle = \langle d\omega, \mathcal{X} \rangle = \langle \omega, \partial\mathcal{X} \rangle = 1.$$

This proves the claim. □

Conversely, Corollary 1 yields a discrete version of a theorem of André Weil (see **[71]** or **[36, 65]**), which plays a prominent role in the process of prequantization.

Theorem 8. *If $\Omega \in \Omega^2(\mathcal{X}, \mathbb{R})$ is integral, then there exists a hermitian line bundle with curvature Ω.*

Proof. Consider $\Omega^* := \exp(i\Omega)$. Since Ω is integral, $\langle \Omega^*, c \rangle = 1$ for all $c \in \ker \partial_2$. Thus, by Corollary 1, there exists $r \in \Omega^1(\mathcal{X}, \mathbb{S})$ such that $dr = \Omega^*$, which in turn defines a hermitian line bundle with curvature Ω. □

Remark 7: *Moreover Corollary 1 shows that the connections of two such bundles differ by an element of $\mathrm{H}^1(\mathcal{X}, \mathbb{S})$. Thus the space of discrete hermitian line bundles with fixed curvature Ω can be parameterized by $\mathrm{H}^1(\mathcal{X}, \mathbb{S})$.*

8 The Index Formula for Hermitian Line Bundles

Before we define the degree of a discrete hermitian line bundle with curvature or the index form of a section, let us first recall the situation in the smooth setting again. Therefore, let $\mathrm{L} \to \mathrm{M}$ be a smooth hermitian line bundle with connection. Since the curvature tensor R^∇ of ∇ is a 2-form taking values in the skew-symmetric endomorphisms of L, it boils down to a closed real-valued 2-form $\Omega \in \Omega^2(\mathrm{M}, \mathbb{R})$,

$$R^\nabla = -\imath\Omega.$$

The following theorem shows there is an interesting relation between the index sum of a section $\psi \in \Gamma(\mathrm{L})$, the curvature 2-form Ω, and the *rotation form* ξ^ψ *of* ψ:

$$\xi^\psi := \frac{\langle \nabla\psi, \imath\psi\rangle}{\langle \psi, \psi\rangle}.$$

Theorem 9. *Let* $\mathrm{L} \to \mathrm{M}$ *be a smooth hermitian line bundle with connection, let* Ω *be its curvature* 2*-form, and let* $\psi \in \Gamma(\mathrm{L})$ *be a section with a discrete zero set* Z*. If* C *is a finite smooth* 2*-chain such that* $\partial C \cap Z = \emptyset$*, then*

$$2\pi \sum_{p \in C \cap Z} \mathrm{ind}_p^\psi = \int_{\partial C} \xi^\psi + \int_C \Omega.$$

Proof. We can assume that C is a single smooth triangle. Then we can express ψ on C in terms of a complex-valued function z and a pointwise-normalized local section ϕ, i.e. $\psi = z\,\phi$. Since $\mathrm{Im}(\frac{dz}{z}) = d\arg(z)$, we obtain

$$\xi^\psi = \frac{1}{|z|^2}\langle dz\,\phi + z\,\nabla\phi, \imath z\,\phi\rangle = \langle \frac{dz}{z}\,\phi, \imath\phi\rangle + \langle \nabla\phi, \imath\phi\rangle = d\arg(z) + \langle \nabla\phi, \imath\phi\rangle.$$

Moreover, away from zeros, we have

$$d\langle \nabla\phi, \imath\phi\rangle = \langle R^\nabla\phi, \imath\phi\rangle + \langle \nabla\phi \wedge \imath\nabla\phi\rangle = \langle R^\nabla\phi, \imath\phi\rangle = -\Omega.$$

Hence, altogether, we obtain

$$\int_{\partial C} \xi^\psi = \int_{\partial C} d\arg(z) + \int_{\partial C} \langle \nabla\phi, \imath\phi\rangle = 2\pi \sum_{p \in C \cap Z} \mathrm{index}_p(\psi) - \int_C \Omega.$$

This proves the claim. □

Actually, in the case that L is a hermitian line bundle with connection over a closed oriented surface M, then Theorem 9 tells us that $\int_\mathrm{M} \Omega \in 2\pi\mathbb{Z}$, which yields a well-known topological invariant - the *degree of* L:

$$\deg\big(\mathrm{L}\big) := \frac{1}{2\pi}\int_\mathrm{M} \Omega.$$

From Theorem 9 we immediately obtain the famous Poincaré-Hopf index theorem.

Theorem 10. *Let* $\mathrm{L} \to \mathrm{M}$ *be a smooth hermitian line bundle over a closed oriented surface. Then, if* $\psi \in \Gamma(\mathrm{L})$ *is a section with isolated zeros,*

$$\deg\big(\mathrm{L}\big) = \sum_{p \in \mathrm{M}} \mathrm{ind}_p^\psi.$$

Now, let us consider the discrete case. Let $\mathrm{L} \to \mathcal{X}$ be a discrete hermitian line bundle with curvature Ω and let $\psi \in \Gamma(\mathrm{L})$ be a discrete nowhere-vanishing section such that

$$\eta_{ij}(\psi_i) \neq -\psi_j \tag{8.1}$$

for each edge ij of $\mathcal{X}$. Here η denotes the connection of L as usual. The *rotation form* ξ^ψ *of* ψ is then defined as follows:

$$\xi_{ij}^\psi := \arg\big(\frac{\psi_j}{\eta_{ij}(\psi_i)}\big) \in (-\pi, \pi).$$

Remark 8: *Equation (8.1) can be interpreted as the condition that no zero lies in the* 1*-skeleton of* $\mathcal{X}$ *(compare Section 11). Actually, by a consistent choice of the argument on each oriented edge, we can drop this condition. Figuratively speaking, if a section has a zero in the* 1*-skeleton, then we decide whether we push it to the left or the right face of the edge.*

This defined, we can use Theorem 9 to define the *index form* of a discrete section.

Definition 9. *Let $\mathrm{L} \to \mathcal{X}$ be a discrete hermitian line bundle with curvature Ω. For $\psi \in \Gamma(\mathrm{L})$, we define the index form of ψ by*

$$\mathrm{ind}^{\psi} := \frac{1}{2\pi}\big(d\xi^{\psi} + \Omega\big).$$

Theorem 11. *The index form of a nowhere-vanishing discrete section is $\mathbb{Z}$-valued.*

Proof. Let L be a discrete hermitian line bundle with curvature and let η be its connection. Let $\psi \in \Gamma(\mathrm{L})$ be a nowhere-vanishing section. Now, choose a connection form ω, i.e. $\eta = \omega\beta$, where β is a trivial connection on L. Then we can write ψ with respect to a non-vanishing parallel section ϕ of β, i.e. there is a $\mathbb{C}$-valued function z such that $\psi = z\phi$. Then $\xi^{\psi}_{ij} = \arg\big(\frac{z_j}{\omega_{ij} z_i}\big)$ and thus

$$\exp\big(2\pi\imath\, d\xi^{\psi}_{ijk}\big) = \exp\big(\imath\, \arg\big(\frac{z_i}{\omega_{ki} z_k}\big) + \imath\, \arg\big(\frac{z_j}{\omega_{ij} z_i}\big) + \imath\, \arg\big(\frac{z_k}{\omega_{jk} z_j}\big)\big) = \frac{1}{d\omega_{ijk}}.$$

Thus

$$\exp\big(2\pi\imath\, \mathrm{ind}^{\psi}_{ijk}\big) = \frac{\exp\big(\imath\Omega_{ijk}\big)}{d\omega_{ijk}} = 1.$$

This proves the claim. □

If L is a discrete hermitian line bundle with curvature Ω over a closed oriented surface $\mathcal{X}$, then we can define the *degree of* L just as in the smooth case:

$$\deg\big(\mathrm{L}\big) := \frac{1}{2\pi}\int_{\mathcal{X}} \Omega.$$

Here we have identified $\mathcal{X}$ by the corresponding closed 2-chain. From Theorem 7 we immediately obtain the following corollary.

Corollary 2. *The degree of a discrete hermitian line bundle with curvature is an integer:*

$$\deg\big(\mathrm{L}\big) \in \mathbb{Z}.$$

The discrete Poincaré-Hopf index theorem follows easily from the definitions.

Theorem 12. *Let $\mathrm{L} \to \mathcal{X}$ be a discrete hermitian line bundle with curvature Ω over an oriented simplicial surface. If $\psi \in \Gamma(\mathrm{L})$ is a non-vanishing discrete section, then*

$$\deg\big(\mathrm{L}\big) = \sum_{ijk \in \mathcal{X}} \mathrm{ind}^{\psi}_{ijk}.$$

Proof. Since the integral of an exact form over a closed oriented surface vanishs,

$$2\pi \deg\big(\mathrm{L}\big) = \int_{\mathcal{X}} \Omega = \int_{\mathcal{X}} d\xi^{\psi} + \Omega = 2\pi \sum_{ijk \in \mathcal{X}} \mathrm{ind}^{\psi}_{ijk},$$

as was claimed. □

9 Piecewise-Smooth Vector Bundles over Simplicial Complexes

It is well-known that each abstract simplicial complex $\mathcal{X}$ has a geometric realization which is unique up to simplicial isomorphisms. In particular, each abstract simplex is then realized as an affine simplex and hence carries the structure of a manifold with corners. Moreover, each face σ' of a simplex $\sigma \in \mathcal{X}$ comes with an affine embedding

$$\iota_{\sigma'\sigma} \colon \sigma' \hookrightarrow \sigma.$$

Here we use the notion of manifold with corners as presented in **[38]**.

Remark 9: *This actually turns $\mathcal{X}$ into a 'stratified space' in the sense that it is patched together from smooth spaces. There are various notions of stratified spaces all of which are adapted to certain needs - but not to ours, as these spaces come usually with a lot of differential geometric invariants. A quite comprehensive overview is given in e.g.* **[55]**.

In the following, we won't distinguish between the abstract simplicial complex and its geometric realization.

Definition 10. *A piecewise-smooth vector bundle* E *over a simplicial complex* $\mathcal{X}$ *is a topological vector bundle* $\pi\colon \mathrm{E} \to \mathcal{X}$ *such that*

a) for each $\sigma \in \mathcal{X}$ *the restriction* $\mathrm{E}_\sigma := \mathrm{E}|_\sigma$ *is a smooth vector bundle over* σ*,*
b) for each face σ' *of* $\sigma \in \mathcal{X}$*, the inclusion* $\mathrm{E}_{\sigma'} \hookrightarrow \mathrm{E}_\sigma$ *is a smooth embedding.*

Clearly, $\mathcal{X}$ has no tangent bundle. Nonetheless, differential forms survive as collections of smooth differential forms defined on the simplices which are compatible in the sense that they agree on common faces.

Definition 11. *Let* E *be a piecewise-smooth vector bundle over* $\mathcal{X}$*. An* E*-valued differential* k*-form is a collection* $\omega = \{\omega_\sigma \in \Omega^k(\sigma, \mathrm{E}_\sigma)\}_{\sigma\in\mathcal{X}}$ *such that for each face* σ' *of a simplex* $\sigma \in \mathcal{X}$ *the following relation holds:*

$$\iota_{\sigma'\sigma}^* \omega_\sigma = \omega_{\sigma'},$$

where $\iota_{\sigma'\sigma}\colon \sigma' \hookrightarrow \sigma$ *denotes the inclusion. The space of* E*-valued differential* k*-forms is denoted by* $\Omega^k_{ps}(\mathcal{X}, \mathrm{E})$*.*

Remark 10: *Note that a* 0*-form defines a continuous map on the simplicial complex. Hence the definition actually includes the definition of functions and sections in general: A smooth section of* E *is a continuous section* $\psi\colon \mathcal{X} \to \mathrm{E}$ *such that for each simplex* $\sigma \in \mathcal{X}$ *the restriction* $\psi_\sigma := \psi|_\sigma : \sigma \to \mathrm{E}_\sigma$ *is smooth, i.e.*

$$\Gamma_{ps}(\mathrm{E}) := \big\{\psi\colon \mathcal{X} \to \mathrm{E} \mid \psi_\sigma \in \Gamma(\mathrm{E}_\sigma) \text{ for all } \sigma \in \mathcal{X}\big\}.$$

Since the pullback commutes with the wedge-product $\wedge$ and the exterior derivative d of real-valued forms we can define the wedge product and the exterior derivative of piecewise-smooth differential forms by applying it componentwise.

Definition 12. *For* $\omega = \{\omega_\sigma\}_{\sigma\in\mathcal{X}} \in \Omega^k_{ps}(\mathcal{X}, \mathbb{R})$*,* $\eta = \{\eta_\sigma\}_{\sigma\in\mathcal{X}} \in \Omega^\ell_{ps}(\mathcal{X}, \mathbb{R})$*,*

$$\omega \wedge \eta := \{\omega_\sigma \wedge \eta_\sigma\}_{\sigma\in\mathcal{X}}, \quad d\omega := \{d\omega_\sigma\}_{\sigma\in\mathcal{X}}.$$

One easily verifies that all the properties of $\wedge$ and d carry over directly to the piecewise-smooth case.

Definition 13. *A connection on a piecewise-smooth vector bundle* E *over* $\mathcal{X}$ *is a linear map* $\nabla\colon \Gamma_{ps}(\mathrm{E}) \to \Omega^1_{ps}(\mathcal{X}, \mathrm{E})$ *such that*

$$\nabla(f\psi) = df\,\psi + f\,\nabla\psi, \quad \text{for all } f \in \Omega^0_{ps}(\mathcal{X}, \mathbb{R}),\ \psi \in \Gamma_{ps}(\mathrm{E}).$$

Once we have a connection on a smooth vector bundle we obtain a corresponding exterior derivative d^∇ on E-valued forms.

Theorem 13. *Let* E *be a piecewise-smooth vector bundle over* $\mathcal{X}$*. Then there is a unique linear map* $d^\nabla\colon \Omega^k_{ps}(\mathcal{X}, \mathrm{E}) \to \Omega^{k+1}_{ps}(\mathcal{X}, \mathrm{E})$ *such that* $d^\nabla\psi = \nabla\psi$ *for all* $\psi \in \Gamma_{ps}(\mathrm{E})$*, and*

$$d^\nabla(\omega \wedge \eta) = d\omega \wedge \eta + (-1)^k \omega \wedge d^\nabla\eta$$

for all $\omega \in \Omega^k_{ps}(\mathcal{X}, \mathbb{R})$ *and* $\eta \in \Omega^\ell_{ps}(\mathcal{X}, \mathrm{E})$*.*

The curvature tensor survives as a piecewise-smooth End(E)-valued 2-form.

Definition 14. *Let* $\mathrm{E} \to \mathcal{X}$ *be a piecewise-smooth vector bundle. The endomorphism-valued curvature* 2*-form of a connection* ∇ *on* E *is defined as follows:*

$$d^\nabla \circ d^\nabla \in \Omega^2_{ps}(\mathcal{X}, \mathrm{End}(\mathrm{E})).$$

10 The Associated Piecewise-Smooth Hermitian Line Bundle

Let $\tilde{\mathrm{L}} \to \mathcal{X}$ be a piecewise-smooth hermitian line bundle with connection ∇ over a simplicial complex. Just as in the smooth case the endomorphism-valued curvature 2-form takes values in the skew-adjoint endomorphisms and hence is given by a piecewise-smooth real-valued 2-form $\tilde{\Omega}$:

$$d^\nabla \circ d^\nabla = -\imath\tilde{\Omega}.$$

Since each simplex of $\mathcal{X}$ has an affine structure, we can speak of constant forms.

The goal of this section will be to construct for each discrete hermitian line bundle with curvature a piecewise-smooth hermitian line bundle with constant curvature which in a certain sense naturally contains the discrete bundle. Therefore we first prove two preparing lemmata.

Lemma 4. *To each closed discrete real-valued k-form ω there corresponds a unique constant piecewise-smooth k-form $\tilde{\omega}$ such that*

$$\omega(c) = \int_c \tilde{\omega}, \qquad \text{for all } c \in \mathrm{C}_k(\mathcal{X}, \mathbb{Z}).$$

The form $\tilde{\omega}$ will be called the piecewise-smooth form associated to ω.

Proof. Clearly, it is enough to consider just a single n-simplex σ. We denote the space of constant piecewise-smooth k-forms on σ by Ω^k_c and the space of discrete k-forms on σ by Ω^k_d. Consider the linear map $F\colon \Omega^k_c \to \Omega^k_d$ that assigns to $\tilde{\omega} \in \Omega^k_c$ the discrete k-form given by

$$F(\tilde{\omega})_{\sigma'} := \int_{\sigma'} \tilde{\omega}.$$

Clearly, F is injective. Moreover, since each constant piecewise-smooth form is closed, we have that $\operatorname{im} F \subset \ker d_k$, where d_k denotes the discrete exterior derivative. Hence it is enough to show that the space of closed discrete k-forms on σ is of dimension $\binom{n}{k}$. This we can do by induction. Clearly, $\dim\ker d_0 = 1 = \binom{n}{0}$. Now, suppose that $\dim\ker d_{k-1} = \binom{n}{k-1}$. By Lemma 2, we have $\ker d_k = \operatorname{im} d_{k-1}$. Hence,

$$\dim\ker d_k = \dim\operatorname{im} d_{k-1} = \dim\Omega^k_d - \dim\ker d_{k-1} = \tbinom{n+1}{k} - \tbinom{n}{k-1} = \tbinom{n}{k}.$$

Hence for each closed discrete k-form we obtain a unique constant piecewise-smooth k-form which has the desired integrals on the k-simplices. □

It is a classical result that on star-shaped domains $U \subset \mathbb{R}^N$ each closed form is exact, i.e. if $\Omega \in \Omega^k(U, \mathbb{R})$ is closed, then there exists a form $\omega \in \Omega^{k-1}(U, \mathbb{R})$ such that $\Omega = d\omega$. Moreover, the potential can be constructed explicitly by the map $K\colon \Omega^k(U, \mathbb{R}) \to \Omega^{k-1}(U, \mathbb{R})$ given by

$$K(\Omega) = \sum_{i_1<\dots<i_k} \sum_{\alpha=1}^{k} (-1)^{\alpha-1} \Big(\int_0^1 t^{k-1} \Omega_{i_1\dots i_k}(tx) dt \Big)\, x_{i_\alpha}\, dx_{i_1} \wedge \dots \wedge \widehat{dx_{i_\alpha}} \wedge \dots \wedge dx_{i_k},$$

where $\Omega = \sum_{i_1<\dots<i_k} \Omega_{i_1\dots i_k}\, dx_{i_1} \wedge \dots \wedge dx_{i_k}$. One directly checks that

$$K(d\Omega) + d\,K(\Omega) = \Omega.$$

Hence, if $d\Omega = 0$, we get $\Omega = d\,K(\Omega)$. Clearly, the same construction works for piecewise-smooth forms defined on the star of a simplex, which yields the following piecewise-smooth version of the Poincaré-Lemma.

Lemma 5. *On the star of a simplex each closed piecewise-smooth form is exact.*

This at hand we are ready to prove the main result of this section.

Theorem 14. *Let $\mathrm{L} \to \mathcal{X}$ be a discrete hermitian line bundle with curvature Ω over a simplicial complex and let $\tilde{\Omega}$ be the piecewise-smooth constant 2-form associated to Ω. Then there is a piecewise-smooth hermitian line bundle $\tilde{\mathrm{L}} \to \mathcal{X}$ with connection $\tilde{\nabla}$ of curvature $\tilde{\Omega}$, such that $\tilde{\mathrm{L}}_i = \mathrm{L}_i$ for each vertex i and the parallel transports coincide along each edge path. The bundle $\tilde{\mathrm{L}}$ is unique up to isomorphism.*

Proof. First we construct the piecewise-smooth hermitian line bundle. Let $\mathrm{L} \to \mathcal{X}$ be a discrete hermitian line bundle with curvature Ω and let η denote its connection. Let $\mathcal{V}$ be the vertex set of $\mathcal{X}$ and let S_i denote the open vertex star of the vertex i. Further, since Ω is closed, by Lemma 4, there is a piecewise-smooth constant form $\tilde{\Omega}$ associated to Ω. Now, consider the set

$$\hat{\mathrm{L}} := \bigsqcup_{i\in\mathcal{V}} S_i \times \mathrm{L}_i.$$

Note, that $S_i \cap S_j \neq \emptyset$ if and only if ij is an edge of $\mathcal{X}$ or $i = j$. Thus, if we set $\eta_{ii} := \mathrm{id}|_{\mathrm{L}_i}$, we can define an equivalence relation on $\hat{\mathrm{L}}$ as follows:

$$(i,p,u) \sim (j,q,v) :\Longleftrightarrow p = q \text{ and } v = \exp\big(-\imath \int_{\Delta^p_{ij}} \tilde{\Omega}\big)\eta_{ij}(u),$$

where Δ^p_{ij} denotes the oriented triangle spanned by the point i, j and p. Note here that Δ^p_{ij} is completely contained in some simplex of $\mathcal{X}$. Let us check shortly that this really defines an equivalence relation. Here the only non-trivial property is transitivity. Therefore, let $(i,p,u) \sim (j,q,v)$ and $(j,q,v) \sim (k,r,w)$. Thus we have $p = q = r$ and p lies in a simplex which contains the oriented triangle ijk. Clearly, the 2-chain $\Delta^p_{ij} + \Delta^p_{jk} + \Delta^p_{ki}$ is homologous to ijk and since constant forms are closed we get

$$\int_{\Delta^p_{ij}+\Delta^p_{jk}} \tilde{\Omega} = -\int_{\Delta^p_{ki}} \tilde{\Omega} + \int_{ijk} \tilde{\Omega} = \int_{\Delta^p_{ik}} \tilde{\Omega} + \Omega_{ijk}.$$

Hence we obtain

$$\begin{aligned} w &= \exp\big(-\imath \int_{\Delta^p_{jk}} \tilde{\Omega}\big)\eta_{jk}\Big(\exp\big(-\imath \int_{\Delta^p_{ij}} \tilde{\Omega}\big)\eta_{ij}(u)\Big) \\ &= \exp\big(-\imath \int_{\Delta^p_{ij}+\Delta^p_{jk}} \tilde{\Omega}\big)\eta_{jk} \circ \eta_{ij}(u) \\ &= \exp\big(-\imath \int_{\Delta^p_{ik}} \tilde{\Omega} - \imath\Omega_{ijk}\big)\eta_{jk} \circ \eta_{ij}(u) \\ &= \exp\big(-\imath \int_{\Delta^p_{ik}} \tilde{\Omega}\big)\eta_{ik}(u), \end{aligned}$$

and thus $(i,p,u) \sim (k,r,w)$. Hence $\sim$ defines an equivalence relation and one easily checks that the quotient $\tilde{\mathrm{L}} := \hat{\mathrm{L}}/_\sim$ is a piecewise-smooth line bundle over $\mathcal{X}$. The local trivializations are then basically given by the inclusions $S_i \times \mathrm{L}_i \hookrightarrow \tilde{\mathrm{L}}$ sending a point to the corresponding equivalence class. Moreover, all transition maps are unitary so that the hermitian metric of L extends to $\tilde{\mathrm{L}}$ and turns $\tilde{\mathrm{L}}$ into a hermitian line bundle. Clearly, $\tilde{\mathrm{L}}\Big|_{\mathcal{V}} = \mathrm{L}$.

Next, we need to construct the connection. Therefore we will use an explicit system of local sections: Choose for each vertex $i \in \mathcal{V}$ a unit vector $X_i \in \mathrm{L}_i$ and define $\phi_i(p) := [i,p,X_i]$. This yields for each vertex i a piecewise-smooth section ϕ_i define on the star S_i. For each non-empty intersection $S_i \cap S_j \neq \emptyset$ we then obtain a function $g_{ij} \colon S_i \cap S_j \to \mathbb{S}$. By the above construction, we find that, if $\eta_{ij}(X_i) = r_{ij}X_j$,

$$g_{ij}(p) = r_{ij} \exp\big(-\imath \int_{\Delta^p_{ij}} \tilde{\Omega}\big). \tag{10.1}$$

Since $\tilde{\Omega}$ is closed, Lemma 5 tells us that $\tilde{\Omega}|_{S_i}$ is exact. Hence there is a piecewise-smooth 1-form ω_i defined on S_i such that $d\omega_i = \tilde{\Omega}|_{S_i}$. In general, the form ω_i is only unique up to addition of an exact 1-form, but among those there is a unique form ω_i which is zero along the radial directions originating from i. To see this, just choose some potential $\tilde{\omega}_i$ of $\Omega|_{S_i}$ and define a function $f \colon S_i \to \mathbb{R}$ as follows:

For $p \in S_i$, let $f(p) := \int_{\gamma_i^p} \tilde{\omega}_i$, where γ_i^p denote the linear path from the vertex i to the point p. Then $\omega_i := \tilde{\omega}_i - df$ is a piecewise-smooth potential of $\Omega|_{S_i}$ and vanishes on radial directions. For the uniqueness, let $\hat{\omega}_i$ be another such potential. Then, the difference $\omega_i - \hat{\omega}_i$ is closed and hence exact on S_i, i.e. there is $f \colon S_i \to \mathbb{R}$ such that $df = \omega_i - \hat{\omega}_i$. Since df vanishes on radial directions f is constant on radial lines starting at i and hence constant on S_i. Thus $\omega_i = \hat{\omega}_i$.

Suppose that for each edge ij the forms ω_i and ω_j are *compatible*, i.e., wherever both are defined,

$$\imath\omega_j = \imath\omega_i + d\log g_{ij}.$$

Then we can define a connection ∇ as follows: Let $\psi \in \Gamma\big(\tilde{\mathrm{L}}\big)$ and let $X \in \mathrm{T}_p\sigma$ for some simplex σ of $\mathcal{X}$, then there is some $S_i \ni p$. On S_i we can express ψ with respect to ϕ_i, i.e. $\psi = z\,\phi_i$ for

some piecewise-smooth function $z\colon S_i \to \mathbb{C}$. Then we define

$$\nabla_X \psi := \big(dz(X) - \imath\omega_i(X) z\big)\phi_i. \tag{10.2}$$

In general there are several stars that contain the point p. From compatibility easily follows that the definition does not depend on the choice of the vertex. Hence we have constructed a piecewise smooth connection ∇. One easily checks that ∇ is unitary and since $d\omega_i = \tilde{\Omega}|_{S_i}$ we get $d^\nabla \circ d^\nabla = -\imath\tilde{\Omega}$ as desired.

So it is left to check the compatibility of the forms ω_{ij} constructed above. Let ij be some edge and let p_0 be a point in its interior. Since $\omega_i - \omega_j$ is closed, we can define $\varphi\colon S_i \cap S_j \to \mathbb{R}$ by $\varphi(p) := \int_{\gamma_p} \omega_i - \omega_j$, where γ_p is some path in $S_i \cap S_j$ from the point p_0 to the point p. Then, for $p \in S_i \cap S_j$,

$$\int_{\Delta_p} \Omega = \int_{\partial\Delta_p} \omega_j = \int_{ij+\gamma_j^p-\gamma_i^p} \omega_j = -\int_{\gamma_i^p} \omega_j = \int_{\gamma_i^p} \omega_i - \omega_j = \varphi(p),$$

where as above γ_i^p denotes the linear path from i to p and, similarly, γ_j^p denotes the linear path from j to the point p. From this we obtain

$$\omega_i - \omega_j = d\varphi = d\int_{\Delta_p} \Omega$$

and in particular $\imath\omega_j = \imath\omega_i + d\log g_{ij}$. This shows the existence.

Now suppose there are two such piecewise-smooth bundles $\tilde{\mathrm{L}}$ and $\hat{\mathrm{L}}$ with connection $\tilde{\nabla}$ and $\hat{\nabla}$, respectively. We want to construct an isomorphism between $\tilde{\mathrm{L}}$ and $\hat{\mathrm{L}}$. Therefore we again use local systems. Explicitly, we choose a discrete direction field $X \in \mathrm{L}$. This yields for each vertex i a vector $X_i \in \tilde{\mathrm{L}}_i = \hat{\mathrm{L}}_i$ which extends by parallel transport along rays starting at i to a local sections $\tilde{\phi}_i$ of $\tilde{\mathrm{L}}$ and, similarly, to a local section $\hat{\phi}_i$ of $\hat{\mathrm{L}}$ defined on S_i.

Now we define $F\colon \tilde{\mathrm{L}} \to \hat{\mathrm{L}}$ to be unique map which is linear on the fibers and satisfies $F(\tilde{\phi}_i) = \hat{\phi}_i$ on S_i. To see that F is well-defined, we need to check that it is compatible with the transition maps. But by construction both systems have equal transition maps, namely the the functions g_{ij} from Equation (10.1) with r_{ij} given by $\eta_{ij}(X_i) = r_{ij}X_j$. Now, if $z_i\,\tilde{\phi}_i = z_j\,\tilde{\phi}_j$, then $z_i = z_j\,g_{ij}$ and hence

$$F(z_i\tilde{\phi}_i) = z_i\,\hat{\phi}_i = z_i\,g_{ij}\hat{\phi}_j = z_j\,\hat{\phi}_j = F(z_j\,\tilde{\phi}_j).$$

Using Equation (10.2) one similarly shows that $F \circ \tilde{\nabla} = \hat{\nabla} \circ F$. Thus $\tilde{\mathrm{L}} \cong \hat{\mathrm{L}}$. □

11 Finite Elements for Hermitian Line Bundles With Curvature

In this section we want to present a specific finite element space on the associated piecewise-smooth hermitian line bundle of a discrete hermitian line with curvature. They are cooked up from the local systems that played such a prominent role in the proof of Theorem 14 and the usual piecewise-linear hat function.

Let $\tilde{\mathrm{L}}$ be the associated piecewise-smooth bundle of a discrete hermitian line bundle $\mathrm{L} \to \mathcal{X}$ and let $x_i\colon \mathcal{X} \to \mathbb{R}$ denote the barycentric coordinate of the vertex $i \in \mathcal{V}$, i.e. the unique piecewise-linear function such that $x_i(j) = \delta_{ij}$, where δ is the Kronecker delta. Clearly,

$$\Gamma(\mathrm{L}) = \bigoplus_{i\in\mathcal{V}} \mathrm{L}_i.$$

To each $X \in \mathrm{L}_i$ we now construct a piecewise-smooth section $\tilde{\psi}$ as follows: First, we extend X to the vertex star S_i of the vertex i using the parallel transport along rays starting at i. To get a global section $\tilde{\psi} \in \Gamma_{ps}(\mathrm{L})$ we use x_i to scale $\tilde{\phi}$ down to zero on ∂S_i and extend it by zero to $\mathcal{X}$, i.e.

$$\tilde{\psi}_p := \begin{cases} x_i(p)\tilde{\phi}_p & \text{for } p \in S_i, \\ 0 & \text{else.} \end{cases}$$

One easily checks that the above construction yields a linear map $\iota\colon \Gamma(\mathrm{L}) \to \Gamma_{ps}(\tilde{\mathrm{L}})$. Clearly, ι is injective - a left-inverse is just given by the restriction map

$$\Gamma_{ps}(\tilde{\mathrm{L}}) \ni \tilde{\psi} \mapsto \tilde{\psi}\big|_{\mathcal{V}} \in \Gamma(\mathrm{L}).$$

Definition 15. *The space of piecewise-linear sections is given by* $\Gamma_{pl}(\tilde{\mathrm{L}}) := \operatorname{im} \iota$.

Thus we identified each section of a discrete hermitian line bundle with curvature with a piecewise-linear section of the associated piecewise-smooth bundle. This allows to define a discrete hermitian inner product and a discrete Dirichlet energy on $\Gamma(\mathrm{L})$, which will finally lead to a generalization of the well-known cotangent Laplace operator for discrete functions on triangulated surfaces. Before we come to the Dirichlet energy, we define Euclidean simplicial complexes.

Similarly to piecewise-smooth form we can define piecewise-smooth (kontravariant) k-tensors as collections of compatible k-tensors: A *piecewise-smooth k-tensor* is a collection $T = \{T_\sigma\}_{\sigma\in\mathcal{X}}$ of smooth kontravariant k-tensors T_σ on σ such that

$$\iota^*_{\sigma'\sigma} T_\sigma = T_{\sigma'},$$

whenever σ' is a face of σ. A *Riemannian simplicial complex* is then a simplicial complex $\mathcal{X}$ equipped with a *piecewise-smooth Riemannian metric*, i.e. a piecewise-smooth positive-definite symmetric 2-tensor g on $\mathcal{X}$.

The following lemma tells us that the space of constant piecewise-smooth symmetric tensors is isomorphic to functions on 1-simplices.

Lemma 6. *Let $\mathcal{X}$ be a simplicial complex and let $\mathcal{E}$ denote the set of its 1-simplices. For each function $f\colon \mathcal{E} \to \mathbb{R}$ there exists a unique constant piecewise-smooth symmetric 2-tensor S such that for each 1-simplex $e = \{i, j\}$*

$$S_e(j - i, j - i) = f(e).$$

Proof. It is enough to consider a single affine n-simplex $\sigma = \{i_0, \dots, i_n\}$ with vector space V. Consider the map F that sends a symmetric 2-tensor S on V to the function given by

$$F(S)(e) := S(i_k - i_j, i_k - i_j), \quad e = \{i_j, i_k\} \subset \sigma.$$

Clearly, F is linear. Moreover, if Q denotes the quadratic form corresponding to S, i.e. $Q(X) := S(X, X)$, then

$$S(X, Y) = \tfrac{1}{2}\big(Q(X) + Q(Y) - Q(X - Y)\big).$$

Hence, from $F(S) = 0$ follows $S = 0$. Thus F is injective. Clearly, the space of symmetric bilinear forms is of dimension $n(n + 1)/2$, which equals the number of 1-simplices. Thus F is an isomorphism. This proves the claim. □

It is also easy to write down the corresponding symmetric tensor in coordinates: Let $\sigma = \{i_0, \dots, i_n\}$ be a simplex. The vectors $e_j := i_j - i_0$, $j = 1, \dots, n$, then yield a basis of the corresponding vector space. Let f be a function defined on the unoriented edges of σ and let x_{i_j} denote the barycentric coordinates of its vertices i_j, then the corresponding symmetric bilinear form S^f_σ is given by

$$S^f_\sigma = \sum_{1\le j\le n} f_{i_0 i_j}\, dx_{i_j} \otimes dx_{i_j} + \sum_{1\le j,k\le n,\, j\neq k} \tfrac{1}{2}\big(f_{i_0 i_j} + f_{i_0 i_k} - f_{i_j i_k}\big)\, dx_{i_j} \otimes dx_{i_k}. \tag{11.1}$$

Thus starting with a positive function f, by Sylvester's criterion, it has to satisfy on each n-simplex $n - 1$ inequalities to determine a positive-definite form. If the corresponding piecewise-smooth form is positive-definite, we call f a discrete metric.

Definition 16. *A Euclidean simplicial complex is a simplicial complex $\mathcal{X}$ equipped with a discrete metric, i.e. a map ℓ that assigns to each 1-simplex e a length $\ell_e > 0$ such that for each simplex σ the symmetric tensor S^ℓ_σ is positive-definite.*

Now, let $\mathcal{X}$ be a Euclidean simplicial manifold of dimension n and denote by $\mathcal{X}_n$ the set of its top-dimensional simplices. Since each simplex of $\mathcal{X}$ is equipped with a scalar product it comes with a corresponding density and hence we know how to integrate functions over the simplices of $\mathcal{X}$. Now, we define the *integral over* $\mathcal{X}$ as follows:

$$\int_{\mathcal{X}} f := \sum_{\sigma\in\mathcal{X}_n} \int_\sigma f_\sigma, \quad f \in \Omega^0_{ps}(\mathcal{X}, \mathbb{R}^n).$$

Moreover, given a piecewise-smooth hermitian line bundle $\tilde{\mathrm{L}} \to \mathcal{X}$ with curvature, then there is a canonical hermitian product $\langle\!\langle .,. \rangle\!\rangle$ on $\Gamma_{ps}(\tilde{\mathrm{L}})$: If $\tilde\psi, \tilde\phi \in \Gamma_{ps}(\tilde{\mathrm{L}})$, then

$$\langle\!\langle \tilde\psi, \tilde\phi \rangle\!\rangle = \int_{\mathcal{X}} \langle \tilde\psi, \tilde\phi \rangle.$$

In particular, if $\tilde{\mathrm{L}}$ is the associated piecewise-smooth bundle of a discrete hermitian line bundle L with curvature Ω, then we can use ι to pull $\langle\!\langle .,. \rangle\!\rangle$ back to $\Gamma(\mathrm{L})$. Since ι is injective this yields a hermitian product on $\Gamma(\mathrm{L})$.

Now, we want to compute this metric explicitly in terms of given discrete data.

Definition 17. *A piecewise-linear section $\tilde\psi \in \Gamma_{pl}(\tilde{\mathrm{L}})$ is called concentrated at a vertex i, if it is of the form $\tilde\psi = \iota(\psi_i)$ for some vector $\psi_i \in \mathrm{L}_i$.*

It is basically enough to compute the product of two such concentrated sections. Therefore, let $\psi_i \in \mathrm{L}_i$ and $\psi_j \in \mathrm{L}_j$ and let $\tilde\psi^i$ and $\tilde\psi^j$ denote the corresponding piecewise-linear concentrated sections.

Now consider their product $\langle \tilde\psi^i, \tilde\psi^j \rangle$. Clearly, this product has support $S_i \cap S_j$. For simplicity, we extend the discrete connection η to arbitrary pairs ij in such way that $\eta_{ii} = \mathrm{id}$ and $\eta_{ij} \colon \mathrm{L}_i \to \mathrm{L}_j$ is zero whenever $\{i,j\} \notin \mathcal{X}$. With this convention, Equation (10.1) yields

$$\langle \tilde\psi^j, \tilde\psi^i \rangle = \langle \psi_j, \eta_{ij}(\psi_i) \rangle \, x_i x_j \, \exp\bigl(-\imath \int_{\Delta^p_{ij}} \tilde\Omega\bigr), \tag{11.2}$$

where $\tilde\Omega$ denotes the constant piecewise-smooth curvature form associated to Ω.

Now, let us express the integral over Δ^p_{ij} on a given n-simplex. Therefore consider an n-simplex $\sigma = \{i_0, \dots, i_n\}$. The hat functions $x_{i_1}, \dots, x_{i_n}$ yield affine coordinates on σ and we can express any 2-form with respect to the basis forms $dx_{i_j} \wedge dx_{i_k}$. One easily shows that

$$\int_{\sigma'} dx_{i_j} \wedge dx_{i_k} = \begin{cases} \pm\frac{1}{2} & \text{for } \sigma' = \pm i_j i_k i_\ell, \\ 0 & \text{else.} \end{cases}$$

Thus we obtain

$$\tilde\Omega = \sum_{1 \le j < k \le n} 2\,\Omega_{i_0 i_j i_k}\, dx_{i_j} \wedge dx_{i_k}.$$

Now we want to compute the integral over the triangle $\Delta^p_{i_0 i_1} \subset \sigma$. By Stokes theorem,

$$\int_{\Delta^p_{i_0 i_1}} dx_{i_j} \wedge dx_{i_k} = \int_{i_0}^{i_1} x_{i_j}\, dx_{i_k} + \int_{i_1}^{p} x_{i_j}\, dx_{i_k} + \int_{p}^{i_0} x_{i_j}\, dx_{i_k},$$

where the integrals are computed along straight lines. A small computation shows

$$\int_{\Delta^p_{i_0 i_1}} dx_{i_j} \wedge dx_{i_k} = \frac{1}{2}\bigl(\delta_{1j}\, x_{i_k}(p) - \delta_{1k}\, x_{i_j}(p)\bigr),$$

Thus, for $j < k$, we get $\int_{\Delta^p_{i_0 i_1}} dx_{i_j} \wedge dx_{i_k} = \frac{1}{2}\delta_{1j}\, x_{i_k}(p)$ and hence

$$\int_{\Delta^p_{i_0 i_1}} \tilde\Omega = \sum_{1 \le j < k \le n} 2\,\Omega_{i_0 i_j i_k} \int_{\Delta^p_{i_0 i_1}} dx_{i_j} \wedge dx_{i_k} = \sum_j \Omega_{i_0 i_1 i_j} x_{i_j}(p),$$

where we have used the convention that Ω vanishes on all triples not representing an oriented 2-simplex of $\mathcal{X}$. With this convention Equation (11.2) becomes

$$\langle \tilde\psi^j, \tilde\psi^i \rangle = \langle \psi_j, \eta_{ij}(\psi_i) \rangle \, x_i x_j \, \exp\bigl(-\imath \sum_k \Omega_{ijk} x_k\bigr). \tag{11.3}$$

In particular, using Equation (11.3), we can compute the norm of a piecewise-linear section $\tilde\psi$ on a given triangle ijk. Therefore we distinguish one of its vertices, say i, and write $\tilde\psi$ with respect to a section which is radially parallel with respect to i. Now, one easily checks that

$$|\tilde\psi| = |c_i + x_j(c_j e^{i\Omega_{ijk} x_k} - c_i) + x_k(c_k e^{-i\Omega_{ijk} x_j} - c_i)|,$$

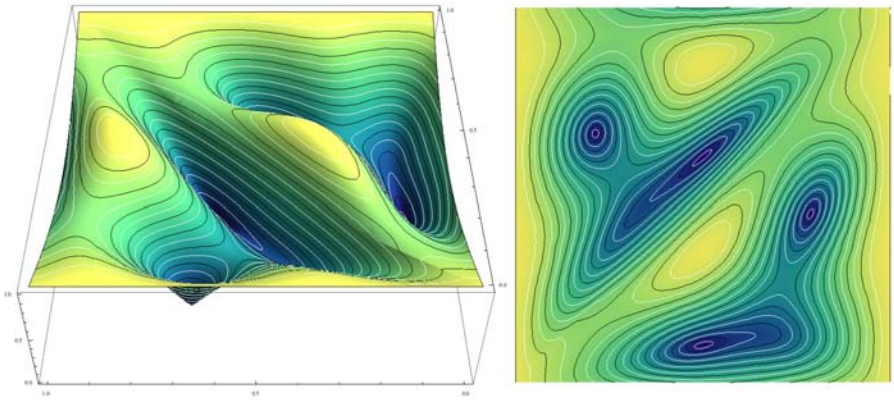

Figure 7. The norm of a piecewise-linear section of a bundle over a torus consisting of two triangles. Its two smooth parts fit continuously together along the diagonal. In this example the curvature of the bundle over each triangles is equal to 4π. The section has 4 zeros - just as predicted by the Poicaré-Hopf index formula.

where $c_i, c_j, c_k \in \mathbb{C}$ are constants depending on the explicit form of $\tilde{\psi}$. An example of the norm of a piecewise-linear section is shown in Figure 7.

As the next proposition shows, the identification of discrete and piecewise-linear sections perfectly fits together with the definitions in Section 8.

Proposition 7. *Let $\psi \in \Gamma(\mathrm{L})$ be a discrete section and let $\tilde{\psi} \in \Gamma_{pl}(\tilde{\mathrm{L}})$ be the corresponding piecewise-linear section, i.e. $\tilde{\psi} = \iota(\psi)$. Then, if $\tilde{\psi}$ has no zeros on edges, the discrete rotation form ξ^{ψ} and the piecewise-smooth rotation form $\xi^{\tilde{\psi}}$ are related as follows: For each oriented edge ij,*

$$\xi^{\psi}_{ij} = \int_{ij} \xi^{\tilde{\psi}}.$$

Proof. The claim follows easily by expressing $\tilde{\psi}$ with respect to some non-vanishing parallel section along the edge ij. □

In particular, by Theorem 9, the index form of a non-vanishing section of a discrete hermitian line bundle with curvature counts the number of (signed) zeros of the corresponding piecewise-linear section of the associated piecewise-smooth bundle.

Let us continue with the computation of the metric on $\Gamma(\mathrm{L})$. To write down the formula we give the following definition.

Definition 18. *Let $\mathcal{X}$ be an n-dimensional simplicial manifold and let $\Omega \in \Omega^2(\mathcal{X}, \mathbb{R})$. To an n-simplex σ and vertices i, j, k, l of $\mathcal{X}$ we assign the value*

$$\Theta^{\Omega}_{\sigma,i,j}(k,l) := \frac{1}{\operatorname{vol}(\sigma)} \int_{\sigma} x_k x_l \, \exp\bigl(-\imath \sum_m \Omega_{ijm} x_m\bigr),$$

where have chosen for integration an arbitrary discrete metric on $\mathcal{X}$.

Remark 11: *Note that the functions $\Theta^{\Omega}_{\sigma,i,j}$ are indeed well-defined. On a simplex, any two such measures induced by a discrete metric differ just by a constant.*

With Definition 18 and Equation (11.3) we obtain the following form of the metric:

Theorem 15 (Product of Discrete Sections)**.** *Let L be a discrete hermitian line bundle with curvature Ω over an n-dimensional Euclidean simplicial manifold $\mathcal{X}$, then the product on $\Gamma(\mathrm{L})$ induced by the associated piecewise-smooth hermitian line bundle is given as follows: Given two discrete sections $\psi = \sum_i \psi_i, \phi = \sum_i \phi_i$,*

$$\langle\!\langle \psi, \phi \rangle\!\rangle = \sum_{i,j} \mu^{ij}_{\Omega} \, \langle \psi_j, \eta_{ij}(\phi_i) \rangle, \quad \textit{where} \quad \mu^{ij}_{\Omega} = \sum_{\{i,j\} \supset \sigma \in \mathcal{X}_n} \Theta^{\Omega}_{\sigma,i,j}(i,j) \operatorname{vol}(\sigma).$$

Note that $\Theta^{\Omega}_{\sigma,i,j}(k,l)$, and hence μ^{ij}_{Ω}, can be computed explicitly using Fubini's theorem and the following small lemma one easily proves by induction.

Lemma 7. *Let $c \in \mathbb{C}_*$, $n \in \mathbb{N}$ and $[a,b] \subset \mathbb{R}$ be an interval. Then*

$$\int_a^b x^n \exp(cx)dx = \frac{n!}{c^{n+1}}\Big(\sum_{k=0}^{n}(-1)^k\frac{(cx)^{n-k}}{(n-k)!}\Big)\exp\big(cx\big)\Big|_a^b.$$

Next, we would like to compute the *Dirichlet energy* of a section $\tilde{\psi} \in \Gamma_{pl}(\tilde{\mathrm{L}})$, i.e.

$$\mathrm{E}_D(\tilde{\psi}) = \int_{\mathfrak{X}} \big|\nabla\tilde{\psi}\big|^2.$$

Note, that the Dirichlet energy comes with a corresponding positive-semidefinite hermitian form $\langle\!\langle .,. \rangle\!\rangle_D$ - called the *Dirichlet product*. Clearly, like the metric, the Dirichlet product is completely determined by the values it takes on concentrated sections.

In general, if $\tilde{\psi} \neq 0$ is piecewise-linear section concentrated at i, it is given on the vertex star S_i as a product $\tilde{\psi} = x_i\,\tilde{\phi}$, where x_i denotes the barycentric coordinate of the vertex i and $\tilde{\phi}$ is a local section radially parallel with respect to i. Clearly,

$$\nabla\tilde{\psi} = dx_i\,\tilde{\phi} + \imath\, x_i\,\omega_i\,\tilde{\phi},$$

where ω_i denotes the rotation form of $\tilde{\phi}$, i.e. $\nabla\tilde{\phi} = \imath\omega_i\,\tilde{\phi}$. Note here that ω_i does not depend on the actual value of $\tilde{\psi}$ at i, but is the same for all non-vanishing piecewise-linear sections concentrated at i.

To compute the rotation form ω_i at a given point $p_0 \in S_i$, we use a local section ζ which is radially parallel with respect to p_0 such that $\zeta_{p_0} = \tilde{\phi}_{p_0}$. Then we can express $\tilde{\phi}$ in terms of ζ, i.e.

$$\tilde{\phi} = z\,\zeta,$$

for some piecewise-smooth $\mathbb{C}_*$-valued function z defined locally at p_0. Clearly, $|z|$ is constant, and hence

$$\imath\omega_i\big|_{p_0}\ \tilde{\phi}_{p_0} = \nabla\tilde{\phi}\Big|_{p_0} = dz|_{p_0}\ \zeta_{p_0} + z(p_0)\ \nabla\zeta|_{p_0} = d\log z|_{p_0}\,\tilde{\phi}_{p_0} = \imath d\arg z|_{p_0}\ \tilde{\phi}_{p_0}.$$

The clue is that we can now use the relation of parallel transport and curvature to obtain an explicit formula for z. If p is sufficiently close to p_0, then the three points p, i and p_0 determine an oriented triangle Δ^p which is contained in a simplex of $\mathfrak{X}$. Its boundary curve γ_p consists of three line segments $\gamma_1, \gamma_2, \gamma_3$ connecting p to i, i to p_0 and p_0 back to p. Hence on each of these segments either $\tilde{\phi}$ of ζ is parallel and

$$\zeta_p = P_{\gamma_p}(\tilde{\phi}_p) = \exp\big(\imath\int_{\Delta^p}\tilde{\Omega}\big)\tilde{\phi}_p.$$

Thus we obtain that $z(p) = \exp\big(-\imath\int_{\Delta^p}\tilde{\Omega}\big)$ and hence

$$\omega_i\Big|_{p_0} = -d\big(\int_{\Delta^p}\tilde{\Omega}\big)\Big|_{p_0}.$$

Now, if Δ^p is contained in a simplex $\sigma = \{i_0,\dots,i_n\}$, one easily verifies that

$$\int_{\Delta^p} dx_{i_j}\wedge dx_{i_k} = \frac{1}{2}\big(x_{i_j}(p_0)x_{i_k}(p) - x_{i_k}(p_0)x_{i_j}(p)\big).$$

Thus,

$$\begin{aligned} d\big(\int_{\Delta^p}\tilde{\Omega}\big)\Big|_{p_0} &= \sum_{1\le j<k\le n} 2\,\Omega_{i_0i_ji_k}\,d\big(\int_{\Delta^p} dx_{i_j}\wedge dx_{i_k}\big)\Big|_{p_0} \\ &= \sum_{1\le j<k\le n}\Omega_{i_0i_ji_k}\left(x_{i_j}dx_{i_k} - x_{i_k}dx_{i_j}\right)\Big|_{p_0}, \\ &= \sum_{1\le j\le n}\big(\sum_{k\neq j}\Omega_{i_0i_ji_k}x_{i_k}\big)dx_{i_j}\Big|_{p_0} \end{aligned}$$

and, using the convention on Ω from above, we find the following simple formula:

$$\omega_i = \sum_j \big(\sum_k \Omega_{ijk}\, x_k\big)\, dx_j\Big|_{S_i}, \tag{11.4}$$

where we sum over the whole vertex set of $\mathcal{X}$.

Now, given this local form expressions, we can finally return to the computation of the products which we are actually interested in. Therefore we consider two piecewise-linear sections concentrated at the vertices i and j:

$$\tilde{\psi}^i := \iota(\psi_i), \quad \tilde{\psi}^j := \iota(\psi_j),$$

for some $\psi_i \in \mathrm{L}_i$ and $\psi_j \in \mathrm{L}_j$. On their common support $S_i \cap S_j$ both section can be expressed, just as above, as products of a real-valued piecewise-linear hat functions x_i and x_j and radially parallel local sections $\tilde{\phi}^i$ and $\tilde{\phi}^j$:

$$\tilde{\psi}^i = x_i\, \tilde{\phi}^i, \quad \tilde{\psi}^j = x_j\, \tilde{\phi}^j.$$

Clearly,

$$\begin{aligned}\langle\!\langle \tilde{\psi}^j, \tilde{\psi}^i \rangle\!\rangle_D &= \int_{S_i \cap S_j} \langle dx_j \tilde{\phi}^j + \imath x_j \omega_j\, \tilde{\phi}^j, dx_i \tilde{\phi}^i + \imath x_i \omega_i\, \tilde{\phi}^i \rangle \\ &= \int_{S_i \cap S_j} \langle dx_j + \imath x_j \omega_j, dx_i + \imath x_i \omega_i \rangle\, \langle \tilde{\phi}^j, \tilde{\phi}^i \rangle.\end{aligned}$$

With Equation (11.3) we see that

$$\langle \tilde{\phi}^j, \tilde{\phi}^i \rangle = \langle \psi_j, \eta_{ij}(\psi_i) \rangle\, \exp\big(-\imath \sum_m \Omega_{ijm} x_m\big).$$

Moreover, by Equation (11.4),

$$\begin{aligned}\langle dx_j + \imath x_j \omega_j, dx_i + \imath x_i \omega_i \rangle =& \Big[\langle dx_j, dx_i \rangle + \sum_{k',k'',l',l''} \Omega_{ik'l'} \Omega_{jk''l''} x_j x_i x_{l'} x_{l''} \langle dx_{k'}, dx_{k''} \rangle\Big] \\ &+ \imath \Big[\sum_{k',l'} (\Omega_{ik'l'} x_i x_{l'} \langle dx_j, dx_{k'} \rangle - \Omega_{jk'l'} x_j x_{l'} \langle dx_{k'}, dx_i \rangle)\Big].\end{aligned}$$

The constants $\langle dx_{k'}, dx_{l'} \rangle$ are basically provided by the following lemma.

Lemma 8. *Let $\sigma = \{v_0, \dots, v_n\}$ be a Euclidean simplex of dimension $n > 0$ and let x_i denote its barycentric coordinate functions. Then*

$$\operatorname{grad} x_i = -\frac{1}{h_i} N_i,$$

where h_i denotes the distance between v_i and $\sigma_i = \sigma \setminus \{v_i\}$ and N_i denotes the outward-pointing unit normal of σ_i.

Proof. This immediately follows from two basic facts: First, $dx_i(v_j - v_0) = \delta_{ij}$ for $i, j > 0$. Second, $h_i = \langle v_0 - v_i, N_i \rangle$. □

Moreover, Lemma 8 yields almost immediately a higher dimensional analogue of the well-known cotangent formula for surfaces.

Theorem 16 (Cotangent Formula)**.** *Let σ be a simplex of a Euclidean simplicial complex $\mathcal{X}$ and let $\dim \sigma > 1$. If $i \neq j$,*

$$c_\sigma^{ij} := \int_\sigma \langle dx_i, dx_j \rangle = \begin{cases} -\frac{1}{n(n-1)} \cot \alpha_\sigma^{ij} \operatorname{vol}\big(\sigma \setminus \{i,j\}\big), & \text{if } \{i,j\} \subset \sigma, \\ 0 & \text{else.} \end{cases}$$

Here α_σ^{ij} denotes the angle between the faces $\sigma \setminus \{i\}$ and $\sigma \setminus \{j\}$. Moreover,

$$c_\sigma^{ii} := \int_\sigma |dx_i|^2 = \begin{cases} \frac{1}{n\, h_i} \operatorname{vol}\big(\sigma \setminus \{i\}\big), & \text{if } i \in \sigma, \\ 0 & \text{else,} \end{cases}$$

where h_i denotes the distance between the vertex i and the face $\sigma \setminus \{i\}$.

PROOF. Clearly, if $\{i,j\} \not\subset \sigma$, then $\int_\sigma \langle dx_i, dx_j\rangle = 0$. Now, let $\{i,j\} \subset \sigma$, $i \neq j$. With the notation of Lemma 8, we have

$$\int_\sigma \langle dx_i, dx_j\rangle = \langle \operatorname{grad} x_i, \operatorname{grad} x_j\rangle \operatorname{vol}\sigma = \frac{\langle N_i, N_j\rangle}{h_i h_j} \operatorname{vol}\sigma.$$

Clearly, $\cos\alpha_\sigma^{ij} = -\langle N_i, N_j\rangle$ and $n!\operatorname{vol}\sigma = (n-2)!\, h_i h_j \sin\alpha_\sigma^{ij} \operatorname{vol}\big(\sigma\setminus\{i,j\}\big)$, which yields the first part of the theorem. Similarly, $n\operatorname{vol}\sigma = h_i \operatorname{vol}\big(\sigma\setminus\{i\}\big)$. Setting $i=j$ then immediately yields the second part. □

DEFINITION 19. *Let $\mathcal{X}$ be an n-dimensional simplicial manifold and let $\Omega \in \Omega^2(\mathcal{X},\mathbb{R})$. Let σ be an n-simplex and i,j,k,l be vertices of $\mathcal{X}$. Then, let*

$$\Lambda^\Omega_{\sigma,i,j} := \frac{1}{\operatorname{vol}(\sigma)} \int_\sigma \exp\big(-\imath \sum_m \Omega_{ijm} x_m\big),$$

$$\Xi^\Omega_{\sigma,i,j}(k,l) := \frac{1}{\operatorname{vol}(\sigma)} \int_\sigma x_i x_j x_k x_l \, \exp\big(-\imath \sum_m \Omega_{ijm} x_m\big),$$

where we choose for the integration an arbitrary discrete metric on $\mathcal{X}$.

Remark 12: *Just like the functions $\Theta^\Omega_{\sigma,i,j}$, the values $\Lambda^\Omega_{\sigma,i,j}$ and the functions $\Xi^\Omega_{\sigma,i,j}$ and are well-defined (compare Remark 11).*

Now, with these definitions, we can summarize the above discussion by the following theorem.

THEOREM 17 (Discrete Dirichlet Energy)**.** *Let* L *be a discrete hermitian line bundle with curvature Ω over an n-dimensional Euclidean simplicial manifold $\mathcal{X}$, then the Dirichlet product on $\Gamma(\mathrm{L})$ induced by the associated piecewise-smooth hermitian line bundle is given as follows: If $\phi = \sum_i \phi_i$ and $\psi = \sum_i \psi_i$ are two discrete sections,*

$$\langle\!\langle \phi, \psi\rangle\!\rangle_D = \sum_{i,j} w^{ij}_\Omega \, \langle \phi_j, \eta_{ij}(\psi_i)\rangle, \quad w^{ij}_\Omega = \sum_{\{i,j\}\supset\sigma\in\mathcal{X}_n} W^\Omega_{\sigma,i,j},$$

where

$$W^\Omega_{\sigma,i,j} = \Big[c^{ij}_\sigma \, \Lambda^\Omega_{\sigma,i,j} + \sum_{k',k'',l',l''} \Omega_{ik'l'}\Omega_{jk''l''}\, c^{k'k''}_\sigma \, \Xi^\Omega_{\sigma,i,j}(l',l'')\Big] + \imath\Big[\sum_{k',l'}\big(\Omega_{ik'l'}\, c^{jk'}_\sigma\, \Theta^\Omega_{\sigma,i,j}(i,l') - \Omega_{jk'l'}\, c^{ik'}_\sigma\, \Theta^\Omega_{\sigma,i,j}(j,l')\big)\Big]. \tag{11.5}$$

12 Discrete Energies on Surfaces - An Example

While the computation of the Dirichlet product $\langle\!\langle .,. \rangle\!\rangle_D$ and the metric $\langle\!\langle .,. \rangle\!\rangle$ of discrete sections is quite complicated and tedious for higher dimensional simplicial manifolds, it is manageable for the 2-dimensional case. We are going to compute it explicitly.

Throughout this section let L denote a discrete hermitian line bundle with curvature Ω over a Euclidean simplicial surface $\mathcal{X}$ and let $\sigma = \{i,j,k\}$ be one of its triangles.

The metric $\langle\!\langle .,. \rangle\!\rangle$ is easily obtained. We basically just need to compute the values $\Theta^\Omega_{\sigma,i,i}(i,i)$ and $\Theta^\Omega_{\sigma,i,j}(i,j)$, which can be done over the standard triangle. We get

$$\Theta^\Omega_{\sigma,i,i}(i,i) = \frac{1}{6}, \quad \Theta^\Omega_{\sigma,i,j}(i,j) = 2\frac{\exp(-\imath\Omega_{ijk}) - 1 + \imath\Omega_{ijk} + \frac{1}{2}\Omega^2_{ijk} - \imath\frac{1}{6}\Omega^3_{ijk}}{\Omega^4_{ijk}}. \tag{12.1}$$

Now, we compute the Dirichlet product $\langle\!\langle .,. \rangle\!\rangle_D$ on $\mathcal{X}$. For $n=2$, the expressions $W^\Omega_{\sigma,i,i}$ and $W^\Omega_{\sigma,i,j}$ simplify drastically. First, we look at the diagonal terms. We have

$$\sum_{k',k'',l',l''} c^{k'k''}_\sigma \Omega_{ik'l'}\Omega_{ik''l''}\, \Xi^\Omega_{\sigma,i,i}(l',l'') = \Big(c^{jj}_\sigma \Xi^\Omega_{\sigma,i,i}(k,k) - 2c^{jk}_\sigma \Xi^\Omega_{\sigma,i,i}(j,k) + c^{kk}_\sigma \Xi^\Omega_{\sigma,i,i}(j,j)\Big)\Omega^2_{ijk},$$

and with

$$\Lambda_{\sigma,i,i} = 1, \quad \Xi_{\sigma,i,i}(j,j) = \frac{1}{90} = \Xi_{\sigma,i,i}(k,k), \quad \Xi_{\sigma,i,i}(j,k) = \frac{1}{180}$$

we get the following formula:

$$W^{\Omega}_{\sigma,i,i} = c^{ii}_{\sigma} + \frac{c^{jj}_{\sigma} - c^{jk}_{\sigma} + c^{kk}_{\sigma}}{90}\Omega^2_{ijk}.$$

Now we would like to obtain a similar formula for the off-diagonal terms. Since $dx_i + dx_j = -dx_k$, we have $c^{jk}_{\sigma} + c^{ki}_{\sigma} = -c^{kk}_{\sigma}$. Hence,

$$\sum_{k',k'',l',l''} c^{k'k''}_{\sigma}\Omega_{ik'l'}\Omega_{jk''l''}\,\Xi^{\Omega}_{\sigma,i,j}(l',l'') = -\Big(c^{ij}_{\sigma}\Xi^{\Omega}_{\sigma,i,j}(k,k) + c^{kk}_{\sigma}\big(\Xi^{\Omega}_{\sigma,i,j}(i,j) + \Xi^{\Omega}_{\sigma,i,j}(j,k)\big)\Big)\Omega^2_{ijk}.$$

This time the expressions become more complicated. We get

$$\Xi^{\Omega}_{\sigma,i,j}(k,k) = \tfrac{2}{\Omega^6_{ijk}}\Big(20 - 12\imath\Omega_{ijk} - 3\Omega^2_{ijk} + \tfrac{1}{3}\imath\Omega^3_{ijk} + \big(-20 - 8\imath\Omega_{ijk} + \Omega^2_{ijk}\big)\exp\big(-\imath\Omega_{ijk}\big)\Big),$$

$$\Xi^{\Omega}_{\sigma,i,j}(i,j) + \Xi^{\Omega}_{\sigma,i,j}(j,k) = \tfrac{2}{\Omega^6_{ijk}}\Big(-6 + 4\imath\Omega_{ijk} + \Omega^2_{ijk} + \tfrac{1}{12}\Omega^4_{ijk} - \tfrac{1}{30}\imath\Omega^5_{ijk} + \big(6 + 2\imath\Omega_{ijk}\big)\exp\big(-\imath\Omega_{ijk}\big)\Big).$$

Thus,

$$\sum_{k',k'',l',l''} c^{k'k''}_{\sigma}\Omega_{ik'l'}\Omega_{jk''l''}\,\Xi^{\Omega}_{\sigma,i,j}(l',l'') =$$
$$\frac{2}{\Omega^4_{ijk}}\Big(\big[6c^{kk}_{\sigma} - 20c^{ij}_{\sigma}\big] + \big[12c^{ij}_{\sigma} - 4c^{kk}_{\sigma}\big]\imath\Omega_{ijk} + \big[3c^{ij}_{\sigma} - c^{kk}_{\sigma}\big]\Omega^2_{ijk} - \tfrac{c^{ij}_{\sigma}}{3}\imath\Omega^3_{ijk} - \tfrac{c^{kk}_{\sigma}}{12}\Omega^4_{ijk}$$
$$+ \tfrac{c^{kk}_{\sigma}}{30}\imath\Omega^5_{ijk} + \big(\big[20c^{ij}_{\sigma} - 6c^{kk}_{\sigma}\big] + \big[8c^{ij}_{\sigma} - 2c^{kk}_{\sigma}\big]\imath\Omega_{ijk} - c^{ij}_{\sigma}\Omega^2_{ijk}\big)\exp\big(-\imath\Omega_{ijk}\big).\Big)$$

Now, let us look at the second sum in Equation (11.5). We have

$$\imath\sum_{k',l'}\big(\Omega_{ik'l'}c^{jk'}_{\sigma}\Theta^{\Omega}_{\sigma,i,j}(i,l') - \Omega_{jk'l'}c^{ik'}_{\sigma}\Theta^{\Omega}_{\sigma,i,j}(j,l')\big) = \Big(c^{ii}_{\sigma}\Theta^{\Omega}_{\sigma,i,j}(j,k) + c^{jj}_{\sigma}\Theta^{\Omega}_{\sigma,i,j}(k,i) + c^{kk}_{\sigma}\Theta^{\Omega}_{\sigma,i,j}(i,j)\Big)\imath\Omega_{ijk}.$$

The formula for $\Theta^{\Omega}_{\sigma,i,j}(i,j)$ is already given in Equation (12.1). Further, we have

$$\Theta^{\Omega}_{\sigma,i,j}(j,k) = \frac{2}{\Omega^4_{ijk}}\Big(3 - 2\imath\Omega_{ijk} - \tfrac{1}{2}\Omega^2_{ijk} + \big(-3 + \imath\Omega_{ijk}\big)\exp\big(-\imath\Omega_{ijk}\big)\Big) = \Theta^{\Omega}_{\sigma,i,j}(k,i).$$

Thus we get

$$\imath\sum_{k',l'}\big(\Omega_{ik'l'}c^{jk'}_{\sigma}\Theta^{\Omega}_{\sigma,i,j}(i,l') - \Omega_{jk'l'}c^{ik'}_{\sigma}\Theta^{\Omega}_{\sigma,i,j}(j,l')\big) =$$
$$\frac{2}{\Omega^4_{ijk}}\Big(\big[3(c^{ii}_{\sigma} + c^{jj}_{\sigma}) - c^{kk}_{\sigma}\big]\imath\Omega_{ijk} + \big[2(c^{ii}_{\sigma} + c^{jj}_{\sigma}) - c^{kk}_{\sigma}\big]\Omega^2_{ijk} + \tfrac{1}{2}\big[c^{kk}_{\sigma} - c^{ii}_{\sigma} - c^{jj}_{\sigma}\big]\imath\Omega^3_{ijk}$$
$$+\tfrac{c^{kk}_{\sigma}}{6}\Omega^4_{ijk} + \big(\big[c^{kk}_{\sigma} - 3(c^{ii}_{\sigma} + c^{jj}_{\sigma})\big]\imath\Omega_{ijk} + \big[c^{ii}_{\sigma} + c^{jj}_{\sigma}\big]\Omega^2_{ijk}\big)\exp\big(-\imath\Omega_{ijk}\big)\Big).$$

Hence, with

$$\Lambda^{\Omega}_{\sigma,i,j} = \frac{2}{\Omega^4_{ijk}}\Big(\Omega^2_{ijk} - \imath\Omega^3_{ijk} - \Omega^2_{ijk}\exp\big(-\imath\Omega_{ijk}\big)\Big),$$

Equation (11.5) becomes

$$W^{\Omega}_{\sigma,i,j} = \frac{2}{\Omega^4_{ijk}}\Big(\big[6\sigma^{kk} - 20c^{ij}_{\sigma}\big] + \big[12c^{ij}_{\sigma} + 3(c^{ii}_{\sigma} + c^{jj}_{\sigma}) - 5c^{kk}_{\sigma}\big]\imath\Omega_{ijk} + \big[4c^{ij}_{\sigma} + 2(c^{ii}_{\sigma} + c^{jj}_{\sigma} - c^{kk}_{\sigma})\big]\Omega^2_{ijk}$$
$$+\tfrac{1}{6}\big[3(c^{kk}_{\sigma} - c^{ii}_{\sigma} - c^{jj}_{\sigma}) - 8c^{ij}_{\sigma}\big]\imath\Omega^3_{ijk} + \tfrac{1}{12}c^{kk}_{\sigma}\Omega^4_{ijk} + \tfrac{1}{30}c^{kk}_{\sigma}\Omega^4_{ijk}$$
$$+\big(\big[20c^{ij}_{\sigma} - 6c^{kk}_{\sigma}\big] + \big[8c^{ij}_{\sigma} - 3(c^{ii}_{\sigma} + c^{jj}_{\sigma}) - c^{kk}_{\sigma}\big]\imath\Omega_{ijk} + \big[c^{ii}_{\sigma} - 2c^{ij}_{\sigma} + c^{jj}_{\sigma}\big]\big]\Omega^2_{ijk}\big)\exp\big(-\imath\Omega_{ijk}\big)\Big).$$

Since $n = 2$, the weights c^{ij}_{σ} are just given as follows:

$$c^{ij}_{\sigma} = -\frac{\cot\alpha^{ij}_{\sigma}}{2}, \quad c^{kk}_{\sigma} = \frac{\ell_{ij}}{2h_k},$$

where ℓ_{ij} denotes the edge length. We would like to express them explicitly in terms of the Euclidean metric g of σ. In fact, we can distinguish the vertex k as origin and use the hat functions x_i and x_j as coordinates on σ. With respect to these coordinates, the metric is given by a matrix:

$$g = \begin{pmatrix} g_{11} & g_{12} \\ g_{21} & g_{22} \end{pmatrix}.$$

In terms of g the cotangent weights are given as follows:

$$c_\sigma^{ij} = -\frac{g_{12}}{2\sqrt{\det g}}, \qquad c_\sigma^{jk} = -\frac{g_{11}-g_{12}}{2\sqrt{\det g}}, \qquad c_\sigma^{ki} = -\frac{g_{22}-g_{12}}{2\sqrt{\det g}},$$
$$c_\sigma^{kk} = \frac{g_{11}-2g_{12}+g_{22}}{2\sqrt{\det g}}, \qquad c_\sigma^{ii} = \frac{g_{22}}{2\sqrt{\det g}}, \qquad c_\sigma^{jj} = \frac{g_{11}}{2\sqrt{\det g}},$$

and we have rederived the formulas in **[33]**:

$$\begin{aligned} W^\Omega_{\sigma,i,j} =& \frac{1}{\operatorname{vol}(\sigma)\Omega^4_{ijk}}\Big(\big[3g_{11}+4g_{12}+3g_{22}\big] - \big[g_{11}+g_{12}+g_{22}\big]\imath\Omega_{ijk} + \tfrac{g_{12}}{6}\imath\Omega^3_{ijk} \\ &+ \tfrac{g_{11}-2g_{12}+g_{22}}{24}\Omega^4_{ijk} + \tfrac{g_{11}-2g_{12}+g_{22}}{60}\Omega^4_{ijk} - \big(\big[3g_{11}+4g_{12}+3g_{22}\big] \\ &+ \big[2g_{11}+3g_{12}+2g_{22}\big]\imath\Omega_{ijk} - \tfrac{1}{2}\big[g_{11}+2g_{12}+g_{22}\big]\Omega^2_{ijk}\big)\exp\big(-\imath\Omega_{ijk}\big)\Big). \end{aligned}$$

List of Figures

Complex Line Bundles over Simplicial Complexes and their Applications

Bibliography

[1] Y. Aharonov and D. Bohm. Significance of electromagnetic potentials in the quantum theory. *Phys. Rev.*, 115:485–491, 1959.

[2] L. V. Ahlfors. *Complex Analysis*. McGraw-Hill, 2nd edition, 1966.

[3] J. Avron, D. Osadchy, and R. Seiler. A topological look at the quantum Hall effect. *Physics Today*, 56:38–42, 2003.

[4] M. Ben-Chen, A. Butscher, J. Solomon, and L. Guibas. On Discrete Killing Vector Fields and Patterns on Surfaces. *Comp. Graph. Forum*, 29(5):1701–1711, 2010.

[5] D. Bommes, H. Zimmer, and L. Kobbelt. Mixed-Integer Quadrangulation. *ACM Trans. Graph.*, 28(3), 2009.

[6] D. Bommes, H. Zimmer, and L. Kobbelt. Practical Mixed-Integer Optimization for Geometry Processing. In *Proc. 7th Int. Conf. Curves & Surfaces*, pages 193–206, 2012. Project page: http://www.graphics.rwth-aachen.de/software/comiso.

[7] R. Bott. On some recent interactions between mathematics and physics. *Canad. Math. Bull.*, 28:129–164, 1985.

[8] R. G. Chambers. Shift of an electron interference pattern by enclosed magnetic flux. *Phys. Rev. Lett.*, 5:3–5, 1960.

[9] Y. Chen, T. A. Davis, W. W. Hager, and S. Rajamanickam. CHOLMOD, Supernodal Sparse Cholesky Factorization, and Update/Downdate. *ACM Trans. Math. Softw.*, 35, 2008.

[10] Y. Chen, T. A. Davis, W. W. Hager, and S. Rajamanickam. Algorithm 887: CHOLMOD, Supernodal Sparse Cholesky Factorization and Update/Downdate. *ACM Trans. Math. Softw.*, 35(3):22:1–22:14, 2009.

[11] A. Chern, U. Pinkall, and P. Schröder. Close-to-conformal deformations of volumes. *ACM Trans. Graph.*, 34, 2015.

[12] S. Christiansen and T. Halvorsen. A gauge invariant discretization on simplicial grids of the Schrödinger eigenvalue problem in an electromagnetic field. *SIAM J. Numer. Anal.*, 49:331Ð–345, 2011.

[13] S. Christiansen and T. Halvorsen. A simplicial gauge theory. *J. Math. Phys.*, 53, 2012.

[14] D. Cohen-Steiner and J.-M. Morvan. Restricted Delaunay Triangulations and Normal Cycle. In *Proc. Symp. Comp. Geom.*, pages 312–321, 2003.

[15] K. Crane, F. de Goes, M. Desbrun, and P. Schröder. Digital Geometry Processing with Discrete Exterior Calculus. In *ACM SIGGRAPH 2013 courses*, 2013.

[16] K. Crane, M. Desbrun, and P. Schröder. Trivial Connections on Discrete Surfaces. *Comp. Graph. Forum*, 29(5):1525–1533, 2010.

[17] M. Desbrun, E. Kanso, and Y. Tong. Discrete differential forms for computational modeling. In *Discrete Differential Geometry*, pages 287–324. Birkhäuser Basel, 2008.

[18] M. Desbrun, E. Kanso, and Y. Tong. Discrete Differential Forms for Computational Modeling. In A. I. Bobenko, P. Schröder, J. M. Sullivan, and G. M. Ziegler, editors, *Discrete Differential Geometry*, volume 38 of *Oberwolfach Seminars*, pages 287–324. Birkhäuser Verlag, 2008.

[19] M. Desbrun, M. Meyer, and P. Alliez. Intrinsic Parameterizations of Surface Meshes. *Comp. Graph. Forum*, 21(3):209–218, 2002.

[20] O. Diamanti, A. Vaxman, D. Panozzo, and O. Sorkine-Hornung. Designing N-PolyVector Fields with Complex Polynomials. *Comp. Graph. Forum*, 33(5), 2014.

[21] P. A. M. Dirac. Quantised singularities in the electromagnetic field. *Phys. Roy. Soc. Lond.*, 133:60–72, 1931.

[22] M. Fisher, P. Schröder, M. Desbrun, and H. Hoppe. Design of Tangent Vector Fields. *ACM Trans. Graph.*, 26(3), 2007.

[23] A. Gil, J. Segura, and N. M. Temme. *Numerical Methods for Special Functions*. SIAM, 2007.

[24] X. Gu and S.-T. Yau. Global Conformal Surface Parameterization. In *Proc. Symp. Geom. Proc.*, pages 127–137, 2003.

[25] T. Halvorsen and T. Sørensen. Simplicial gauge theory and quantum gauge theory simulation. *Nuclear Physics B*, 854:166–183, 2012.

[26] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.

[27] A. Hertzmann and D. Zorin. Illustrating Smooth Surfaces. In *Proc. ACM/SIGGRAPH Conf.*, pages 517–526, 2000.

[28] B. Jobard and W. Lefer. Creating evenly-spaced streamlines of arbitrary density. In W. Lefer and M. Grave, editors, *Vis. Sci. Comp.* Springer Vienna, 1997.

[29] F. Kälberer, M. Nieser, and K. Polthier. QuadCover - Surface Parameterization using Branched Coverings. *Comp. Graph. Forum*, 26(3):375–384, 2007.

[30] F. Kälberer, M. Nieser, and K. Polthier. Stripe parameterization of tubular surfaces. *Top. Meth. Data Anal. Vis*, 2010.
[31] S. Kalpakjian and S. R. Schmid. *Manufacturing Engineering and Technology*. Prentice Hall, New York, 6th edition, 2009.
[32] M. Kass and A. Witkin. Analyzing Oriented Patterns. *Comp. Vis., Graph., Im. Proc.*, 37(3):362–385, 1987.
[33] F. Knöppel, K. Crane, U. Pinkall, and P. Schröder. Globally Optimal Direction Fields. *ACM Trans. Graph.*, 32(4), 2013.
[34] F. Knöppel, K. Crane, U. Pinkall, and P. Schröder. Stripe patterns on surfaces. *ACM Trans. Graph.*, 34, 2015.
[35] S. Kobayashi. La connexion des variétés fibrés I and II. *Comptes Rendus de l'Académie Sciences, Paris*, 54:318–319, 443–444, 1954.
[36] B. Kostant. Quantization and unitary representations. In *Lectures in Modern Analysis and Applications III*, volume 170, pages 87–208. Springer, 1970.
[37] C. Kreft and R. Seiler. Models of the Hofstadter type. *J. Math. Phys.*, 37:5207–5243, 1996.
[38] J. M. Lee. *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics. Springer, 2012.
[39] S. Lefebvre and H. Hoppe. Appearance-space Texture Synthesis. *ACM Trans. Graph.*, 25(3):541–548, 2006.
[40] S. Lefebvre and H. Hoppe. Appearance-Space Texture Synthesis. *ACM Trans. Graph.*, 25(3), 2006.
[41] T. Levi-Civita. Nozione di parallelismo in una varietá qualunque e conseguente specificazione geometrica della curvatura riemanniana. *Rendiconti del Circolo Matematico di Palermo*, 42:173–204, 1916.
[42] B. Lévy, S. Petitjean, N. Ray, and J. Maillot. Least Squares Conformal Maps for Automatic Texture Atlas Generation. *ACM Trans. Graph.*, 21(3):362–371, 2002.
[43] C. Ling, J. Nie, L. Qi, and Y. Ye. Biquadratic Optimization Over Unit Spheres and Semidefinite Programming Relaxations. *SIAM J. on Opt.*, 20(3):1286–1310, 2009.
[44] R. Ling, J. Huang, B. Jüttler, F. Sun, H. Bao, and W. Wang. Spectral quadrangulation with feature curve alignment and element size control. *ACM Trans. Graph.*, 34(1), 2014.
[45] R. MacNeal. *The Solution of Partial Differential Equations by means of Electrical Networks*. PhD thesis, Caltech, 1949.
[46] A. Mebarki, P. Alliez, and O. Devillers. Farthest Point Seeding for Efficient Placement of Streamlines. *Proc. IEEE Vis*, 2005.
[47] K. Morrison. Yang-Mills connections on surfaces and representations of the path group. *Proceedings of the American Mathematical Society*, 112:1101–1106, 1991.
[48] P. Mullen, Y. Tong, P. Alliez, and M. Desbrun. Spectral Conformal Parameterization. *Comp. Graph. Forum*, 27(5):1487–1494, 2008.
[49] J. R. Munkres. *Elements of Algebraic Topology*. Advanced book classics. Perseus Books, 1984.
[50] A. Myles and D. Zorin. Global parameterization by incremental flattening. *ACM Trans. Graph.*, 31(4), 2012.
[51] A. Myles and D. Zorin. Controlled-distortion constrained global parameterization. *ACM Trans. Graph.*, 32(4), 2013.
[52] T. Napier and M. Ramachandran. *An Introduction to Riemann Surfaces*. Birkhäuser, 2011.
[53] M. Nieser, J. Palacios, K. Polthier, and E. Zhang. Hexagonal Global Parameterization of Arbitrary Surfaces. *IEEE Trans. Vis. Comp. Graph.*, 18(6):865–878, 2012.
[54] J. Palacios and E. Zhang. Rotational Symmetry Field Design on Surfaces. *ACM Trans. Graph.*, 26(3), 2007.
[55] M. J. Pflaum. *Analytic and Geometric Study of Stratified Spaces: Contributions to Analytic and Geometric Aspects*. Analytic and Geometric Study of Stratified Spaces. Springer, 2001.
[56] U. Pinkall and K. Polthier. Computing Discrete Minimal Surfaces and Their Conjugates. *Experiment. Math.*, 2(1):15–36, 1993.
[57] K. Polthier and M. Schmies. Straightest Geodesics on Polyhedral Surfaces. In H.-C. Hege and K. Polthier, editors, *Mathematical Visualization: Algorithms, Applications and Numerics*, pages 135–152. Springer Verlag, 1998.
[58] E. Praun, A. Finkelstein, and H. Hoppe. Lapped Textures. In *Proc. ACM/SIGGRAPH Conf.*, pages 465–470, 2000.
[59] N. Ray, W. C. Li, B. Lévy, A. Sheffer, and P. Alliez. Periodic Global Parameterization. *ACM Trans. Graph.*, 25(4):1460–1485, 2006.
[60] N. Ray, B. Vallet, L. Alonso, and B. Lévy. Geometry Aware Direction Field Processing. *ACM Trans. Graph.*, 29(1), 2009.
[61] N. Ray, B. Vallet, W. C. Li, and B. Lévy. N-Symmetry Direction Field Design. *ACM Trans. Graph.*, 27(2):10:1–10:13, 2008.
[62] N. Ray, B. Vallet, W. C. Li, and B. Lévy. N-symmetry direction field design. *ACM Trans. Graph.*, 27(2), 2008.
[63] R. J. Rost. *OpenGL(R) Shading Language (2nd Edition)*. Addison-Wesley Professional, 2005.
[64] K. Sasaki, Y. Kawazoe, and R. Saito. Aharonov-Bohm effect in higher genus materials. *Physica A*, 321:369–375, 2004.
[65] D. J. Simms and N. M. J. Woodhouse. *Lectures on Geometric Quantization*. Lecture Notes in Physics. Springer Berlin Heidelberg, 1976.
[66] B. Spencer, R. S. Laramee, G. Chen, and E. Zhang. Evenly Spaced Streamlines for Surfaces: An Image-Based Approach. *Comp. Graph. Forum*, 28(6), 2009.
[67] Y. Tong, P. Alliez, D. Cohen-Steiner, and M. Desbrun. Designing Quadrangulations with Discrete Harmonic Forms. In *Proc. Symp. Geom. Proc.*, 2006.
[68] G. Turk. Generating Textures on Arbitrary Surfaces using Reaction Diffusion. *Proc. ACM/SIGGRAPH Conf.*, 1991.
[69] G. Turk. Texture Synthesis on Surfaces. In *Proc. ACM/SIGGRAPH Conf.*, 2001.

[70] L.-Y. Wei and M. Levoy. Texture Synthesis over Arbitrary Manifold Surfaces. In *Proc. ACM/SIGGRAPH Conf.*, 2001.
[71] A. Weil. *Variétés kählériennes*. Actualités Scientifiques et Industrielles. Hermann, 1958.
[72] S. Weißmann, U. Pinkall, and P. Schröder. Smoke Rings from Smoke. *ACM Trans. Graph.*, 33(4), 2014.
[73] H. Weyl. Gravitation und Elektrizität. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin*, pages 465–480, 1918.
[74] H. Weyl. Elektron und Gravitation I. *Zeitschrift für Physik*, 56:330–352, 1929.
[75] W. Wirtinger. Zur formalen Theorie der Funktionen von mehr komplexen Veränderlichen. *Math. Ann.*, 97(1):357–375, 1927.
[76] A. Witkin and M. Kass. Reaction-Diffusion Textures. *Proc. ACM/SIGGRAPH Conf.*, 1991.
[77] C. N. Yang and R. L. Mills. Conservation of isotopic spin and isotopic gauge invariance. *Phys. Rev.*, 96:191–195, 1954.
[78] L. Ying, A. Hertzmann, H. Biermann, and D. Zorin. Texture and Shape Synthesis on Surfaces. *Proc. EG W. Rend*, 2001.
[79] E. Zhang, K. Mischaikow, and G. Turk. Vector Field Design on Surfaces. *ACM Trans. Graph.*, 25(4):1294–1326, 2006.
[80] M. Zhang, J. Huang, X. Liu, and H. Bao. Wave-based anisotropic quadrangulation. *ACM Trans. Graph.*, 29(4), 2010.