N-Symmetry Direction Field Design

Nicolas Ray, Bruno Vallet, Wan-Chiu Li, Bruno Lévy

N-Symmetry Direction Field Design

Nicolas Ray
Bruno Vallet
Wan Chiu Li
Bruno Lévy
INRIA - ALICE

Many algorithms in computer graphics and geometry processing use two orthogonal smooth direction fields (unit tangent vector fields) defined over a surface. For instance, these direction fields are used in texture synthesis, in geometry processing or in non-photorealistic rendering to distribute and orient elements on the surface. Such direction fields can be designed in fundamentally different ways, according to the symmetry requested: inverting a direction or swapping two directions may be allowed or not.

Despite the advances realized in the last few years in the domain of geometry processing, a unified formalism is still lacking for the mathematical object that characterizes these generalized direction fields. As a consequence, existing direction field design algorithms are limited to use non-optimum local relaxation procedures.

In this paper, we formalize N-symmetry direction fields, a generalization of classical direction fields. We give a new definition of their singularities to explain how they relate with the topology of the surface. Namely, we provide an accessible demonstration of the Poincaré-Hopf theorem in the case of N-symmetry direction fields on 2-manifolds. Based on this theorem, we explain how to control the topology of N-symmetry direction fields on meshes. We demonstrate the validity and robustness of this formalism by deriving a highly efficient algorithm to design a smooth field interpolating user defined singularities and directions.

Categories and Subject Descriptors: I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*Color, shading, shadowing, and texture*; I.3.5 [**Computer Graphics**]: Computational Geometry and Object Modeling; G.1.6 [**Numerical Analysis**]: Optimization; J.6 [**Computer Aided Engineering**]:

General Terms: Algorithms

Additional Key Words and Phrases: Vector field design, N-symmetry direction field, topology

Author's address: N. Ray, B. Vallet, W.C. Li, B. Lévy, Project ALICE, INRIA, 54500 Vandoeuvre, France,
e-mail: {ray | vallet | wan-chiu.li | levy }@loria.fr

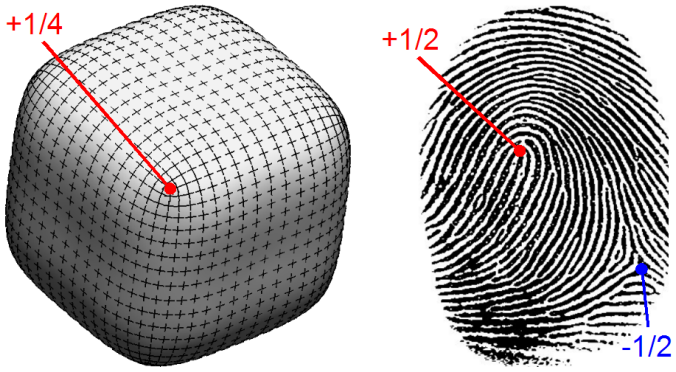

Fig. 1. Examples of singularities with fractionary index. **Left**: Singularity of index 1/4 of a natural smooth cross field on a cube like object. **Right**: Singularities of index 1/2 and -1/2 on a fingerprint.

1. INTRODUCTION

In Computer Graphics, a wide class of applications requires a smooth direction field over a surface to be defined. For instance, such a direction field was used in [Hertzmann and Zorin 2000] to place hatch strokes in NPR applications, in [Turk 2001] to steer the orientation of features in texture synthesis, or in [Alliez et al. 2003] to remesh a surface with cells aligned with a pair of direction fields. A closer look reveals that these applications use objects of higher symmetry than simple direction fields, i.e. objects invariant by rotation of π or $\pi/2$. For instance, the "natural" result one would expect when designing two smooth orthogonal direction fields on the cube-like object shown on Figure 1(left) connects the two direction fields around the corners of the cube with a $\pi/2$ rotation. To the best of our knowledge, this problem was first mentioned in the Computer Graphics field by Hertzmann and Zorin in their non-photorealistic paper [Hertzmann and Zorin 2000], where they intuitively introduce the notion of "cross-field". They mentioned that (1) characterizing singularities (points where a direction cannot be defined) and especially (2) controlling their number, type and placement are key issues and interesting mathematical problems involving an analog of Euler formula.

A solution to issue (1) appears in [Li et al. 2006] which proposes a data structure for direction fields based on controlling its behavior between discretization points through integers called "period jump" which characterized singularities. The authors then noticed that allowing period jumps to be multiple of $1/2$ or $1/4$ allowed to handle the direction symmetries, thus singularities of fractionary indices (see Figure 1). They used this to build an interpolation scheme for such symmetric fields, allowing to visualize arbitrary singularities, but examples had to be hand made or converted from an existing vector field (in analytic form or in a classical data structure).

This paper provides a solution to the much more complex issue (2). We start by formalizing N-symmetry directions in the continuous setting as the sets of directions invariant by $2\pi/N$ rotations, such that 1,2 and 4-symmetry direction fields correspond to classical direction fields, line fields and cross fields. We then provide a comprehensive study of their topology, and link it with the data structure and notion of period jumps from [Li et al. 2006]. This

allows finally to derive an efficient algorithm to generate a smooth N-symmetry direction field that interpolates a user-defined set of singularities (number, type and placement), as well as a set of geometric constraints on directions.

Before entering the heart of the matter, we give an overview of the previous work in direction field processing.

Previous work

Our work explains how to design smooth N-symmetry direction fields on mesh surfaces under both geometric and topological constraints, based on the data structure presented in [Li et al. 2006]. This problem has already been tackled for several applicative purposes:

—**non-photorealistic rendering:** In [Hertzmann and Zorin 2000], a smooth cross field is constructed to place strokes for a non-photorealistic rendering application. The construction method shares some common points with ours: the variables used to represent the directions are angles measured relative to a given arbitrary direction field. This paper follows the possibilities of future work and mathematical investigations suggested at the end of their paper.

—**texture synthesis:** In *lapped textures* [Praun et al. 2000] and *texture particles* [Dichler et al. 2002], a direction field is used to control the orientations and sizes of texture patches distributed over the object. In [Praun et al. 2000], a smooth direction field is generated through radial basis functions, with geodesic distances computed over the surface. The method presented in [Turk 2001] operates at texel level, by using the direction field to steer the anisotropy of a texture synthesizer. The direction smoothing procedure they used is inspired by [Gortler et al. 1996], which is based on a multi-resolution Laplacian smoother. A similar procedure is described in [Ohtake et al. 2001], with the addition of non-linear weights that preserve important direction field discontinuities. In [Wei and Levoy 2001], 2 and 4-symmetry direction fields are used to steer synthesizing using 2 and 4-symmetry texture samples. Finally, [Zelinka and Garland 2004] steer their texture generation method using a direction field defined as the gradient of a fair Morse function (it has the same singular points as the function). Based on the study of the Morse complex of smooth harmonic functions [Ni et al. 2004], this allows the user to control the number and configuration of singularities.

—**remeshing and global parameterization:** In [Alliez et al. 2003], a method is proposed to generate quadrangles aligned with two orthogonal direction fields, obtained by smoothing an estimation of the curvature tensor. The refinement presented in [Marinov and Kobbelt 2004] operates without a global parameterization and uses, in a certain sense, an explicit version of Ohtake et al.'s non-linear weights [2001] to preserve important features. More recently, [Ray et al. 2006] presented a quad remeshing method based on a periodic global parameterization following a direction field. The global parameterization method presented in [Gu and Yau 2003] also follows a direction field defined through a pair of holomorphic functions.

Our paper shares some common points with [Zelinka and Garland 2004] and [Gu and Yau 2003], in particular, the ability of controlling singularities. The main difference is that in these two methods, the direction field is defined to be the gradient of a scalar field, so it is

necessarily curl-free. Moreover, we handle symmetry so we can represent a wider class of singularities, with *arbitrary* indices.

The works that handle symmetries usually focus on the specific case of "cross-fields" $\vec{v}_1, \vec{v}_2 = R_{\vec{n}}(\vec{v}_1, \pi/2)$, where $\vec{v}_1$ and $\vec{v}_2$ can be swapped (see the cube-like object of Figure 1), most of the smoothing algorithms mentioned above use a *local* relaxation procedure, updating values $\vec{v}_1(p_i)$ and $\vec{v}_2(p_i)$ *vertex by vertex*. During these computations, $\vec{v}_1(p_i)$ will be influenced either by $\vec{v}_1(p_j)$ or $\vec{v}_2(p_j)$. The algorithm chooses the direction nearest to $\vec{v}_1(p_i)$. As a consequence, singularities may appear without control, and the convergence is slow. In contrast, based on a combined topological analysis of both the surface and the direction field, we derive a *global* formulation of the problem as a quadratic form minimization, yielding a more efficient optimization procedure (conjugate gradient), which computes directly an optimal solution.

Some more recent papers directly address problems related with direction field processing. For instance, in [Polthier and Preuss 2002] and [Tong et al. 2003] a procedure for computing the Hodge decomposition of a direction field is described. This decomposition isolates some features of the field, and makes it possible to filter or to enhance them. In [Tricoche et al. 2003], a method is presented to simplify the topology of symmetric, second order 2D direction fields. Recent work on vector field design include [Wang et al. 2006] where subdivision schemes are used to define bases for discrete differential 0- and 2-forms, and [Fisher et al. 2007] that uses Whitney forms and the Hodge decomposition. Both approaches allow to design vector fields under user constraints, but do not provide any guarantee that new singularities will not appear, nor tackle the problem of symmetric fields. Finally, a complete toolkit for interactive vector field design based on moving singularities and pair cancellation is presented in [Zhang et al. 2004], and generalized to tensors in [Zhang et al. 2005], then to N-rotational symmetry fields (N-RoSy) in [Palacios and Zhang 2007]. As in this paper, [Palacios and Zhang 2007] describes a tool to design a smooth N-symmetry direction field from a user-defined set of singularities. However, they start by generating a base field from the user constraints without controlling the topology, then "repair" the field *a posteriori* by moving or removing (by pair cancellation) the undesired singularities that appeared. In contrast, we ensure topological constraints in the design process *a priori*, by construction.

Contributions

—We provide a comprehensive study of the topology of N-symmetry direction fields by introducing the concept of *turning number*. This generalizes the notion of singularity *index* (see Sections 2.5 and 2.6);

—We prove an analog of the Poincaré-Hopf theorem based on turning numbers, adapt it to N-symmetry direction fields, and generalize it to manifolds with borders. We also prove a theorem relating turning numbers to direction field topology. With both theorems we can fully explain the links between the topology of a surface and the topology of a direction field defined over it. (see Appendix A);

—We apply this theory to the period jumps based structure presented in [Li et al. 2006] for representing direction fields. Thus we explain how period jumps control the topology of the direction field (see Section 3);

Fig. 2. Homotopy for direction fields: the middle field is homotopic to the one on the right but not to the one on the left (they can not be transformed continuously one into another)

—We describe an algorithm for constrained N-symmetry direction field design. From a user-defined set of singularities and an optional set of points with fixed directions, our algorithm constructs a smooth direction field that satisfies all the constraints. If the indices of the user-defined singularities sum up to the Euler characteristic $(2g-2)$, the construction is guaranteed to generate no other singularity (see Section 5). To our knowledge, this is the first algorithm that performs provably correct direction field design.

2. DIRECTION FIELDS ON SURFACES WITH BORDERS

This section presents the fundamental tools for studying direction fields defined over surfaces with borders. more specifically, we focus on their topology. Topology is the study of properties that are invariant by continuous deformations (without cutting or gluing anything), called homotopies. In other words, topology tries to answer the following question: under which conditions are two objects homotopic (i.e. can be continuously transformed one into another) ? For oriented surfaces with borders, the answer is that they need to have the same genus g (number of handles) and number of borders b. In other terms, if $\mathscr{S}_1$ and $\mathscr{S}_2$ are two surfaces with borders, we have $S_1 \equiv_t S_2 \Leftrightarrow g(S_1) = g(S_2)$ and $b(S_1) = b(S_2)$, where $\equiv_t$ denotes the homotopic equivalence (t stands for topology). What is even more interesting is the structure of the set of homotopy classes (classes of all objects with same topology). For oriented surfaces with borders again, this set is isomorphic to $\mathbb{N}^2$, since any pair of non negative integers (g,b) can be associated to the class of all surfaces with genus g and b borders. For this reason, we say that these two integers are the two *topological degrees of freedom* (TDoF) of surfaces. This section addresses the same questions for N-symmetry direction fields defined over a 2-manifold, which means that we aim at studying under which conditions two N-symmetry direction fields are homotopic. In other words, we want to exhibit the TDoF of these direction fields. To answer these questions, we will introduce the concept of turning numbers of cycles in direction fields. As we will show, the TDoF of a direction field are the turning numbers of a homology basis of cycles, such that these turning numbers capture all the topological structure of the field. Moreover, turning numbers generalize singularities, and we notice that singularities alone do not control all of the field topology.

The intuitive idea behind turning numbers is that when following a closed cycle, a N-symmetry direction might do an arbitrary number of N^{th} of turns before coming back to its starting point. Imagine you are traveling on earth along a cycle with a compass giving you the north direction, then you can count the number of turns of the compass while

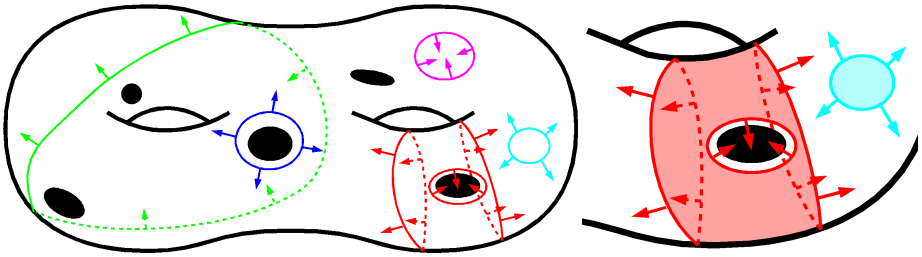

Fig. 3. **Left**: Cycles on a surface with borders (black areas are holes). Only the cycles shown on **Right** are exact because they are the boundaries (with conormal pointing outwards) of submanifolds of $\mathscr{S}$ (in light color).

following the cycle. If you turned around a tree (there are no trees at earth poles...) you will get one turn, but if you followed the equator or turned around a pole, you will get zero turns. We can do the same on any surface, and with any direction field defined over it. We call this quantity the "turning number" of the field along the cycle (it depends on both the cycle and the direction field). We will show that turning numbers capture the topology of the direction field, because two direction fields are homotopic iff they have the same turning numbers along the $2g+b-1$ cycles of a homology basis (a basis for cycles on a 2 manifold). This shows that homotopy classes of direction fields are isomorphic to $\mathbb{Z}^{2g+b-1}$, or in other words that the topology of a direction field is entirely defined by $2g+b-1$ integers. These integers are the TDoF (topological degrees of freedom) of direction fields.

In this section, we will first give the definitions of the objects that we will be using throughout this paper: surfaces (Section 2.1), cycles (Section 2.2) and direction fields (Section 2.3). Then we define the curvature for cycles and direction fields. We use these curvatures to formally define turning numbers, and exhibit their fundamental properties. We finally explain how turning numbers relate to field singularities.

2.1 Surfaces with borders

In this paper, we call surface (or 2-manifold) $\mathscr{S}$ a topological space where each point has a neighborhood homeomorphic to the plane or half plane. The surface $\mathscr{S}$ considered here is compact, connected and oriented, so that each point has a unique unit normal vector $\vec{n}$. The surface $\mathscr{S}$ has g handles and a boundary $\partial\mathscr{S}$ which is the collection of its b borders, for some (non-negative) integers g called the *genus* and b called the *number of borders* of $\mathscr{S}$.

2.2 Chains and Cycles on $\mathscr{S}$

A *chain* γ on $\mathscr{S}$ is an oriented 1-manifold embedded in $\mathscr{S}$ ($\gamma \subset \mathscr{S}$). Because γ is oriented, it has a unique tangent vector $\vec{t}_\gamma$ in each point which is also tangent to $\mathscr{S}$. Using this tangent vector, along with the surface normal $\vec{n}$, we can define a unique conormal vector $\vec{n}_\gamma = \vec{n} \times \vec{t}_\gamma$ on the chain, which ensures that $(\vec{t}_\gamma, \vec{n}_\gamma, \vec{n})$ form a natural local orthonormal basis called the *Darboux frame* (see Figure 4). Notice that chains are not necessarily connected, so the term "set of chains" would be more appropriate (but more cumbersome). We define the following notions on chains (see Figure 3 left):

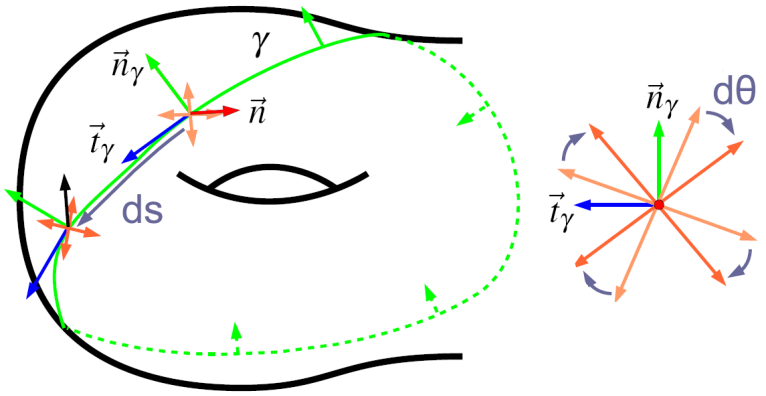

Fig. 4. A Darboux frame on a chain (γ, green) consists of the tangent ($\vec{t}_\gamma$, blue), conormal ($\vec{n}_\gamma$, green) and normal ($\vec{n}$, red). When moving along an infinitesimal portion ds of γ, the direction turns of an infinitesimal angle $d\theta$ in the Darboux Frame.

—The *reversal* $-\gamma$ of a chain γ is the chain with opposite orientation: $\vec{t}_{-\gamma} = -\vec{t}_\gamma$. We use the notations $\gamma_0 + \gamma_1$ and $\gamma_0 - \gamma_1$ in place of $\gamma_0 \cup \gamma_1$ and $\gamma_0 \cup -\gamma_1$ as it is more practical to handle unions of cycles with various orientations.

—A *cycle* on $\mathscr{S}$ is a chain without endpoints ($\partial\gamma = 0$).

—We call ∂ the boundary operator, such that $\partial\mathscr{S}$ is the subset of points of $\mathscr{S}$ with neighborhoods homeomorphic to the half plane. This subset is a cycle for which we can choose an orientation by requiring its conormal to point outwards $\mathscr{S}$ (this is why we represent the conormal instead of tangent in our figures).

—A chain γ is called *exact* if there exist a sub-manifold S of $\mathscr{S}$ such that $\gamma = \partial S$. An exact chain is necessarily a cycle (a boundary cannot have endpoints).

—Two cycles γ_0 and γ_1 are *homological* iff $\gamma_0 - \gamma_1$ is exact (see Figure 5). We will denote $\equiv_l$ the homologic equivalence (for cycles), when $\equiv_t$ is the homotopic equivalence (for surfaces and direction fields). More formally, the homology for cycles is defined as the quotient set of closed chains (cycles) over exact chains. In particular the 0 of homology is the class of exact cycles, so γ is exact $\Leftrightarrow \gamma \equiv_l \emptyset$

—$H(\mathscr{S}) = \{\gamma_i^H\}_{i=1..n}$ is called a homology basis on $\mathscr{S}$ if:

—Linear independence: $\sum a_i\gamma_i^H \equiv_l \emptyset \Leftrightarrow a_1 = ... = a_n = 0$

—Spanning: $\forall\gamma\ \exists a \in \mathbb{Z}^n$ such that $\gamma \equiv_l \sum_{i=1}^n a_i\gamma_i^H$ (any cycle on $\mathscr{S}$ is homological to a formal sum of the basis cycles).

An important result from manifold topology is that a homology basis on $\mathscr{S}$ has $2g+b-1$ basis vectors, where g is the genus, and b the number of borders of $\mathscr{S}$.

2.3 Direction field

A unit tangent vector $\vec{u}$ on $\mathscr{S}$ is a vector that satisfies $||\vec{u}|| = 1$ and $\vec{u}\cdot\vec{n} = 0$. We call N-symmetry direction on $\mathscr{S}$ a set of N unit tangent vectors on $\mathscr{S}$ invariant by rotation

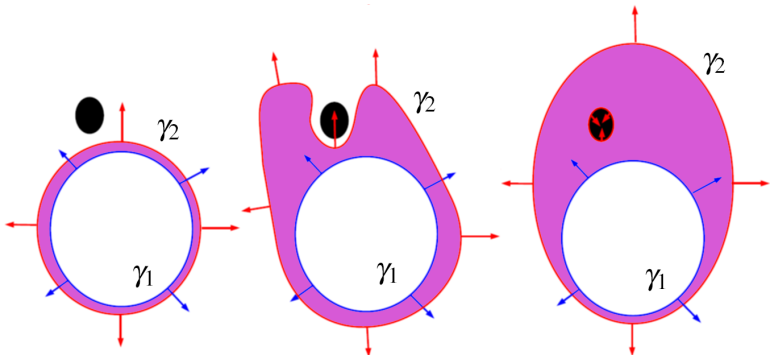

Fig. 5. γ_2 stays homological to γ_1 because $\gamma_2 - \gamma_1$ is exact(it is the border of a 2-manifold). Note that homological cycles might have a different number of connected components.

of $2\pi/N$ around the normal. Hence, based on a unit tangent vector $\vec{u}_0$ we can build a N-symmetry direction $\vec{d} = \{\vec{u}_k = R_{\vec{n}}(\vec{u}_0, 2k\pi/N) | k \in \mathbb{Z}\}$ where $R_{\vec{n}}$ is the rotation around $\vec{n}$. We call direction field $\vec{d}$ on $\mathscr{S}$ a mapping that associates a N-symmetry direction $\vec{d}(P)$ to each point $P \in \mathscr{S}$. In what follows, we will omit the term N-symmetry for brevity.

We call $\mathscr{D}_N(\mathscr{S})$ the set of differentiable direction fields, which are direction fields which curvature (see definition below) is defined almost everywhere. Two direction fields $\vec{d}_0$, $\vec{d}_1 \in \mathscr{D}_N(\mathscr{S})$ are called homotopic if there exists a continuous function $\Gamma : [0,1] \to \mathscr{D}_N(\mathscr{S})$ such that $\Gamma(0) = \vec{d}_0$, $\Gamma(1) = \vec{d}_1$. In other terms, two direction fields are homotopic (or have the same topology) if they can be continuously transformed one into another. As this is a relation of equivalence, we can define homotopy classes as the sets of all direction fields equivalent to a given one. Based on the intuition of what is a direction field in the plane, it would seem that all direction fields are homotopic, then there would be only one homotopy class. However, if the direction field is defined on a surface of arbitrary topology, we will show that it is no longer the case. Homotopy classes of direction fields can be characterized by the *turning numbers* of the field along some cycles. They correspond intuitively to the number of times the direction turns in a local Darboux frame while moving along the cycle. We are now going to define the notion of curvature for both cycles and direction fields, which will be used for a rigorous definition of the turning number.

2.4 Curvature

The curvature of a chain expresses the angular variation of its tangent along the chain. If we call s the arclength on a chain γ, we can define the curvature of γ using the decomposition:

$$\frac{\partial \vec{t}_\gamma}{\partial s} = \kappa_\gamma \vec{n}_\gamma + \kappa_{\mathscr{S}} \vec{n} \qquad \kappa_{\mathscr{S}} = \frac{\partial \vec{t}_\gamma}{\partial s} \cdot \vec{n} \qquad \kappa_\gamma = \frac{\partial \vec{t}_\gamma}{\partial s} \cdot \vec{n}_\gamma \tag{1}$$

—$\kappa_{\mathscr{S}}$ measures the normal curvature of $\mathscr{S}$ in direction $\vec{t}_\gamma$.

—κ_γ measures the curvature of γ in the tangent plane of $\mathscr{S}$.

—$\partial \vec{t}_\gamma / \partial s \cdot \vec{t}_\gamma = 0$ by derivation of $\vec{t}_\gamma \cdot \vec{t}_\gamma = 1$ ($\vec{t}_\gamma$ is a unit vector)

We can similarily define the curvature $\kappa_{\vec{d}}(\vec{t}_\gamma)$ for the direction field $\vec{d} = \{\vec{u}_k\}$ in direction $\vec{t}_\gamma$ as:

$$\kappa_{\vec{d}}(\vec{t}_\gamma) = \frac{\partial \vec{u}_k}{\partial s} \cdot \vec{u}_k^\perp \tag{2}$$

where $\vec{u}_k^\perp = \vec{n} \times \vec{u}_k$ is a unit vector orthogonal to $\vec{u}_k$ in the tangent plane, such that $(\vec{u}_k, \vec{u}_k^\perp, \vec{n})$ is an orthonormal basis (it is the Darboux frame of the streamlines of $\vec{u}_k$). The curvature $\kappa_{\vec{d}}$ is the same for all $k \in \mathbb{N}$ so it can be called the curvature of $\vec{d}$ in direction $\vec{t}_\gamma$. The quantity $(\kappa_{\vec{d}} - \kappa_\gamma)ds$ can be interpreted as the infinitesimal angular variation $d\theta$ of the direction field in the Darboux frame when moving along the cycle (see Figure 4). We are now going to integrate this variation to define the turning numbers.

2.5 Turning number

The turning number of a direction field along a cycle corresponds to the number of rotations of the field along this cycle. Turning numbers are characteristic of homotopy classes of direction fields, hence of their topology.

Let γ_{PQ} be a chain with endpoints P and Q, and $\theta_\gamma(P)$ be an angle (in $\mathbb{R}$) between $\vec{d}(P)$ and $\vec{t}_\gamma(P)$. Then:

$$\theta_\gamma(Q) = \theta_\gamma(P) + \int_{\gamma_{PQ}} d\theta = \theta_\gamma(P) + \int_{\gamma_{PQ}} (\kappa_{\vec{d}} - \kappa_\gamma)ds \tag{3}$$

is an angle (in $\mathbb{R}$) between $\vec{d}(Q)$ and $\vec{t}_\gamma(Q)$. Especially, for a cycle γ, we have $P = Q$, so we will have equality of the angles modulo $2\pi/N$. Therefore, the quantity:

$$T_{\vec{d}}(\gamma) = \frac{1}{2\pi} \oint_\gamma (\kappa_{\vec{d}} - \kappa_\gamma)ds \tag{4}$$

is necessarily an integer multiple of $1/N$ which we call the *turning number* of $\vec{d}$ along γ, and corresponds intuitively to the number of N^{th} of turns $\vec{d}$ does in the Darboux frame along γ (see Figure 6).

In the case $T_{\vec{d}}(\gamma) = 0$, there exists a continuous scalar function θ defined on γ such that $(\kappa_{\vec{d}} - \kappa_\gamma)ds = d\theta$. The scalar θ is defined up to a constant which can be chosen such that θ is an angle between $\vec{t}_\gamma$ and $\vec{d}$. In the case $T_{\vec{d}}(\gamma) \neq 0$, θ cannot be defined in a way to be continuous (we need $2\pi/N$ jumps).

Turning numbers have two fundamental properties (see Appendix A for a proof) which make them useful for studying direction field topology:

THEOREM 2.1 BOUNDARY TURNING NUMBER. *Let S be a surface (2-manifold with borders) embedded in $\mathbb{R}^3$, then:*

$$T_{\vec{d}}(\partial S) = -\chi(S) \tag{5}$$

where $\chi(S) = 2 - 2g(S) - b(S)$ is the Euler characteristic of S.

Theorem 2.1 (Boundary turning number) is equivalent to the Poincaré Hopf theorem with a proper definition for the index of a singularity, which will be developed in next subsection. It links the turning numbers of homological cycles as $\gamma_1 \equiv_l \gamma_2 \Leftrightarrow \exists S$ such that $\gamma_1 - \gamma 2 = \partial S$

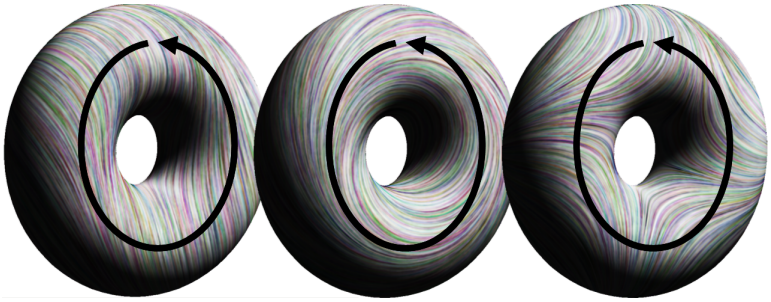

Fig. 6. The turning number associated to a generator defines topology that cannot be captured by singularity indices. The difference of topology of theses direction fields without any singular points is defined by the turning numbers of the generator (black cycle) which are respectively -1, 0 and 1

THEOREM 2.2 TOPOLOGICAL EQUIVALENCE. *Two direction fields defined over a surface $\mathscr{S}$ are homotopic iff they have the same turning numbers along the cycles of a homology basis $H(\mathscr{S})$ of $\mathscr{S}$:*

$$\vec{d}_1 \equiv_t \vec{d}_2 \Leftrightarrow \forall \gamma \in H(\mathscr{S}), T_{\vec{d}_1}(\gamma) = T_{\vec{d}_2}(\gamma)$$

Theorem 2.2 shows that direction fields have homotopy classes isomorphic to $\mathbb{Z}^{2g+b-1}$. This comes from the fact that the $2g$ generators of the surface and $b-1$ surface borders define a homology basis. Hence a direction field has $2g+b-1 = 1-\chi(\mathscr{S})$ topological degrees of freedom (TDoF) on a genus g surface with b borders. In the case $\chi(\mathscr{S}) = 2$ ($\mathscr{S}$ is a topological sphere), this number is negative because it is impossible to build a continuous direction field without introducing at least a singularity (a hole), which brings $\chi(\mathscr{S})$ to 1.

Both theorems are proved in Appendix A. With these two theorems, one can exhibit, understand and control all the TDoF of a direction field.

We are now going to explain how turning numbers generalize the usual notion of singularity index.

2.6 Singularities

Let $\vec{v}: \mathbb{R}^2 \to \mathbb{R}^2$ be a vector field. It is usually assumed that the zero set of $\vec{v}$: $\{x \in \mathbb{R}^2 | \vec{v}(x) = 0\}$ consists of a finite number of distinct points P_i which are called the singularities of $\vec{v}$. The singularities can then be classified by their index:

$$I_{\vec{d}}(P_i) = \frac{1}{2\pi} \int_{\partial\Omega(P_i)} d\theta \tag{6}$$

where $\Omega(P_i)$ is a small neighborhood of P containing no other singularities (zeros) of the vector field, and θ is the angle formed by the vector field and a reference vector. The singularity index is necessarily an integer that equals 1 for sources, vortices and sinks, -1 for saddles. An index of 0 corresponds to a degenerate singularity, which means the corresponding singularity can be removed without creating another singularity.

The notion of index applies naturally to direction fields and keeps the same definition (6) as it depends only on the direction of the vector. As singularities are zeros, the direction cannot be defined there, hence we identify singularities of a direction field with "holes" in its domain of definition,

For N-symmetry vector fields, indices can still be defined by (6), but they are now a multiple of $1/N$ because θ is defined modulo $2\pi/N$.

This 2D definition for direction field singularity index cannot be directly extended to surfaces embedded in $\mathbb{R}^3$ because it lacks a reference vector to define θ. However, it can be shown that in $\mathbb{R}^2$:

$$T_{\vec{d}}(\partial\Omega(P)) = I_{\vec{d}}(P) - 1 \tag{7}$$

As the turning number can be extended to surfaces in $\mathbb{R}^3$ because the reference vector is the tangent to the cycle, this allows us to extend the definition of the index of a singularity on a surface in $\mathbb{R}^3$ to:

$$I_{\vec{d}}(P) = 1 + T_{\vec{d}}(\partial\Omega(P)) \tag{8}$$

With this definition for the index, we can generalize the Poincaré-Hopf theorem to N-symmetry direction fields:

THEOREM 2.3 POINCARÉ-HOPF. *The sum of singularity indices on a closed surface $\mathscr{S}$ equals its Euler characteristic:*

$$\sum_{i=1}^{b} I_i = \chi(\mathscr{S}) = 2 - 2g$$

Proof: Notice first that this theorem holds for surfaces without borders, so the number of borders is absent in the Euler characteristic. Let us call P_i the point of index I_i and $\mathscr{S}_h = \mathscr{S} \backslash \{\Omega(P_i)\}$ (surface with borders in place of singularities), on which the field is continuous because it does not contain the singularities. Applying the boundary relation (5) to $\mathscr{S}_h$ yields:

$$T(\partial\mathscr{S}_h) = T(\partial\mathscr{S}) - \sum_{i=1}^{b} T(\partial\Omega(P_i)) = -\sum_{i=1}^{b}(I_i - 1) = -\chi(\mathscr{S}_h) = -\chi(\mathscr{S}) + b$$

we have $T(\partial\mathscr{S}) = 0$ because $\mathscr{S}$ has no borders, and we get the Poincaré-Hopf theorem by subtracting b in these equalities □

The idea of replacing singularities with holes allows singularities to be handled as borders. Moreover borders are topological objects that are very well understood through cycle homology. In what follows, we will use again this idea of replacing singularities with holes. We will then characterize the singularity by the behavior of the field along the border of the hole.

Notice that singularities do not capture all the topological degrees of freedom, since a homology basis also contains generators of the surface, that do not enclose any singularity (see Figure 6). This justifies why we need the concept of turning number to capture the direction field topology, and not only the singularity indices.

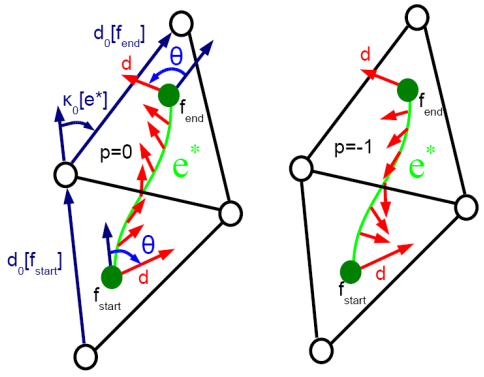

Fig. 7. p solves the ambiguity for equal angles with the reference direction $\vec{d}_0$ at endpoints of a dual edge e^*.

We will now explain how the turning numbers of a direction field in the data structure from [Li et al. 2006] depend only on the period jumps. As we have proved in this section that the topology of a direction field is completely defined by its turning number, this means we will have a total control over the direction field topology through period jumps.

3. DISCRETE DIRECTION FIELD

We will now briefly recall the period jump based discrete representation for N-symmetry direction fields presented in [Li et al. 2006], and explain its coherence with the continuous theory presented above.

3.1 Period jumps based discretization

The first step of the discretization of a direction field $\vec{d}$ such as described in [Li et al. 2006] is to sample it at the center of each facet. This is done by choosing a reference direction $\vec{d}_0$ in each triangle (for instance one of the oriented edges of the triangle), and defining $\vec{d}$ by the angle θ if forms with $\vec{d}_0$.

The discrete curvature of $\vec{d}$ can be defined as the integral of the continuous curvature along each dual edge e^*, which is given by applying (3) to e^*:

$$\kappa_{\vec{d}} = \kappa_0 + \theta(f_{end}) - \theta(f_{start}) + \frac{2\pi p}{N} \tag{9}$$

where the curvature of the base field κ_0 is defined as the oriented angle in $(-\pi, \pi]$ between $\vec{d}_0(f_{start})$ and $\vec{d}_0(f_{end})$, and p is an integer called period jump that corresponds to the ambiguity in the interpolation from $\vec{d}_0(f_{start})$ to $\vec{d}_0(f_{start})$ (see Figure 7). Once p is chosen, the ambiguity is solved, and the direction can be interpolated on the whole mesh (except at vertices) by the interpolation proposed in [Li et al. 2006]. Hence we can define any direction field $\vec{d}$ through:

(1) The choice of a reference direction $\vec{d}_0$ on each facet (i.e. an oriented edge of the facet),

(2) A scalar θ that represents the angle between $\vec{d}$ and $\vec{d}_0$ on each facet, such that $\vec{d} = R(\vec{d}_0, \theta)$ on dual vertices.

(3) An integer p that represents the behavior of $\vec{d}$ between discretization points, such that the discrete curvature of $\vec{d}$ is $\kappa_{\vec{d}} = \kappa_0 + d\theta + \frac{2\pi p}{N}$

Using the definition of the turning number (4) and the discrete curvature (9), we derive an expression for the turning number of a cycle γ in the discrete setting:

$$T_{\vec{d}}(\gamma) = \frac{1}{2\pi}\sum_{e^*\in\gamma} \kappa_{\vec{d}}(e^*) - \kappa_\gamma(e^*) = \frac{1}{2\pi}\sum_{e^*\in\gamma} \kappa_0(e^*) - \kappa_\gamma(e^*) + \frac{1}{N}\sum_{e^*\in\gamma} p(e^*) \tag{10}$$

We used the fact that the terms in $\theta(f_{end}) - \theta(f_{start})$ all cancel out because each θ appears once with each sign in the sum since γ is a cycle. As the first term is independent of the geometry θ of the field, this means that the turning numbers of $\vec{d}$ are defined entirely by the period jumps p. Finally, $\sum_\gamma \kappa_\gamma$ is computed as the sum of oriented angles between successive edges along all triangles crossed by γ. To compute the singularity index at a vertex v, we can simply choose $\gamma = \partial v^*$ (v^* is a dual facet), in which case $\sum_\gamma \kappa_\gamma$ is simply the sum of angles incident to v, and we get the expression for the index:

$$I_{\vec{d}}(v^*) = 1 + \frac{1}{2\pi}\sum_{e^*\in v^*} (\kappa_0 - \kappa_\gamma) + \sum_{e^*\in v^*} \frac{p(e^*)}{N} = I_0(v) + \sum_{e^*\in v^*} \frac{p(e^*)}{N} \tag{11}$$

We interpret $I_0(v)$ as the index of the base field, that is computed by writing it:

$$2\pi I_0(v) = \sum_{e^*\in v^*} \big(\kappa_0(e^*) - \kappa_\gamma(e^*)\big) + 2\pi = \sum_{e^*\in v^*} \kappa_0(e^*) + A_d(v) \tag{12}$$

where $A_d(v)$ is the angle defect at v.

This discretization will now allow us to derive a simple algorithm to compute the period jumps p such that the topology of the field matches user constraints.

4. ZIPPING ALGORITHM: CONSTRAINING SINGULARITY INDICES

When designing a direction field, we can choose the period jumps $p(e^*)$ on each edge. This completely defines the topology of the represented direction field. We propose in this section an algorithm that computes the period jumps such that the singularities and their indices are exactly the ones defined by a user. However, the period jumps control all TDoF, including those that do not correspond to singularities, such as turning numbers of generators (see Figure 6). We choose to let these TDoF free, and let our optimization play on these additional TDoF to find a smoother solution. As a result the algorithm will not explicitly compute all period jumps, but find some free period jumps and express all the other period jumps as a combination of the free ones.

4.1 Problem setting

The problem we tackle in this section is the following:

Given constrained singularity indices $I_c(v_i)$ at each interior vertex v_i of the mesh, find the integers p such that:

$$I_{\vec{d}}(v_i) = I_c(v_i) \quad \forall v_i \tag{13}$$

Using (11), the above expression can be rewritten as:

$$\sum_{e^* \in v^*} p(e^*) = N(I_c(v) - I_0(v)) = \Delta I(v) \tag{14}$$

in other words, we are trying to find integers p on oriented dual edges that sum up to $\Delta I(v)$ around each vertex v. We will now propose a greedy algorithm to compute a solution to (14) in linear time, which works by classifying edges into three sets.

4.2 Algorithm

As in the continuous setting, the homology basis $H(M)$ of M consists of $2g$ generators and $b-1$ borders (we will discuss later the case $b = 0$). As the turning numbers of the cycles in $H(M)$ do not correspond to singularities, we will leave the turning numbers of these cycles free. As the direction cannot be defined at vertices, the direction field is defined over M_h, which is the mesh M with holes at every interior vertices v_i allowing a singularity to occur at v_i. Each removed interior vertex v_i adds a cycle ∂v_i^* around it to the homology basis, except one of them that can be chosen arbitrarily (we call it v_{last}). Hence M_h has a homology basis $H(M_h) = H(M) \cup \bigcup_{v_i \neq v_{last}} \{\partial v_i^*\}$. The turning numbers along these additional cycles in $H(M_h)$ correspond to singularities located on interior vertices, so we want to control them explicitly. Hence, the basic idea in Zipping is to progressively classify the dual edges in three sets:

(1) The set $\mathscr{E}_0^*$ of *null* edges, whose associated period jumps we can set to 0 without constraining any turning number. Hence we have $p(e_0^*) = 0$

(2) The set $\mathscr{E}_{free}^*$ of *free* edges, whose associated period jumps correspond to the turning number of a cycle homological to a cycle in $H(M)$, hence corresponding to a TDoF but not to a singularity (see Figure 9). We choose the name *free* because these TDoF will be left free in the subsequent optimization. item The set $\mathscr{E}_{dep}^*$ of *dependent* edges (they may depend on free edges), whose associated period jumps constrain the turning numbers of cycles ∂v_i^* around a single vertex v_i, hence constraining the index of a the singularity at v_i.

Zipping is a greedy algorithm that computes this classification iteratively. Zipping also computes an expression of the period jumps of edges in $\mathscr{E}_{dep}^*$ as a function of the free period jumps of edges in $\mathscr{E}_{free}^*$, such that the indices satisfy their constrained value.

(1) Fill $\mathscr{E}_0^*$: While it is possible, add to $\mathscr{E}_0^*$ edges e_0^* that do not close any cycle, and set their period jumps to 0: $p_0(e_0^*) = 0$, $\mathbf{c}(e_0^*) = 0$ This is equivalent to building a spanning tree of the dual graph G^*.

(2) Fill $\mathscr{E}_{dep}^*$: While it is possible, add to $\mathscr{E}_{dep}^*$ edges that close dual cells (cycles around a single primal vertex) and compute the corresponding period jumps such that they satisfy (14).

(3) Add an edge to $\mathscr{E}_{free}^*$: any remaining edge necessarily closes a cycle (else it would have been added to $\mathscr{E}_0^*$) that does not enclose a single vertex (else it would have been added to $\mathscr{E}_{dep}^*$). Hence the corresponding cycle is homological to a border or to a generator of the mesh, which turning numbers we want to keep free, so any such edge

Algorithm 1 Zipping algorithm (see Figure 8)

Build a recovering tree $\mathscr{E}_0^*$ of G^* // *grow* black edges
$\forall e_0^* \in \mathscr{E}_0^*$ set $p_0(e_0^*) \leftarrow 0$, $\mathbf{c}(e_0^*) \leftarrow 0$
$\mathscr{E}_{dof}^* \leftarrow \mathscr{E}^* \backslash \mathscr{E}_0^* \quad \mathscr{V}_{zip} \leftarrow \mathscr{V}$
$i_{free} \leftarrow 0$ // Number of free variables found
while $\mathscr{E}_{dof}^* \neq \emptyset$ **do**
  while $\mathscr{V}_{zip} \neq \emptyset$ **do**
    Take $v_z \in \mathscr{V}_{zip}$ and remove it from $\mathscr{V}_{zip}$
    if $v_z \notin \partial M$ and $\exists!$ unset edge $e_z^* \in \partial v_z^*$ **then**
      // *zip* red (then blue) edges
      Move e_z^* from $\mathscr{E}_{dof}$ to $\mathscr{E}_{dep}$
      Set $p_0(e_z^*) \leftarrow \Delta I(v_z) - \sum_{e^* \in \partial v_z^* \backslash e_z^*} p_0(e^*)$,
      Set $\mathbf{c}(e_z^*) \leftarrow -\sum_{e^* \in \partial v_z^* \backslash e_z^*} \mathbf{c}(e^*)$,
      Add the face opposite to v_z^* across e_z^* to $\mathscr{V}_{zip}$
    end if
  end while
  // *free* green edge
  $i_{free} \leftarrow i_{free} + 1$
  Take $e_{free}^* \in \mathscr{E}_{dof}^*$ and move it to $\mathscr{E}_{free}^*$
  Set $p_0(e_{free}^*) \leftarrow 0$, $c_i(e_{free}^*) \leftarrow \delta_{i,i_{free}}$
  Add the 2 faces adjacent to e_{free}^* to $\mathscr{V}_{zip}$
end while

can be added to $\mathscr{E}_{free}^*$. These edges have trivially: $p_0(e_{free}^*) = 0$. The vector $\mathbf{c}(e_{free}^*)$ only contains zeros except for a 1 that corresponds to e_{free}^*.

(4) After freeing an edge, it becomes possible again to find edges that close dual cells, such that steps 2 and 3 may be iterated until no edge remain. The period jumps computed in step 2 will depend on the period jumps of the edges that have been freed in a step 3 (see Figure 9).

The complete Zipping algorithm is detailed in Algorithm 1. To handle period jumps dependencies, we store all period jumps under the form:

$$p(e^*) = p_0(e^*) + \mathbf{c}(e^*)\mathbf{p}_{free}$$

where p_0 is an integer, and $\mathbf{c}$ a vector of integers which size equals the number of edges already freed.

4.3 Topological issues

If M has no border, we have:

$$\sum_{v \in \mathscr{V}} \Delta I(v) = \sum_{v \in \mathscr{V}} \sum_{e^* \in v^*} p(e^*) = 0$$

as each edge is taken once in each direction. The base field singularities naturally satisfy the Poincaré-Hopf theorem 2.3 ($\sum_{\mathscr{V}} I_{\vec{d}_0}(v) = 2 - 2g$). Therefore, the index at the last

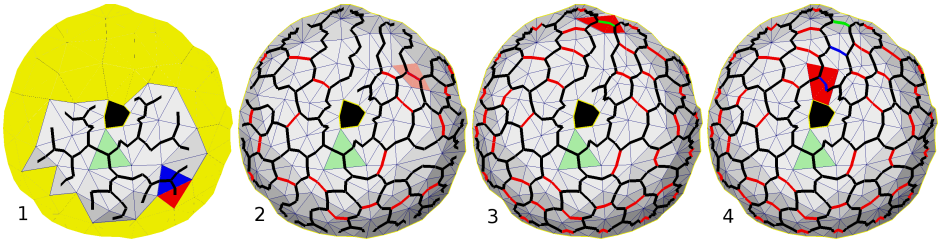

Fig. 8. Zipping: 1-**Grow** black edges (width first search) 2-**Zip** red edges 3-**Free** green edge 4-**Zip** blue edges (blue edges depend on the freed green edge)

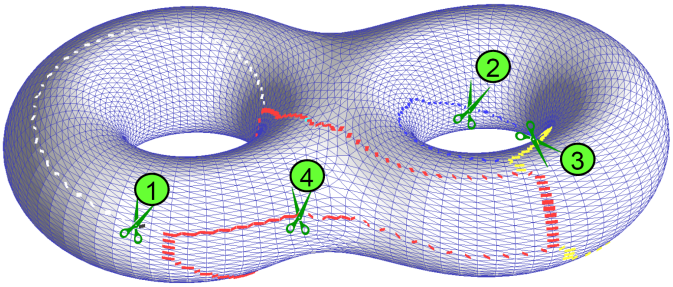

Fig. 9. A Genus g surface without borders requires $2g$ edges to be freed. The image shows in different colors the edges which period jumps effectively depend on free period jumps (scissors).

traversed vertex v_{last} satisfies:

$$\Delta I(v_{last}) = -\sum_{\mathscr{V}\setminus v_{last}} \Delta I(v) \quad , \quad I_c(v_{last}) = 2-2g-\sum_{\mathscr{V}\setminus v_{last}} I_c(v) \tag{15}$$

Hence the index of v_{last} adapts to ensure the Poincaré-Hopf theorem is satisfied. If the constraints on indices have been chosen such that:

$$\sum_{\mathscr{V}} I_c(v) = 2-2g$$

then the index $I_{\vec{d}}(v_{last})$ will be 0, which proves that no new singularity will appear. As the position of v^*_{last} depends on some choices arbitrarily made by the greedy algorithm, it cannot be easily determined, hence it is highly recommended to run the algorithm with constrained indices that sum up to $2-2g$ to avoid the appearance of a random (but necessary) singularity.

If M has borders, they are also handled by the Zipping algorithm, but their turning numbers are not constrained. In fact, we can also leave some vertices unconstrained by simply declaring them as border vertices before Zipping. However, in this case, we cannot guarantee that no undesired singularity appears, but only authorize where they can appear. A simple way to ensure that no singularity appear when M has borders, is to triangulate the holes of M before running our Zipping algorithm. The index of a border on such a field will be equal to the sum of the indices of the border vertices. In this case again, no singularity will appear iff the sum of indices of constrained vertices (including border vertices) sum up to $2-2g$.

5. APPLICATION TO DISCRETE DIRECTION FIELD DESIGN

We now show a simple use of our algorithm to design smooth direction fields with strong topological constraints.

5.1 Problem formulation

Using the notions above, the problem we want to solve is:

Given:

—a mesh $M =< \mathscr{V}, \mathscr{E}, \mathscr{F} >$

—a constrained singularity index $I_c(v)/N$ given on each vertex $v \in \mathscr{V}$.

—a constrained direction $\vec{d_c}(f)$ given on each facet of a subset $\mathscr{F}_c \subset \mathscr{F}$.

Interpolate the constraints as smoothly as possible. More formally, minimize:

$$E_{G^*}(\theta, p) = ||\kappa_{\vec{d}}||^2 = \sum_{e^* \in \mathscr{E}^*} \kappa_{\vec{d}}(e^*)^2 = \sum_{e^* \in \mathscr{E}^*} \left(\theta(f_{end}) - \theta(f_{start}) + \kappa_0(e^*) + \frac{2\pi p(e^*)}{N} \right)^2 \tag{16}$$

subject to the constraints:

$$I_{\vec{d}}(v) = I_{\vec{d}_0}(v) + \sum_{e^* \in \partial v^*} (e^*) = I_c(v) \quad \forall v \in \mathscr{V} \tag{17}$$

$$\vec{d}(f) = \vec{d_c}(f) \quad \forall f \in \mathscr{F}_c \tag{18}$$

As explained in the previous section, we build a representation for the direction field that implicitly enforces the topological constraints, by expressing all the period jumps $p(e^*)$ as a function of a limited number of free period jumps. We now describe a design algorithm that uses this structure to reduce the problem to a simple quadratic minimization procedure.

5.2 Algorithm

The design algorithm links between user inputs and Zipping inputs by computing the indices of the base field and making the difference with user constrained indices, and uses the output of the Zipping algorithm to ensure that the field topology is constrained during the creation of the direction field. The minimization problem is then solved in two passes:

(1) Minimize E_{G^*} with respect to θ and p by assuming that p is continuous

(2) Minimize with respect to θ only, with the p being set to their rounded value of the first pass

We use for both passes a standard formula [Levy 2005] to solve the problem of minimizing $([A_f, A_l][x_f, x_l]^t - B)^2$ where x_f are variables and x_l are set:

$$x_f = (A_f^t A_f)^{-1} A_f^t (B - A_l x_l) \tag{19}$$

This method is not guaranteed to find the global minimum with respect to the discrete variables $p(e^*)$, but offers good results in practice that satisfy all the constraints. For the

Algorithm 2 Design algorithm

(1) Set $I_c(v) = 0$ on $\mathscr{V} \setminus \mathscr{V}_c$ // no singularities outside $\mathscr{V}_c$
(2) Choose a direction $\vec{d_0}(f)$ on each facet // base field
(3) Compute $\kappa_0(e^*)$ as the angles between $\vec{d_0}(f_{end})$ and $\vec{d_0}(f_{start})$ for each dual edge.
(4) Compute the angle defects $A_d(v)$ at all vertices.
(5) Compute the base field indices $I_{\vec{d_0}}(v)$ using (12)
(6) Compute $\Delta I(v) = N(I_c(v) - I_{\vec{d_0}}(v))$
(7) Apply the Zipping Algorithm
(8) Build the linear system $[A_f, A_l, C][\theta_f, \theta_l, k_{free}]^t = B$ corresponding to (16). Each line of the system corresponds to an edge e^*: $[A_f, A_l]$ contains $+1$ and -1 at the indices corresponding to the θ at the 2 extremities of e^*, C contains $\mathbf{c}(e^*)$ corresponding to k_{free}. B contains the $\kappa_{\vec{d_0}}(e^*) - k_0(e^*)$.
(9) Pass 1:

$$[\theta_f^1, k_{free}^1]^t = ([A_f, C]^t [A_f, C])^{-1} [A_f, C]^t (B - A_l \theta_l)$$

(10) Pass 2:

$$[\theta_f^2]^t = (A_f^t A_f)^{-1} A_f^t (B - [C, A_l][rnd(k_{free}^1), \theta_l]^t)$$

where *rnd* is the rounding to the nearest integer value.
(11) Apply rotations θ_f^2 to $\vec{d_0}$ to get a direction $\vec{d}$ on each facet of M.

continuous variables $\theta(f)$, if at least one directional constraint is set, A_f is of maximal rank, therefore, by Gramm's theorem, $A_f^t A_f$ is non-degenerate and our algorithm finds the unique minimum. We also noticed that we could improve the visual aspect of the direction field near the constrained directions by adding a Laplacian smoothing term to the energy (16).

6. RESULTS

Our framework allows to create and edit direction fields on a surface via topological and geometric constraints. As illustrated in Figure 10, only few constraints are required to create the desired field. The same picture also demonstrates that our method is not affected by the complex geometry of the hairs of the David, whereas the previous methods only based on relaxation ([Ray et al. 2006] or [Hertzmann and Zorin 2000]) are trapped by this geometry and generate many singularities in this zone.

We have also tested this algorithm on large models from the Stanford 3D Scanning Repository called the statue and Lucy. These two models are treated in less than a minute using

	Lucy	The statue
genus	45	13
#borders	47	9
#triangles	125000	300000
time	12s	35s

Table I. Timings obtained on a Pentium IV 1.7Ghz

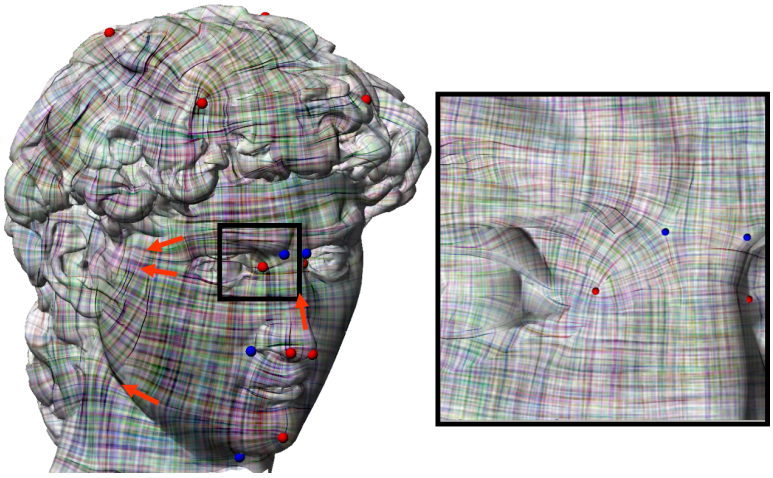

Fig. 10. Direction field edited by topological and geometric constraints. Red (resp. blue) dots are constrained singularities with positive (resp. negative) index. Orange arrows are directional constraints.

TAUCS for the quadratic minimization, and smooth fields are obtained (see Figure 11). Note that these two models have a complex topology (large genus g and number of borders b). The turning number around each border is used to counter the effect of each handle (a small handle is equivalent to a singular point of index 2).

Our algorithm also deals nicely with important constraints that can be applied on geometry (see the rotation constraints in Figure 12) and on topology (see Figure 6).

Finally, we compare with recent work in direction field design. Figure 13 shows that starting from a user defined set of constraints (singularities and directional constraints), our method generates a direction field with no other singularity (Right). To our knowledge, this is the first algorithm that achieves direction field design with provably correct singularity control. Figure 14 shows a more difficult configuration, successfully handled by our algorithm.

7. CONCLUSION

We have introduced a generalization of direction field that is general enough to manipulate direction fields with different kind of symmetries, which applications range from texture synthesis to quad remeshing.

As the representation we use separates the topology and the geometry of the direction fields, the singularities can be constrained through the topology by using a simple greedy algorithms, and the geometry can be smoothed by simply minimizing a quadratic form. This has been made possible by a complete analysis of the direction field topology including the characterization of the field behavior along cycles, and an extension of the Poincaré-Hopf theorem to rational indices. A new direction field smoothing algorithm has been presented to demonstrate the benefits of our structure. With our algorithm, one can for the first time provably control the placement of singularities. We are confident this new

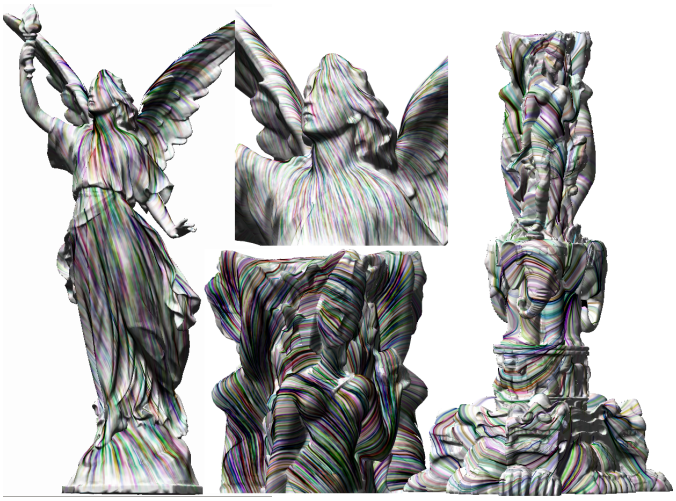

Fig. 11. Large models with many borders and high genus efficiently processed by our framework.

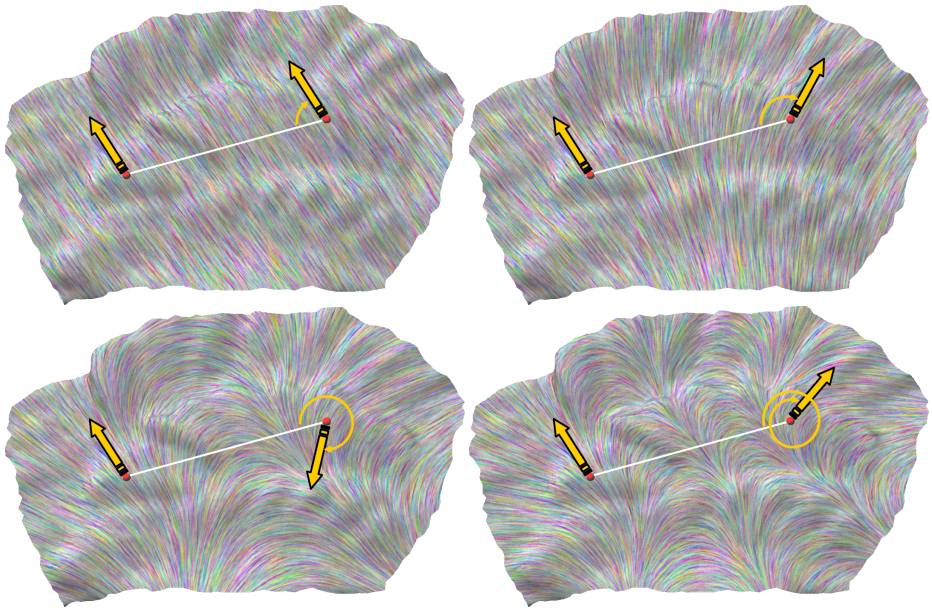

Fig. 12. Directional constraints are applied to a direction field. Notice that the direction is given by a rotation of the base field that can be greater than 2π.

way to consider and manipulate direction fields will set a clean and practical basis for a broad range of applications in geometry processing.

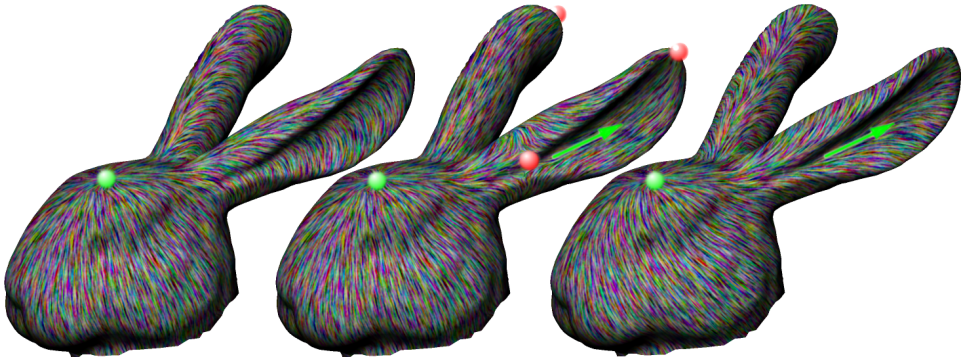

Fig. 13. **Left**: result of [Fischer et.al 2007] with a single singularity constraint (green dot); **Center**: result of [Fischer et.al 2007] with the same singularity constraint and directional constraint (green arrow), additional singularities appear (red dots); **Right**: our result with the same set of constraints. As predicted by our theorem, no additional singularity appears.

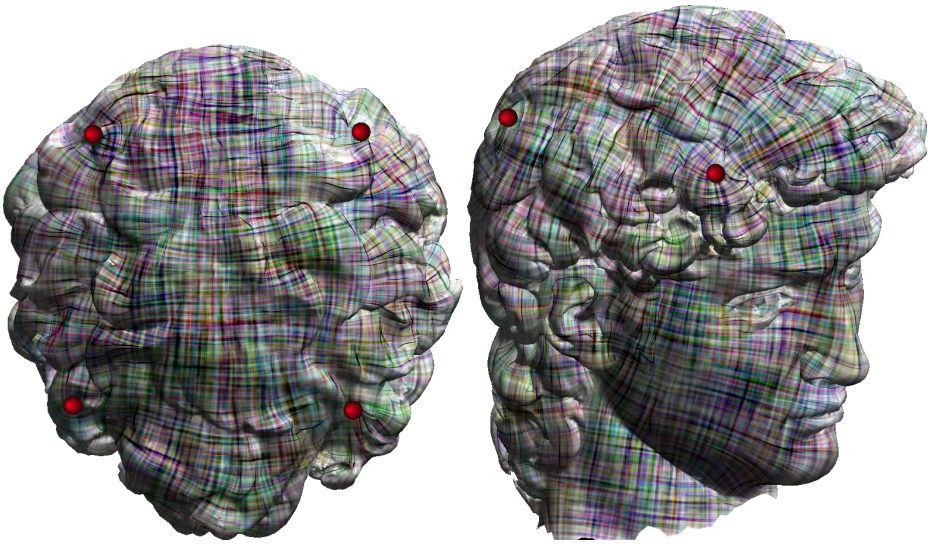

Fig. 14. A particularly challenging data set for direction field design algorithms. Our algorithm successfully satisfies the four user-defined singularities (green dots) and does not generate any additional singularity.

Acknowledgments

We want to thank the Associate Editor and the anonymous reviewers for their detailed reviews and numerous comments that helped us improve the quality of this paper. We also thank the Stanford Scanning Repository and Digital Michelangelo Project for the 3D models. This work was supported by Microsoft Research Europe (Geometric Intelligence Grant)

REFERENCES

ALLIEZ, P., COHEN-STEINER, D., DEVILLERS, O., LEVY, B., AND DESBRUN, M. 2003. Anisotropic polygonal remeshing. *ACM Transactions on Graphics (SIGGRAPH '03) 22,* 3, 485–493.

DICHLER, J.-M., MARITAUD, K., LEVY, B., AND CHAZANFARPOUR, D. 2002. Texture particles. *Computer Graphics Forum 21,* 3.

FISHER, M., SCHRODER, P., DESBRUN, M., AND HOPPE, H. 2007. Design of tangent vector fields. *ACM Transactions on Graphics (SIGGRAPH '07), to appear*.

GORTLER, S., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M.-F. 1996. The lumigraph. In *Proceedings of ACM SIGGRAPH*. 43–54.

GU, X. AND YAU, S.-T. 2003. Global conformal surface parameterization. In *Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*. 127–137.

HERTZMANN, A. AND ZORIN, D. 2000. Illustrating smooth surfaces. In *Proceedings of ACM SIGGRAPH*. 517–526.

LEVY, B. 2005. Numerical methods for digital geometry processing. In *Israel Korea Bi-National Conference - Invited talk, extended abstract*.

LI, W.-C., VALLET, B., RAY, N., AND LEVY, B. 2006. Representing higher-order singularities in vector fields on piecewise linear surfaces. *IEEE Transactions on Visualization and Computer Graphics (Proceedings Visualization '06)*.

MARINOV, M. AND KOBBELT, L. 2004. Direct anisotropic quad-dominant remeshing. In *Proceedings of Pacific Graphics*. 207–216.

NI, X., GARLAND, M., AND HART, J. C. 2004. Fair morse functions for extracting the topological structure of a surface mesh. *ACM Transactions on Graphics (SIGGRAPH '04) 23,* 3, 613–622.

OHTAKE, Y., HORIKAWA, M., AND BELYAEV, A. 2001. Adaptive smoothing tangential direction fields on polygonal surfaces. In *Proceedings of Pacific Graphics*. 189.

PALACIOS, J. AND ZHANG, E. 2007. Rotational symetry field design on surfaces. *ACM Transactions on Graphics (SIGGRAPH '07), to appear*.

POLTHIER, K. AND PREUSS, E. 2002. Identifying vector fields singularities using a discrete hodge decomposition. In *Visualization and Mathematics III*, H. Hege and K. Polthier, Eds. Springer Verlag, 113–134.

PRAUN, E., FINKELSTEIN, A., AND HOPPE, H. 2000. Lapped textures. In *Proceedings of ACM SIGGRAPH*. 465–470.

RAY, N., LI, W. C., LEVY, B., SHEFFER, A., AND ALLIEZ, P. 2006. Periodic global parameterization. *ACM Transactions on Graphics*.

TONG, Y., LOMBEYDA, S., HIRANI, A., AND DESBRUN, M. 2003. Discrete multiscale vector field decomposition. *ACM Transactions on Graphics (SIGGRAPH '03) 22,* 3, 445–452.

TRICOCHE, X., SCHEUERMANN, G., AND HAGEN, H. 2003. Topology simplification of symmetric, second-order 2D tensor fields. In *Hierarchical and Geometrical Methods in Scientific Visualization*.

TURK, G. 2001. Texture synthesis on surfaces. In *Proceedings of ACM SIGGRAPH*. 347–354.

WANG, K., WEIWEI, TONG, Y., DESBRUN, M., AND SCHRÖDER, P. 2006. Edge subdivision schemes and the construction of smooth vector fields. *ACM Transactions on Graphics (SIGGRAPH '06)*.

WEI, L.-Y. AND LEVOY, M. 2001. Texture synthesis over arbitrary manifold surfaces. In *Proceedings of ACM SIGGRAPH*. 355–360.

ZELINKA, S. AND GARLAND, M. 2004. Jump map-based interactive texture synthesis. *ACM Transactions on Graphics 23,* 4, 930–962.

ZHANG, E., HAYS, J., AND TURK, G. 2005. Interactive design and visualization of tensor fields on surfaces. Tech. rep., Oregon State University.

ZHANG, E., MISCHAIKOW, K., AND TURK, G. 2004. Vector field design on surfaces. *ACM Transactions on Graphics*.

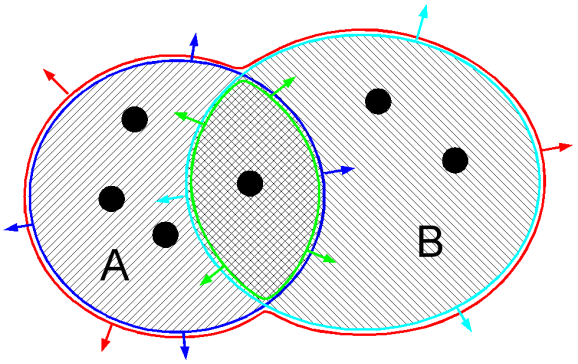

Fig. 15. Equivalence between $\partial A \cup \partial B$ and $\partial(A \cup B) \cup \partial(A \cap B)$

A. TURNING NUMBERS FUNDAMENTAL PROPERTIES

We give here an outline of the proof of two fundamental turning number properties.

A.1 Boundary property

The boundary property states that for any direction field on a 2-manifold S, we have $T(\partial S) = -\chi(S)$ which generalizes the Poincaré Hopf theorem to N-symmetry direction fields. We first establish three simple results, then prove the equation by structural induction.

LEMMA A.1. *The reversal of a cycle has opposite turning number:*

$$T(-\gamma) = -T(\gamma) \tag{20}$$

Proof: integrating in opposite direction changes ds in $-ds$ □

LEMMA A.2. *If we call A and B two sub-manifolds of $\mathscr{S}_h$ (see Figure 15), the turning numbers of ∂A and ∂B are linked by the equation:*

$$T(\partial A) + T(\partial B) = T(\partial(A \cup B)) + T(\partial(A \cap B)) \tag{21}$$

Proof: The term in $\kappa_{\vec{d}}$ is preserved because integrated on the same set $\partial A \cup \partial B = \partial(A \cup B) \cup \partial(A \cap B)$. The term in κ_γ, is preserved by application of the Gauss-Bonnet formula to the equality $\int_A + \int_B = \int_{A \cup B} + \int_{A \cap B}$ □

LEMMA A.3. *Boundaries of topological disks have a turning number of -1.*

Proof: Let γ^C be the boundary of a topological disk D. By definition, there exists a continuous bijective mapping $m : D \to \mathscr{D}_2$ where $\mathscr{D}_2$ is the unit disk in $\mathbb{R}^2$. Using this application, we can define the image field $\vec{d'} = m(\vec{d})$ of the direction field on $\mathscr{D}_2$. The turning number is preserved by m as it is continuous, so $T_{\vec{d}}(\partial D) = T_{\vec{d'}}(\partial \mathscr{D}_2)$. Because we are on the unit disk, the arclength s is equivalent to the angle on the disk boundary, so we have

$\vec{t}_\gamma = (\cos(s), \sin(s))^t$ and $\vec{n}_\gamma = (-\sin(s), \cos(s))^t$, which gives:

$$\kappa_\gamma = \frac{\partial \vec{t}_\gamma}{\partial s} \cdot \vec{n}_\gamma = 1$$

Let $\vec{x} = (x = \cos(s), y = \sin(s))^t$ be the position vector on the disk boundary. We have:

$$\kappa_{\vec{d}'}(\vec{t}_\gamma) = \frac{\partial \vec{d}'}{\partial s} \cdot \vec{d}'^{\perp} = \left(\frac{\partial \vec{d}'}{\partial x} \cdot \vec{d}'^{\perp}, \frac{\partial \vec{d}'}{\partial y} \cdot \vec{d}'^{\perp} \right)^t . \vec{t}_\gamma = \vec{\kappa}_{\vec{d}'} . \vec{t}_\gamma$$

one can verify that the rotational $\nabla \wedge \vec{\kappa}_{\vec{d}'} = 0$, so $\vec{\kappa}_{\vec{d}'}$ is a gradient, or in other words there exist a scalar field θ such that $\vec{\kappa}_{\vec{d}'} = \nabla\theta$. Thus its integral over a closed loop is zero, so the turning number evaluates to:

$$T_{\vec{d}'}(\partial \mathscr{D}_2) = \frac{1}{2\pi} \oint_{\partial \mathscr{D}_2} (\kappa_{\vec{d}'} - \kappa_\gamma) ds = \frac{1}{2\pi} \oint_{\partial \mathscr{D}_2} \nabla\theta . \vec{t}_\gamma ds - \frac{1}{2\pi} \int_{s=0}^{2\pi} 1 \cdot ds = -1 \square$$

These lemmas allow us to give a proof of Theorem 2.1:

$$T_{\vec{d}}(\partial S) = -\chi(S) \tag{22}$$

Proof: It is well known in topology that any orientable surface (with borders) can be cut along g cycles to obtain a sphere (with borders). Hence we start by proving (22) for a sphere with $b > 0$ borders.

If S is a topological disk ($b = 1, \chi = 1$), its boundary is contractible so its turning number is -1 by Lemma A.3 which proves (22) for $b = 1$.

If S is a topological cylinder ($b = 2, \chi = 0$), its boundary is composed of 2 borders: $\partial S = \gamma_1 + \gamma_2$. γ_1 and $-\gamma_2$ are homotopic so they have the same turning numbers, so by (20) we have:

$$T(\partial S) = T(\gamma_1) + T(\gamma_2) = T(\gamma_1) - T(-\gamma_2) = 0$$

which proves (22) for $b = 2$.

For higher numbers of borders, we prove the property by induction: assume (22) is true $\forall b \leq B$, and let S have $B+1$ borders $\gamma_1 ... \gamma_{B+1}$. Let A_B and A_{B+1} be two connected submanifolds which boundary contains γ_B and γ_{B+1} (respectively) and intersect in a topological disk. $S \backslash \{A_B \cup A_{B+1}\}$ has at most B borders, so it satisfies (22):

$$T(\partial(S \backslash \{A_B \cup A_{B+1}\})) = \sum_{i=1}^{B-1} T(\gamma_i) + T(\partial\{A_B \cup A_{B+1}\}) =$$

$$\sum_{i=1}^{B-1} T(\gamma_i) + T(\partial A_B) + T(\partial A_{B+1}) - T(\partial\{A_B \cap A_{B+1}\}) =$$

$$\sum_{i=1}^{B+1} T(\gamma_i) + 1 = \chi(S \backslash \{A_B \cup A_{B+1}\}) = -\chi(S) + 1$$

which proves (22) recursively.

We can now get to the most general case where S has genus g and b borders $\gamma_1 ... \gamma_b$. By definition of the genus, there exists a family $\gamma_1^G, ... \gamma_g^G$ of single cycles such that $S \backslash \{\gamma_1^G, ... \gamma_g^G\}$

is connected. This operation does not change the Euler characteristic of the surface ($b \leftarrow b+2g$, $g \leftarrow 0$ preserves $\chi = 2-2g-b$). The result is a sphere S_{cut} with $b+2g$ borders $\gamma_1, ...\gamma_b, \gamma_1^G, -\gamma_1^G, ...\gamma_g^G, -\gamma_g^G$, so we can apply (22):

$$T(\partial S_{cut}) = \sum_{i=1}^{n} T(\gamma_i) + \sum_{i=1}^{g} T(\gamma_i^G) + T(-\gamma_i^G) = 2-2g-b$$

as the second sum is null by (20). This finally proves Theorem 2.1 in the general case □

A.2 Topological equivalence

We will need a simple lemma to prove the topological equivalence:

LEMMA A.4.

$$T_{\vec{d}_1}(\gamma) = T_{\vec{d}_2}(\gamma) \quad \forall \text{ cycle } \gamma \subset \mathscr{S} \Leftrightarrow T_{\vec{d}_1}(\gamma) = T_{\vec{d}_2}(\gamma) \quad \forall \gamma \in H(\mathscr{S})$$

where $H(\mathscr{S})$ is a homology basis on $\mathscr{S}$. In other words, the turning numbers of all cycles on $\mathscr{S}$ depend only on the turning numbers of a homology basis.

Proof: By definition of homology basis, any cycle γ is homological to a cycle $\sum_i a_i \gamma_i^B$ where $\gamma_i^B \in H(\mathscr{S})$ are basis cycles, such that $\sum_i a_i \gamma_i^B - \gamma$ is a boundary. Hence, by Theorem 2.1, we have:

$$T_{\vec{d}_0}(\gamma) = \sum_i a_i T_{\vec{d}_0}(\gamma_i^B) + \chi(\mathscr{S}) = \sum_i a_i T_{\vec{d}_1}(\gamma_i^B) + \chi(\mathscr{S}) = T_{\vec{d}_1}(\gamma)$$

because the turning numbers are equal along the cycles of the homology basis. □

We can now prove the Theorem 2.2 which states that two direction fields are homotopic if and only if they have the same turning numbers along the cycles of a homology basis of $\mathscr{S}$.

Proof: The direct implication is immediate because turning numbers are preserved by homotopy. For the converse, let $\vec{d}_0$ and $\vec{d}_1$ be the two fields with the same turning numbers along the cycles of a homotopy basis. By Lemma A.4, $\vec{d}_0$ and $\vec{d}_1$ have the same turning numbers along any cycle γ:

$$T_{\vec{d}_0}(\gamma) - T_{\vec{d}_1}(\gamma) = \oint_\gamma (\kappa_\gamma - \kappa_{\vec{d}_0}(\vec{t}_\gamma))ds - \oint_\gamma (\kappa_\gamma - \kappa_{\vec{d}_1}(\vec{t}_\gamma))ds = \oint_\gamma (\kappa_{\vec{d}_1} - \kappa_{\vec{d}_0})(\vec{t}_\gamma)ds = 0$$

$(\kappa_{\vec{d}_1} - \kappa_{\vec{d}_0})(\vec{t}_\gamma)$ is a linear application which circulation along any closed curve is zero. This means that there exist a continuous scalar field θ that satisfies:

$$(\kappa_{\vec{d}_1} - \kappa_{\vec{d}_0})(\vec{t}_\gamma) = \nabla\theta.\vec{t}_\gamma$$

As θ is defined up to a constant, we can chose it to match the angle (in $[0, 2\pi[$ for instance) between $\vec{d}_0$ and $\vec{d}_1$. Thus we have $\vec{d}_1 = R(\vec{d}_0, \theta)$ where R is the rotation around the surface normal. As θ is continuous, we can build a continuous function $\Gamma : [0,1] \to \mathscr{D}_N(\mathscr{S})$ defined by $\Gamma(t) = R(\vec{d}_0, t\theta)$, such that $\Gamma(0) = \vec{d}_0$ and $\Gamma(1) = \vec{d}_1$ which proves that $\vec{d}_0$ and $\vec{d}_1$ are homotopic. □