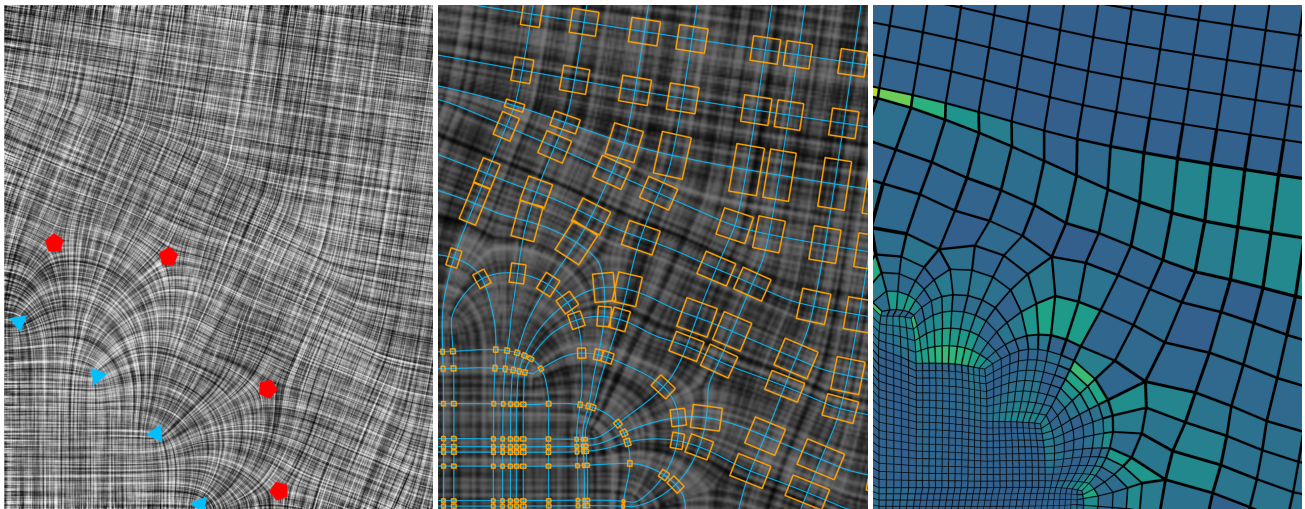# Anisotropy and Cross Fields

L. Simons[1] and N. Amenta[1]

[1]University of California, Davis, Department of Computer Science

Figure 1: *A cross field (left) in the plane, with singular points of valence 3 or 5 (highlighted). We assign two lengths at every point to define an anisotropic orthogonal frame field (visualized as an orange rectangle) at each point, and optimize to minimize the anisotropy (center). A quad mesh constructed with the cross field as a guide using an approach based on [CBK15] shows the effects of the anisotropy inherent in the cross field (right). Brighter shades indicate larger element anisotropy.*

Abstract
We consider a cross field, possibly with singular points of valence 3 or 5, in which all streamlines are finite, and either end on the boundary or form cycles. We show that we can always assign lengths to the two cross field directions to produce an anisotropic orthogonal frame field. There is a one-dimensional family of such length functions, and we optimize within this family so that the two lengths are everywhere as similar as possible. This gives a numerical bound on the minimal anisotropy of any quad mesh exactly following the input cross field. We also show how to remove some limit cycles.

CCS Concepts
• *Computing methodologies → Mesh geometry models; Shape analysis;*

1. Introduction

A *cross field* assigns two orthogonal unoriented directions to every point in a two-dimensional region. A common approach to constructing a quad-mesh is, roughly, to construct a cross field on the input domain, and then sample streamlines from the cross field to produce quads; but it rarely works out to be so easy!

Here is one fundamental issue with this approach. Within any regular infinitesimally small region we can always sample stream-lines to form a mesh in which the quads approach perfect squares as the density increases. But globally, the curvature of the cross field might force some quads to be arbitrarily anisotropic. This is well-known, in the sense that existing methods deal with it; for instance, Quantized Global Parameterization [CBK15] empirically chooses streamlines so as to minimize anisotropy (as in Figure 1). In this paper, we lay out the math describing how the curvature of a cross field may force any mesh directly derived from it to be

be anisotropic, and we give a numerical algorithm to measure the minimal anisotropy inherent in a given cross field.

We only handle a simple case. We assume we are given a smooth cross field, defined on an open or closed domain R in the plane (not necessarily a disk). We allow the cross field to have singular points, but only of valence 3 or 5 (see Figure 2). We require every streamline to have finite length, and we do not allow any streamline to both begin and end at a singular point. Thus, although we allow cycles, we exclude cross fields that contain limit cycles (Figure 2), and also cycles containing singular points.

Given our limitations on R, our conditions on the the input cross field are reasonable. Other kinds of singular points can be (and often are) simulated by clusters of singular points of valence 3 and valence 5. A streamline with two singular endpoints can be eliminated by a local perturbation of the cross field. Limit cycles cause infinite anisotropy and streamlines of infinite length, so we cannot say much about cross fields that contain them.

We show that a cross field meeting our conditions is always 2-*integrable*, by which we mean we can assign two lengths to the two orthogonal directions at each point, consistent with the curvature of the cross field. We explain intuitively what we mean by "consistent" via the following purely theoretical construction. For some large enough n, we select a discrete set of n streamlines that decomposes R into intrinsic rectangles, that is, four-sided regions whose curved sides are segments of streamlines or closed boundaries, so that the region sides meet at right angles, except at singular points. We increase n towards infinity, interpolating, if possible, our initial discrete set of streamlines with a smooth distribution of streamlines. In the limit, the streamlines divide R into infinitesimally small perfect rectangles with straight sides ("tiny bricks"). In this sense, R is the the infinite sum of a set of infinitesimal rectangles. Definition 1, below, is a more technical definition of integrability.

The infinitesimal limit rectangles may be arbitrarily anisotropic (that is, their two side lengths may differ). Some of the anisotropy may come from our choice of the smooth streamline distribution, and some may be forced by the curvature of the cross field. What distribution of streamlines minimizes the anisotropy, so that our infinitesimal rectangles are as square as possible? That is, how much anisotropy is forced by the cross field? This is the question we address.

The infinitesimal rectangles might all be perfect squares, in which case we only need one length function (known as the *conformal factor*); we say such a cross field is 1-*integrable*. It is widely understood that a cross field is 1-integrable if and only if it can be produced by a conformal map from the plane to itself (see Section 5). But not all cross fields can be produced in this way; Figures 1 and 7 show examples. In this more general case, we show that there is a always a one-dimensional space (up to scale) of ways to assign two length functions consistent with the curvature of the cross field.

Next, we describe how we optimize the length functions within the set of feasible assignments so as to minimize anisotropy. Given an input cross field obeying our conditions, we can thus compute and visualize the minimal anisotropy of its best infinitely fine "tiny bricks" mesh. We present some results on measuring the anisotropy of an input cross field.

Finally, we observe that if a limit cycle is simple - that is, if it does not cross itself - it can be removed without adding any new singularities.

We emphasize that producing a finite quad mesh at a reasonable resolution poses additional challenges beyond controlling anisotropy. For instance, the finite mesh must incorporate the separatrices (the streamlines incident to the singular points) and any corners of the closed boundaries, and mesh quality includes other measures such as angle error. Thus the output mesh may depart from the optimal streamline distribution or from the streamlines themselves, as in, for example [TPP*11] or [LCK21], or in Figure 1.
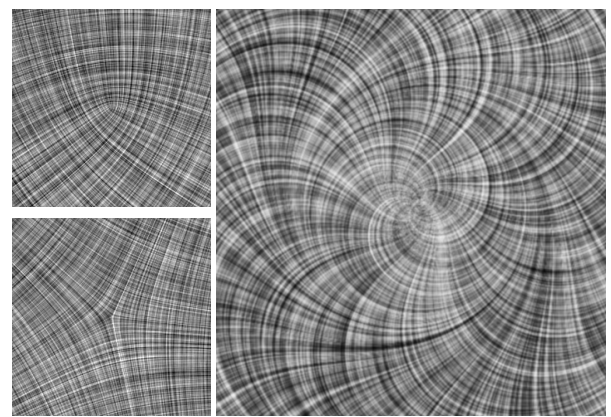
Figure 2: *We allow our cross fields to contain singular points of valence* 3 *(top left) and* 5 *(bottom left), but not limit cycles (right).*

Nonetheless our results are useful in various ways.

- If we construct a quad mesh which fails to meet a target sizing function or minimum anisotropy bound, we can estimate how much of this failure was inherent in the cross field and how much was due to downstream steps of the meshing process. This is useful since the cross fields we compute in practice are almost never exactly 1-integrable.
- Instead of designing cross fields only for smoothness under constraints, we could also optimize for isotropy or targeted anisotropy.
- The idea for removing limit cycles in Section 12 appears to be novel.

2. Related work

There is a great deal of research on using cross fields to construct quad meshes. The design of cross fields was surveyed in [VCD*16], and their use in quad meshing appears in the surveys [BLP*13, Cam17].

Since harmonic angle functions lead to 1-integrable cross fields, most cross field design papers focus on this goal. This difficult computational problem has inspired a great deal of research, including [RVAL09, KCPS13, ACBCO17, VO19, CCS*21]. Much of this

research focuses on quad meshing on a closed orientable surface embedded in $\mathbb{R}^3$, where a cross field might be computed as part of a *seamless parameterization*, eg. [BCE*13, CSZZ19]. In this situation there are stronger constraints on the cross field, which we avoid in our simple planar case; we return to this important topic in Section 14.

All singular point configurations allowed under Euler's Theorem can be quadrangulated, with the exception of the torus with two singular points of valence 3 and valence 5 [SZC*22, JT73]. However, this does not mean that they all admit 2-integrable cross fields or that all smooth cross fields are 2-integrable.

More closely related to our work is research on the explicit construction of 2-integrable cross fields. Constructing a *frame field*, in which the two unoriented directions at a regular point need not be orthogonal and may have different lengths, was explored in [PPTSH14] and [DVPSH15]. *Odeco tensors*, studied in [PSS21], were used recently in a section of [JCR23] to construct 2-integrable frame fields, by solving a PDE with objectives equivalent to our Equations 3 and 4.

Another relevant line of research has approached the problem of constructing a frame field by constructing a metric on the domain. In [JFH*15], a smooth orthogonal frame field is optimized so as match an arbitrary input metric as well as possible. In [CZK*19], they construct a metric that is guaranteed to admit an orthogonal frame at every point, which is equivalent to constructing the orthogonal frame field itself (see Section 3).

We can handle our simple case using fairly accessible mathematics, although we will point out some relevant sophisticated concepts as we go along.

3. Definitions

3.1. Input

We assume that we are given a smooth cross field $X(\alpha)$ represented by an angle-valued function $\alpha : R \to \mathbb{R} \mod \frac{\pi}{2}$ in a region R of the plane; see Figure 3. Specifically, for any real-valued $a : R \to \mathbb{R}$, let $\alpha = a(x,y) \mod \frac{\pi}{2}$, and define the *cross* at (x,y) to be the unordered set of four directions $[\alpha, \alpha + \frac{\pi}{2}, \alpha + \pi, \alpha + \frac{3\pi}{2}]$, or, equivalently, the unordered set of two unoriented directions $[\alpha, \alpha + \frac{\pi}{2}]$. (Often the angle is represented by a complex number, which is easier for computation, but this representation will yield some intuitive equations later on.)

Locally, we can orient the two directions and write the cross at any point as two orthogonal unit vectors, the columns of a rotation matrix:

$$\begin{bmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix}.$$

Cross fields are also called $4-$RoSy fields [PZ07] or $4-$directional symmetry fields [RVLL08].

A streamline in a cross field, like a streamline in a vector field or direction field, is a curve $\gamma(t)$ whose tangent at every point is the direction of the field. In a cross field, two unoriented streamlines pass through every regular point. Each streamline either forms a cycle or begins and ends at a boundary or singular point. In a cross
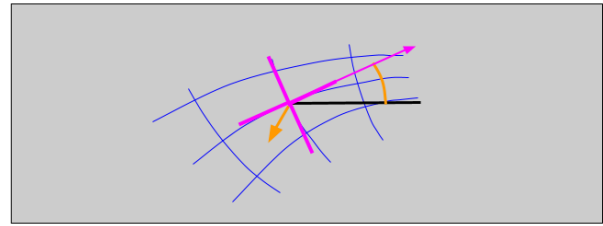
Figure 3: *The cross field is given by a periodic angle-valued function $\alpha(x,y)$, in radians, defined on the region R. The range of $\alpha(x,y)$ is $[0, \pi/2]$. The cross at a point (magenta), is given by four rotations of the vector $(\cos\alpha, \sin\alpha)$. The blue curves are streamlines. The orange arc indicates α, and the orange vector indicates $\nabla\alpha$, the direction in in which the cross rotates counter-clockwise most quickly. We observe that $\nabla\alpha$ is not aligned with the streamlines.*

field, a streamline may cross itself at right angles; that is, cycles are not necessarily *simple*.

The cross field $X(\alpha)$ may contain isolated *singular points*, at which the number of incident streamlines is not two. We only consider singular points of valence 3, incident to three streamlines, and valence 5, incident to five (these are also called singular points of index -1 and 1, respectively). Let S be the set of singular points.
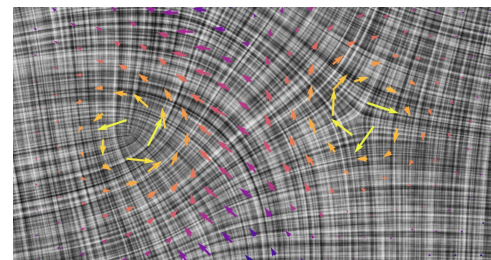

Figure 4: *Some vectors from $\nabla\alpha$ superimposed on the cross field $X(\alpha)$, colored based on distance to the nearest singular points. The magnitude of $\nabla\alpha$ increases near the singular points.*

In a disk-shaped neighborhood of any regular point, we can represent α using a scalar function a which may extend beyond the range $[0, \pi/2]$. In particular, the vector field $\nabla\alpha = \nabla a$ is well-defined, and it forms (like all gradients) a vector field with zero curl, although possibly with non-zero divergence (recall that the curl of vector field $(u(x,y), v(x,y))$ is $v_x - u_y$, and the divergence is $u_x + v_y$). The magnitude of $\nabla\alpha$ increases near a singular point s. The singular points of $\nabla\alpha$ are the singular points of $X(\alpha)$, where α is undefined.

A *limit cycle* is a cycle into which other streamlines "spiral in"; that is, a streamline γ is a cycle and there is another streamline ζ such that $\zeta(t)$ approaches γ as t goes to infinity; see Figure 2. X is smooth at a limit cycle γ. Since we forbid limit cycles, the neighborhood of a cycle streamline is covered by a set of nearly parallel cycles.

The cross field X is defined on a region R of the plane. We can

require X to be aligned with the boundaries of R, or not; but we do not allow periodic boundaries, at which X would have to match.
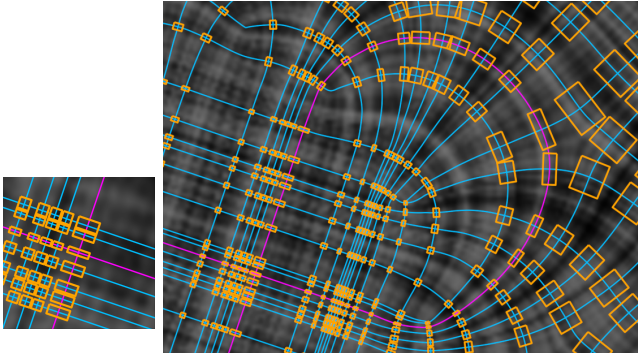
Figure 5: *Sometimes streamlines, like the one highlighted in magenta, can cross themselves, locking the aspect ratio $\sigma : \tau$ at the point of intersection. This shows that some cross fields require anisotropy.*

3.2. Orthogonal frame fields

Our goal will be to assign a length function along every undirected streamline, and hence two lengths at every regular point. Since, for example, a streamline can cross itself at right angles, as in Figure 5, we usually can't define the lengths as two separate continuous scalar functions on R. Instead, we will focus on a local neighborhood of every regular point, and require the length functions to be smooth in every neighborhood. Then we will make a global argument to show that the functions can be globally as well as locally smooth.

Let's define some local notation. In a neighborhood $N(p)$ of any regular point p, the two sets of streamlines, and their length functions, can be separated. Within the neighborhood, we call the two length functions $\sigma(x,y), \tau(x,y)$ (even if, globally, both come from the same set of streamlines). Also locally, we can orient both sets of streamlines. Now we can write a local orthogonal frame as two vectors, or the columns of a matrix:

$$\begin{bmatrix} \sigma\cos\alpha & -\tau\sin\alpha \\ \sigma\sin\alpha & \tau\cos\alpha \end{bmatrix}$$

In the next section, we will see that the change in σ or τ along a streamline is determined by the curvature of the orthogonal set of streamlines. We use these relationships to give the following definition.

Definition 1 A cross field is 2-integrable if we can define a length function for each direction vector at each point, creating an orthogonal anisotropic frame field, such that

1. (smoothness) in the neighborhood of any point $p \in R$, the two scalar length functions are both smooth (say C2-continuous), and
2. (α-consistency) along any streamline, the associated length function is consistent with α, as defined by Equations 3 and 4, below.

4. Local parameterization

What is stopping us from assigning, for example, $\sigma = \tau = 1$ everywhere? Unless we are extremely lucky, this will not correspond to the aspect ratios of any possible set of infinitesimal rectangles aligned with $X(\alpha)$. We capture this idea as follows. Let $P_{x,y}$ and $P_{u,v}$ be two copies of the Euclidean plane. Assume there is a smooth map M from some neighborhood in $P_{u,v}$ to $N(p)$, the neighborhood of $p \in R \subset P_{x,y}$. Under map M, infinitesimal axis-aligned squares in the $P_{u,v}$ are mapped to rectangles aligned with the cross field in $P_{x,y}$. See Figure 6.
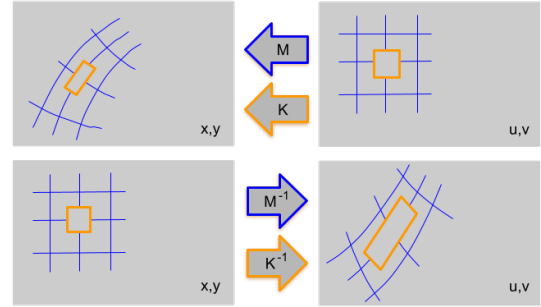
Figure 6: *On any neighborhood, we consider a local map M from plane $P_{u,v}$ into plane $P_{x,y}$. The Jacobian K of M takes infinitesimal axis-aligned squares in (u,v) to infinitesimal rectangles aligned with the cross field in (x,y). The inverse Jacobian takes infinitesimal squares in (x,y) to parallelograms in (u,v).*

We don't need to write down M explicitly in order to understand the local behavior of infinitesimal squares. We are just interested in its derivative, the Jacobian matrix

$$K = \begin{bmatrix} x_u & x_v \\ y_u & y_v \end{bmatrix} = \begin{bmatrix} \sigma\cos\alpha & -\tau\sin\alpha \\ \sigma\sin\alpha & \tau\cos\alpha \end{bmatrix}$$

Here σ, τ and α are functions of (x,y), but $x(u,v), y(u,v)$ are functions of u,v.

Locally, the inverse map M^{-1} always exists, and its Jacobian K^{-1} is the inverse of the Jacobian of M:

$$K^{-1} = \begin{bmatrix} u_x & u_y \\ v_x & v_y \end{bmatrix} = \begin{bmatrix} (1/\sigma)\cos(-\alpha) & -(1/\sigma)\sin(-\alpha) \\ (1/\tau)\sin(-\alpha) & (1/\tau)\cos(-\alpha) \end{bmatrix} \quad (1)$$

It's easier to figure out how σ, τ have to change with α by looking at K^{-1} because its elements u_x, u_y, v_x, v_y are all functions of (x,y).

We note that K^{-1} also represents a Riemannian metric at p. If $\gamma(t)$ is a unit-speed curve in (x,y), then the length of γ using the metric K^{-1} is the length of $M(\gamma(t))$ in (u,v)

$$\int_{t=0}^{\text{length } \gamma} \left[(K^{-1}\gamma_t)^T (K^{-1}\gamma_t) \right]^{1/2} dt \quad (2)$$

We also observe that the rows, rather than the columns, of K^{-1} are orthogonal (unless $\sigma = \tau$), so K^{-1} is not in general an orthogonal frame, and it maps squares to parallelograms.

The functions u, v have to satisfy

$$(u_x)_y = (u_y)_x$$
$$(v_x)_y = (v_y)_x.$$

That is, ∇u and ∇v, like all gradients, have zero curl. From $(u_x)_y = (u_y)_x$ we have

$$((1/\sigma)\cos\alpha)_y = ((1/\sigma)\sin\alpha)_x$$
$$-\sigma_y\cos\alpha - \sigma\sin\alpha\,\alpha_y = -\sigma_x\sin\alpha + \sigma\cos\alpha\,\alpha_x$$
$$\nabla\sigma(-\sin\alpha,\cos\alpha) = -\sigma\,\nabla\alpha\,(\cos\alpha,\sin\alpha)$$

and using $(v_x)_y = (v_y)_x$ we have

$$((-1/\tau)\sin\alpha)_y = ((1/\tau)\cos\alpha)_x$$
$$\tau_y\sin\alpha - \tau\cos\alpha\,\alpha_y = -\tau_x\cos\alpha - \tau\sin\alpha\,\alpha_x$$
$$\nabla\tau(\cos\alpha,\sin\alpha) = \tau\,\nabla\alpha\,(-\sin\alpha,\cos\alpha)$$

We highlight these two partial differential equations.

$$\nabla\sigma(-\sin\alpha,\cos\alpha) = -\sigma\,\nabla\alpha\,(\cos\alpha,\sin\alpha) \tag{3}$$
$$\nabla\tau(\cos\alpha,\sin\alpha) = \tau\,\nabla\alpha\,(-\sin\alpha,\cos\alpha) \tag{4}$$

(One could get equivalent equations by requiring a Lie bracket to be the zero vector, as in [JCR23]. Alternatively, one could relate the metric to the curvature of $\nabla\alpha$ via the Levi-Civita connection.)

5. Isotropic cross fields

If $\tau = \sigma$ everywhere, then the maps M, M^{-1} are *conformal*, and the isotropic metrics K, K^{-1} are called *isothermal*. Most cross field design programs optimize towards conformal maps (by minimizing the Dirichlet energy or the Ginzberg-Landau energy, or following the Ricci flow, etc.). With $\tau = \sigma$, our two differential equations become

$$\nabla\sigma(-\sin\alpha,\cos\alpha) = -\sigma\,\nabla\alpha\,(\cos\alpha,\sin\alpha)$$
$$\nabla\sigma(\cos\alpha,\sin\alpha) = \sigma\,\nabla\alpha\,(-\sin\alpha,\cos\alpha)$$

Thinking of the vectors $(\cos\alpha,\sin\alpha)$ and $(-\sin\alpha,\cos\alpha)$ as an orthonormal coordinate system, we see that

$$\sigma_x = -\sigma\alpha_y \tag{5}$$
$$\sigma_y = \sigma\alpha_x \tag{6}$$

So in a region where $\sigma = \tau$, there is a unique solution for σ given α, up to a global scale factor. We can think of Equations 5 and 6 as a way of writing the usual Cauchy-Riemann equations $u_x = v_y$ and $u_y = -v_x$, which describe an isothermal metric.

In order for these differential equations to be solvable, however, (σ_x, σ_y) has to have zero curl. That is, at every point in R,

$$(\sigma\alpha_x)_x = (-\sigma\alpha_y)_y$$
$$\sigma_x\alpha_x + \sigma\alpha_{xx} = -\sigma_y\alpha_y - \sigma\alpha_{yy}$$
$$-\sigma\alpha_x\alpha_y + \sigma\alpha_{xx} = -\sigma\alpha_x\alpha_y - \sigma\alpha_{yy}$$

which shows that α must be harmonic (that is, $\alpha_{xx} + \alpha_{yy} = 0$), and $\nabla\alpha$ has zero curl and zero divergence.

6. Singular points

In a neighborhood containing a singular point s of odd valence, we cannot separate the streamlines into two distinct sets; so the functions σ, τ can be smooth on the entire neighborhood only if they are identical at s.

Observation 2 At a singular point s, $\sigma(s) = \tau(s)$.

So to understand the relationship of σ and α at singular points we can examine the well-understood isotropic case. A singular point s of valence 3 is similar (scaled, rotated and/or translated) to the singular point at the origin induced by the complex map $M^{-1}(z) = z^{3/4}$, which we can write in terms of (x,y) and (u,v) as

$$M^{-1}(x,y) = (x^2+y^2)^{3/8}e^{(3/4)i\arctan(y/x)} = (u,v)$$

We can similarly write the complex derivative $K^{-1}(z) = 3/4z^{-1/4}$ as

$$K^{-1}(x,y) = 3/4(x^2+y^2)^{-1/8}e^{(-1/4)i\arctan(y/x)} = (u,v)$$

From Equation 1, we see that $1/\sigma(x,y) = 3/4(x^2+y^2)^{-1/8}$ and $-\alpha = (-1/4)\arctan(y/x)$. Recalling that $\nabla\arctan(y/x) = (-y,x)/(x^2+y^2)$, we find

$$\sigma = 4/3\,(x^2+y^2)^{1/8}$$
$$\nabla\sigma = 1/3\,(x^2+y^2)^{-7/8}(x,y)$$
$$\alpha = 1/4\,\arctan y/x$$
$$\nabla\alpha = 1/4\,(x^2+y^2)^{-1}(-y,x)$$

We can verify that these equations obey Equations 5 and 6. They imply a counter-clockwise rotation of $\nabla\alpha$ around s (as in Figure 4), and a zero of σ at s.

Similarly, a singular point s of valence five at the origin is produced by $M^{-1}(z) \to z^{5/4}$, with $K^{-1}(z) = 5/4\,z^{1/4}$. In this case we find that $\nabla\alpha$ rotates clockwise, and s is a pole.

In either case, we avoid tracing streamlines very close to singularities due to the difficulty in accurately estimating σ and τ as $\nabla\alpha$ grows without bound.

7. Integrating σ and τ

When α is not harmonic, we must allow τ and σ to differ. Equations 3 and 4 are not sufficient to determine $\nabla\sigma$ and $\nabla\tau$. But they do completely determine two directional derivatives per point, one for σ and one for τ. (This section is an example of the "method of characteristics" for solving first-order PDEs.)

In particular, for a unit-speed curve $\gamma(t)$ through a point $p \in R$, with unit direction vector $(-\sin\alpha,\cos\alpha)$, the directional derivative

$$\frac{d\sigma}{dt} = \nabla\sigma\,(-\sin\alpha,\cos\alpha) = -\sigma\,\nabla\alpha\,(\cos\alpha,\sin\alpha)$$

So, for a streamline following $(-\sin\alpha,\cos\alpha)$, given a single value of σ somewhere along the streamline, we can find all the other values of σ along the streamline by integrating using Equation 3.

Similarly we can integrate τ along the streamlines with direction $(\cos\alpha,\sin\alpha)$, using Equation 4:

$$\frac{d\tau}{dt} = \nabla\tau\,(\cos\alpha,\sin\alpha) = \tau\,\nabla\alpha\,(-\sin\alpha,\cos\alpha)$$

Fortunately, these are exactly the streamlines we have.

For example, consider an intrinsic rectangle not containing a singular point. We can set arbitrary values of σ along one edge, and arbitrary values of τ along another, and then integrate both into the
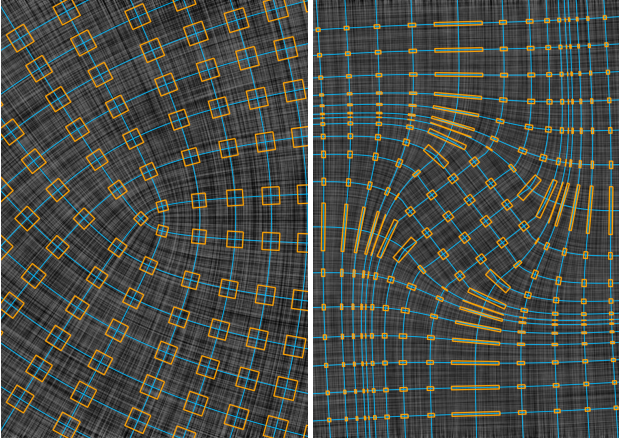
Figure 7: *Given the isotropic cross field derived from by the conformal map $z \to z^{3/4}$ (left) we (of course) find that the minimum anisotropy is zero; that is, we can set $\sigma = \tau$ everywhere. A cross field with nonzero divergence (right) exhibits large anisotropy, even in some regions where $\nabla\alpha = 0$. We computed this twisting cross field by requesting a large isotropic scale in the center and a smaller isotropic scale on the outside, and minimizing the Dirichlet energy. Like most sets of user-defined cross field constraints, there is no harmonic α function that meets them. The non-harmonic $X(\alpha)$ we found does achieve the large isotropic scale in the center and smaller isotropic scale in the corners, but requires large anisotropy. Requesting a larger scale difference in this case produces singular points, which allows the scale change with lower anisotropy.*

interior using equations 3 and 4. If the initial values of σ are smooth along the edge, the two-dimensional σ function across the interior of the intrinsic rectangle will be smooth, and similarly for τ.

Recall that the variables σ and τ are only defined on a small local neighborhood not containing a singularity. As seen in Figure 5, a streamline in a cross field can intersect itself at right angles, and at the intersection point σ and τ are defined by different points along the same field line. At points of self-intersection like these, the ratio between σ and τ is fixed, since any change to one value on the streamline changes the other.

8. Global consistency of σ and τ

The previous section established that we can locally integrate σ and τ along streamlines within a disk-shaped singularity-free neighborhood. So the length function along an entire streamline γ is determined if we set the value at any one point $p \in \gamma$. A connected set of adjacent streamlines whose length values are continuous along any curve orthogonal to them all will have continuous values throughout their length.

Self-intersecting streamlines are not a problem. But what about cycle streamlines? Say γ is a cycle following direction $(\cos\alpha, \sin\alpha)$. If we fix the value of $\tau(p)$ at one point $p \in \gamma$, and we integrate along γ using Equation 4 and return to p, do we always get the same value of τ? Notice that this is not true of every
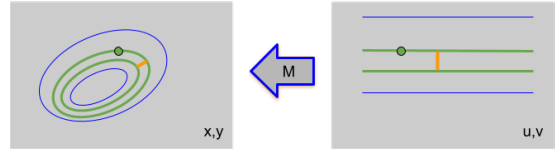
Figure 8: *Illustration for Theorem 3. Mapping a point p on cycle $\gamma(0)$ to parameter space $P_{u,v}$ takes the cycle to a line. A parallel line at distance ϵ maps back (pulls back) to a nearby cycle. The distance between the two cycles in $P_{x,y}$ always returns to the same value as we loop around, and converges to $\epsilon\tau(p)$ as $\epsilon \to 0$.*

function; for instance, integrating $\nabla\alpha$ around a cycle would give us 2π rather than zero, even though $\nabla\alpha$ has zero curl.

Theorem 3 Let γ be a cycle streamline, not a limit cycle, with no singular point in its neighborhood. Let τ be the length function orthogonal to γ. The function τ is smooth everywhere on γ, including at p.

Proof Since γ is not a limit cycle, a small neighborhood of γ is covered by a set of nearby disjoint parallel cycle streamlines. Let $\gamma(0) = \gamma$. Recall the mappings $M : (u,v) \to (x,y)$ and $M^{-1} : (x,y) \to (u,v)$ from Section 4. At an arbitrary point $p \in \gamma$, let $q = M^{-1}(p)$. See Figure 8. From q, a an ϵ step in the v direction takes us to point $q + \epsilon v$, which is mapped by K to step orthogonal to $\gamma(0)$, landing us on a nearby cycle $\gamma(\epsilon)$. As ϵ goes to zero, the distance from $p \in \gamma(0)$ to the nearest point on $\gamma(\epsilon)$ converges to the infinitesimal width $w(p, \epsilon) =$

$$\left[\begin{bmatrix} 0 & \epsilon \end{bmatrix} \begin{bmatrix} \sigma\cos\alpha & \sigma\sin\alpha \\ -\tau\sin\alpha & \tau\cos\alpha \end{bmatrix} \begin{bmatrix} \sigma\cos\alpha & -\tau\sin\alpha \\ \sigma\sin\alpha & \tau\cos\alpha \end{bmatrix} \begin{bmatrix} 0 \\ \epsilon \end{bmatrix} \right]^{1/2} = \tau\epsilon$$

with $w(p, \epsilon)/\epsilon$ converging to τ. Since cycles $\gamma(0)$ and $\gamma(\epsilon)$ are smooth and closed, if we traverse all the way around $\gamma(0)$ and return to p, we get the same infinitesimal width $w(p, \epsilon)$ and converge to the same value of τ. $\square$

Setting the value of $\sigma = \tau$ at a singular point s determines the length function along the (3 or 5) separatrices adjacent to s. We required each of these separatrices to end at the boundary, not at another singular point. Hence,

Observation 4 We can choose the values of $\sigma = \tau$ at each singular point independently of each other.

Theorem 5 There is an infinite one-dimensional set of ways to choose the length functions along each streamline, consistent with σ, τ being smooth on the neighborhood of any point in R.

Proof First, we select the size values $\sigma = \tau$ at the singular points. There is a discrete finite set of such values, and we can choose them independently because of Observation 4. This choice determines the length values along the separatrices. The separatrices divide the remaining streamlines into groups, such that we cannot deform a streamline from one group into one from another group without crossing one or more singular points, or leaving the region (that is, the groups are *homotopic*). Some groups may be unions of cycles, while others begin and

end on the boundary. Within each group, we can pick any one-dimensional curve segment *g* crossing each streamline in the group once (for example, *g* might be a segment of one of the streamlines in the orthogonal direction, or a boundary curve). We pick a smooth length function along *g*, respecting any length values already fixed at the endpoints because they lie on separatrices. This fixes the length values all along the entire group of streamlines. The length values are continuous at the separatrices since they were continuous at *g*. All of the length values along all of the segments *g* form an infinite one-dimensional set, parameterizing all the possible length assignments.   □

9. Divergence of $\tau\nabla\alpha$

In this section we comment on the behavior of $\nabla\alpha$ on a cycle streamline. Let γ be a cycle streamline γ, not a limit cycle, and let's use τ as the length function that is integrated along γ.

The *flux* of the vector field $\nabla\alpha$ across a cycle γ is

$$\int_\gamma \nabla\alpha \, (-\sin\alpha, \cos\alpha).$$

The usual physical intuition is that if α is the flow of some quantity as a function of time, the flux is the net movement of the quantity across the curve γ per unit time.

Observation 6 If $x > 0$, we have $(\ln x)_x = 1/x$. So

$$\nabla\ln\tau = \frac{\nabla\tau}{\tau}.$$

Lemma 7 The flux of the vector field $\nabla\alpha$ is zero across any cycle streamline γ.

Proof The proof of Theorem 3 showed that the function τ is smooth everywhere on cycle γ, and thus, so long as $\tau > 0$ everywhere, the function $\ln\tau$ is also smooth everywhere on γ. So

$$0 =$$
$$\int_\gamma \nabla(\ln\,\tau)\,(\cos\alpha, \sin\alpha) =$$
$$\int_\gamma \frac{\nabla\tau}{\tau}(\cos\alpha, \sin\alpha) =$$
$$\int_\gamma \nabla\alpha\,(-\sin\alpha, \cos\alpha)$$

□

If γ is a simple cycle, the flux is also the right-hand side of Gauss's divergence theorem:

$$\int\int_{\text{interior } \gamma} \nabla\cdot\nabla\alpha = \int_\gamma \nabla\alpha\,(-\sin\alpha, \cos\alpha)$$

That gives us:

Observation 8 Let γ be a simple cycle streamline. The total divergence of the vector field α in the region bounded by γ must be zero.

This tells us, on the one hand, that simple cycles only appear in quite special cross fields, since in general the divergence of $\nabla\alpha$ at an arbitrary point is not zero. The situation in which $\nabla\cdot\nabla\alpha$ is everywhere zero is the useful case in which α is everywhere locally harmonic.

On the other hand, since divergence can be positive or negative, the total divergence of $\nabla\alpha$ in the interior of a cycle streamline γ may be zero without $\nabla\cdot\nabla\alpha = 0$ at any point in the interior.

10. Definition of anisotropy

For non-harmonic α, we need a functional that we can optimize to choose among the valid values of σ, τ. In this section we consider minimizing the anisotropy. Making σ and τ as similar as possible also encourages them to be smooth; if, say, the smoothness of σ along a streamline is given by the smoothness of $\nabla\alpha$, encouraging τ to change similarly to σ also encourages τ to be smooth.

Using Equations 3 and 4, we can solve for the values of σ, τ at the vertices of the mesh formed by the representative streamlines, minimizing the maximum anisotropy over the region. Let A be our measure of anisotropy.

$$\min_{\sigma,\tau} \int_R A(x,y)$$
st.
$$-(\sigma_y)\cos\alpha - \sigma\alpha_y\sin\alpha = -(\sigma_x)\sin\alpha + \sigma\alpha_x\cos\alpha$$
$$(\tau_y)\sin\alpha - \tau\alpha_y\cos\alpha = -(\tau_x)\cos\alpha - \tau\alpha_x\sin\alpha$$
$$\int_R \sigma^2 + \tau^2 = 1$$

Making $\sigma^2 + \tau^2$ integrate to one over the whole region keeps them from going to zero.

There are many ways we could define the anisotropy function A. We use:

$$A = \left[\frac{\sigma}{\tau} + \frac{\tau}{\sigma} - 2\right]^{2k} \tag{7}$$

The exponent $k \geq 1$ is an integer. Raising the anisotropy to an even power ensures it's positive, although the constant offset -2 sufficiently penalizes negative aspect ratios that we do not see negative σ or τ in practice. The greater k, the more it penalizes the maximum. Subtracting two from the sum of the two ratios gives an energy function that equals zero for isotropic size functions.

11. Discretization, optimization and visualization

We select a sparse set L of streamlines to represent the cross field. For each $\ell \in L$, we arbitrarily set $\sigma(p)$ for some point in ℓ, which determines the σ values along the streamline, based on Equation 3. We then define a scalar multiplier for each streamline ℓ, which changes all of its σ values simultaneously. Each intersection between two streamlines in L is a point where both σ and τ are defined. We evaluate the function A from Equation 7 at each of these points.

Minimizing the anisotropy measure is achieved through an iterative process. We resize σ along each chord $\ell \in L$ in turn, so as to minimize the sum of the anisotropy measures of all intersections involving ℓ. If a pass through all chords yields no improvement, or an improvement below a given threshold, the optimization is complete. Since the total sum of the anisotropy measures at all intersections is reduced at each step, this process terminates.

To select the lines in L, we start with a subset T of streamlines

that captures the topology of X. These may be the separatrices (as in Figure 1) or they may be regular streamlines (as in Figure 10). Along the streamlines in T, we select orthogonal streamlines to include in L, until every point on a line in T is sufficiently close to a streamline of L. We do not use separatrices in L, as growth of $\nabla \alpha$ near singularities makes robust size estimation difficult.

As a check, we find that this solve behaves well on conformal cross fields, even though we do not use Equations 5 and 6. The conformal cross field in Figure 7 (left) exhibits a worst-element numerical anisotropy of $1 : 1.0033$.

To visualize cross fields, we use line integral convolution, a technique that "smears" a noise function into strokes along the streamlines, as suggested by [CL93].

We visualize anisotropy in two ways. In low-anisotropy cases, we draw an orange rectangle centered at every intersection point p of two lines in L, reflecting the aspect ratio and sizes given by $\sigma(p), \tau(p)$. One might expect that there is some scale factor that would grow all of the squares to nicely "collide" with their neighbors, creating a rough approximation of a mesh; this is not true because L is chosen before we find σ and τ. When anisotropy is very large, we visualize it instead using a color scale, as in Figure 10. There σ and τ are interpolated over the region, and a per-pixel anisotropy is calculated (except for near singularities along the wing surface where the singularity is not fully enclosed by representative streamlines).

12. Removable limit cycles

Simple limit cycles appear in vector fields as well as cross fields and are studied in dynamical systems. There are two kinds of limit cycles, stable and unstable. Unstable limit cycles do not cause much trouble; they disappear under small smooth perturbations of α, while stable limit cycles persist. A stable limit cycle is one in which, as we traverse γ, the nearby streamlines on either side either all approach or all move away from γ (depending on our direction of travel on γ). See Figure 9.

Some limit cycles in cross fields are forced by the topology, for example, the torus with one singular point of valence 3 and one of valence 5 must have a limit cycle which is also a fundamental cycle of the torus; this is why it cannot be meshed. But limit cycles may also appear where they are not forced.

As we see in Figure 9, a stable simple limit cycle γ in a planar cross field can be removed by replacing a small neighborhood around γ, without adding new singular points. Before removal, on one side of γ, one set of streamlines cycles in towards γ, and the other set crosses γ, and similarly on the other side. To remove γ, we smoothly add a twist of $\pi/2$ to α in the neighborhood of γ, so that each streamline spiraling into the neighborhood is connected to a streamline crossing γ, and visa versa. Since we need two sets of streamlines, this trick does not work on vector fields.

13. Results

The cross field in Figure 1 is the upper-left portion a cross field computed using a method similar to that of [KCPS13]; it was given
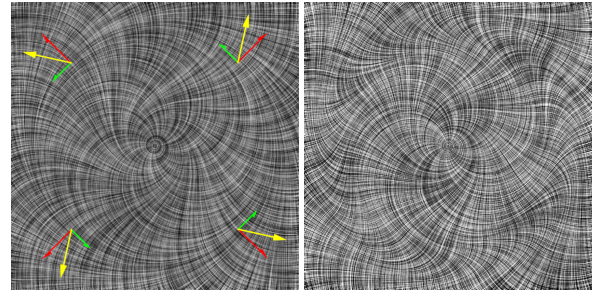
Figure 9: *A simple limit cycle (left) can be smoothly replaced with a region within which a twist of $\pi/2$ swaps the two sets of streamlines (right). The streamlines away from the cycle remain unchanged, and the region in which the swap occurs can be made arbitrarily thin. An exaggerated $\nabla\alpha$ vector (yellow) has been drawn on top of the left image, along with the components tangential (green) and normal (red) to the limit cycle; limit cycles are distinguished by the nonzero normal component (and hence have a non-zero net flux of $\nabla\alpha$.)*

initial conditions requiring small isotropic elements over a region in the center and large isotropic elements on the outside.

In figure 10, we show a cross field generated from a two-dimensional simulation of airflow over a wing, NACA Airfoil 8650 from [SQB21]. The original flow vector field had more than one hundred singularities. We created a cross field by adding the orthogonal direction at each point, and then smoothed. This reduced the number of singularities and had the additional benefit of removing two ellipse-shaped limit cycles corresponding to vortices in the fluid flow. Each limit cycle was replaced by a pair of valence-three singularities, surrounded by a lens-shaped region with less total curvature. The central part of the image has very low anisotropy, as the vertical path lines can be freely adjusted to match the horizontal path lines. The left side of the image has slightly more anisotropy, as those collisions represent streamlines that form cycles around the entire wing. The turbulent airflow to the right of the wing creates extreme anisotropies, as areas of high rotation in the field require tiny chord widths. This could be fixed by adding more singularities as we move away from the turbulent region.

14. Discussion

The main limitation of this work is that it only applies to regions in the plane with boundary. The interesting case of a cross field where all streamlines have finite length on a two-dimensional surface embedded in $\mathbb{R}^3$ is much more structured. The separatrices all begin and end at singular points, and all other streamlines must be cycles. Observation 4 does not hold. Might there be a cycle of separatrices for which there is no set of lengths obeying Equations 3 and 4 and Observation 2? In [CZK*19], they use a version of the "tiny bricks" argument to establish that we can consistently assign two length functions at every point away from the separatrices, but it is not clear that these length functions can or must be continuous at the separatrices or singular points.
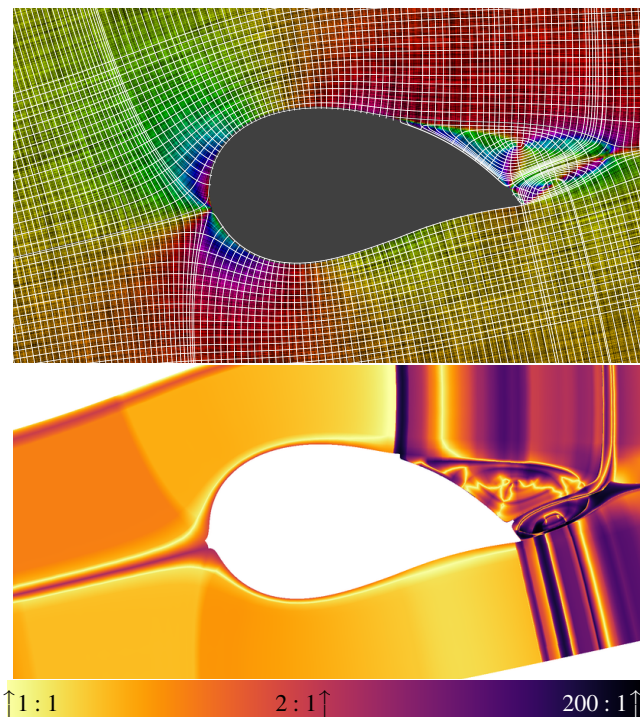
Figure 10: *Airflow over an airfoil, with prominent trailing vortices. We start with a cross field that is aligned with the airflow around the wing (top); the background LIC is colored to match field direction to highlight singularities. Next, representative streamlines (top image, white) are traced. We optimize for a minimal-anisotropy configuration on the representatives, and we visualize the anisotropy by interpolating colors (bottom). The most extreme aspect ratio is 192:1.*

References

[ACBCO17] AZENCOT O., CORMAN E., BEN-CHEN M., OVSJANIKOV M.: Consistent functional cross field design for mesh quadrangulation. *ACM Transactions on Graphics (TOG) 36*, 4 (2017), 1–13. 2

[BCE*13] BOMMES D., CAMPEN M., EBKE H.-C., ALLIEZ P., KOBBELT L.: Integer-grid maps for reliable quad meshing. *ACM Transactions on Graphics (TOG) 32*, 4 (2013), 1–12. 3

[BLP*13] BOMMES D., LÉVY B., PIETRONI N., PUPPO E., SILVA C., TARINI M., ZORIN D.: Quad-mesh generation and processing: A survey. *Computer graphics forum 32*, 6 (2013), 51–76. 2

[Cam17] CAMPEN M.: Partitioning surfaces into quadrilateral patches: A survey. *Computer Graphics Forum 36*, 8 (2017), 567–588. 2

[CBK15] CAMPEN M., BOMMES D., KOBBELT L.: Quantized global parametrization. *ACM Trans. Graph. 34*, 6 (nov 2015). 1

[CCS*21] CAMPEN M., CAPOUELLEZ R., SHEN H., ZHU L., PANOZZO D., ZORIN D.: Efficient and robust discrete conformal equivalence with boundary. *ACM Transactions on Graphics (TOG) 40*, 6 (2021), 1–16. 2

[CL93] CABRAL B., LEEDOM L. C.: Imaging vector fields using line integral convolution. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1993), SIGGRAPH '93, Association for Computing Machinery, p. 263–270. 8

[CSZZ19] CAMPEN M., SHEN H., ZHOU J., ZORIN D.: Seamless parametrization with arbitrary cones for arbitrary genus. *ACM Transactions on Graphics (TOG) 39*, 1 (2019), 1–19. 3

[CZK*19] CHEN W., ZHENG X., KE J., LEI N., LUO Z., GU X.: Quadrilateral mesh generation i: Metric based method. *Computer Methods in Applied Mechanics and Engineering 356* (2019), 652–668. 3, 8

[DVPSH15] DIAMANTI O., VAXMAN A., PANOZZO D., SORKINE-HORNUNG O.: Integrable polyvector fields. *ACM Transactions on Graphics (TOG) 34*, 4 (2015), 1–12. 3

[JCR23] JEZDIMIROVIĆ J., CHEMIN A., REMACLE J.-F.: Integrable cross-field generation based on imposed singularity configuration—the 2d manifold case. In *International Meshing Roundtable* (2023), Springer, pp. 343–369. 3, 5

[JFH*15] JIANG T., FANG X., HUANG J., BAO H., TONG Y., DESBRUN M.: Frame field generation through metric customization. *ACM Transactions on Graphics (TOG) 34*, 4 (2015), 1–11. 3

[JT73] JUCOVIČ E., TRENKLER M.: A theorem on the structure of cell–decompositions of orientable 2–manifolds. *Mathematika 20*, 1 (1973), 63–82. 3

[KCPS13] KNÖPPEL F., CRANE K., PINKALL U., SCHRÖDER P.: Globally optimal direction fields. *ACM Transactions on Graphics (ToG) 32*, 4 (2013), 1–10. 2, 8

[LCK21] LYON M., CAMPEN M., KOBBELT L.: Simpler quad layouts using relaxed singularities. In *Computer Graphics Forum* (2021), vol. 40, Wiley Online Library, pp. 169–180. 2

[PPTSH14] PANOZZO D., PUPPO E., TARINI M., SORKINE-HORNUNG O.: Frame fields: Anisotropic and non-orthogonal cross fields additional material. *Proceedings of the ACM TRANSACTIONS ON GRAPHICS (PROCEEDINGS OF ACM SIGGRAPH)* (2014). 3

[PSS21] PALMER D., STEIN O., SOLOMON J.: Frame field operators. In *Computer Graphics Forum* (2021), vol. 40, Wiley Online Library, pp. 231–245. 3

[PZ07] PALACIOS J., ZHANG E.: Rotational symmetry field design on surfaces. *ACM Transactions on Graphics (TOG) 26*, 3 (2007), 55–es. 3

[RVAL09] RAY N., VALLET B., ALONSO L., LEVY B.: Geometry-aware direction field processing. *ACM Transactions on Graphics (TOG) 29*, 1 (2009), 1–11. 2

[RVLL08] RAY N., VALLET B., LI W. C., LÉVY B.: N-symmetry direction field design. *ACM Transactions on Graphics (TOG) 27*, 2 (2008), 1–13. 3

[SQB21] SCHILLACI A., QUADRIO M., BORACCHI G.: A database of CFD-computed flow fields around airfoils for machine-learning applications, Mar. 2021. URL: https://doi.org/10.5281/zenodo.4106752, doi:10.5281/zenodo.4106752. 8

[SZC*22] SHEN H., ZHU L., CAPOUELLEZ R., PANOZZO D., CAMPEN M., ZORIN D.: Which cross fields can be quadrangulated? global parameterization from prescribed holonomy signatures. *ACM Transactions on Graphics (TOG) 41*, 4 (2022), 1–12. 3

[TPP*11] TARINI M., PUPPO E., PANOZZO D., PIETRONI N., CIGNONI P.: Simple quad domains for field aligned mesh parametrization. In *Proceedings of the 2011 SIGGRAPH asia conference* (2011), pp. 1–12. 2

[VCD*16] VAXMAN A., CAMPEN M., DIAMANTI O., PANOZZO D., BOMMES D., HILDEBRANDT K., BEN-CHEN M.: Directional field synthesis, design, and processing. *Computer graphics forum 35*, 2 (2016), 545–572. 2

[VO19] VIERTEL R., OSTING B.: An approach to quad meshing based on harmonic cross-valued maps and the Ginzburg–Landau theory. *SIAM Journal on Scientific Computing 41*, 1 (2019), A452–A479. 2