

Controlled-Distortion Constrained Global Parametrization

Ashish Myles* and Denis Zorin†
New York University

Abstract

The quality of a global parametrization is determined by a number of factors, including amount of distortion, number of singularities (cones), and alignment with features and boundaries. Placement of cones plays a decisive role in determining the overall distortion of the parametrization; at the same time, feature and boundary alignment also affect the cone placement. A number of methods were proposed for automatic choice of cone positions, either based on singularities of cross-fields and emphasizing alignment, or based on distortion optimization.

In this paper we describe a method for placing cones for seamless global parametrizations with alignment constraints. We use a close relation between variation-minimizing cross-fields and related 1-forms and conformal maps, and demonstrate how it leads to a constrained optimization problem formulation. We show for boundary-aligned parametrizations metric distortion may be reduced by *cone chains*, sometimes to an arbitrarily small value, and the trade-off between the distortion and the number of cones can be controlled by a regularization term. Constrained parametrizations computed using our method have significantly lower distortion compared to the state-of-the-art field-based method, yet maintain feature and boundary alignment. In the most extreme cases, parametrization collapse due to alignment constraints is eliminated.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—[Geometric algorithms, languages, and systems];

Keywords: geometric modeling, parametrization

Links:  DL  PDF

1 Introduction

Global surface parameterization has a variety of uses, including surface quadrangulation, tiling the surface seamlessly with texture maps and solving equations on surfaces. Many applications require constrained parametrizations, with parametric lines aligned with surface boundaries, creases, features, or user-specified directions.

A global parametrization can be defined as a flat metric on the surface, with isolated cones where all Gaussian curvature is concentrated. If the surface is cut to a disk with the cut going through all cones, the metric determines a mapping to the plane up to a rigid transform.

In the absence of constraints, cones may be necessary either for topological reasons (for closed surfaces of positive genus), or to

reduce a measure of parametrization distortion, most commonly deviation from isometry. Additional constraints on parametric lines may require additional cones: for example, an unconstrained parametrization of a planar mesh is perfectly isometric, while boundary alignment may require a parametrization with cones.




Figure 1: *Quadrangulation results for cone locations found using the field optimization of [Bommes et al. 2009] (top), and with our method (bottom) with visualized isometric distortions computed as indicated in Section 7.*

Among recent methods, [Bommes et al. 2009] uses a guidance field and constraints on parametric variables to align the parametrization with features. The problem of determining cone positions based on distortion was considered by a number of authors. Most recently, [Myles and Zorin 2012] proposed *incremental flattening*, a way to evolve the metric of the surface, concentrating curvature at isolated points which become cones. The method produces low-distortion parametrization for a variety of surfaces, but does not allow for feature alignment.

In this paper, we present an algorithm for concentrating curvature of a surface at cones in the presence of feature alignment constraints, driven by distortion minimization. Our framework unifies methods based on field smoothing [Bommes et al. 2009], and conformal map ideas [Myles and Zorin 2012] in a common framework suggested by the notion of a connection 1-form [Crane et al. 2010]. This view makes it possible, on the one hand, to control alignment with features precisely, and, on the other hand, include metric distortion

*e-mail: amyles@cs.nyu.edu

†e-mail: dzorin@cs.nyu.edu

directly in the optimization energy.

We demonstrate that in general if an arbitrary number of cones can be introduced, the metric distortion can be made *arbitrarily low*, if meshes can be refined. We show that a smoothness term in the energy provides a natural control of the trade-off between the number of cones and metric distortion.

2 Related work

A broad overview of quadrangulation and field-guided techniques is given in [Bommes et al. 2012]. Here we focus on the works most relevant to our goals.

The main foundation of our work includes the work on [Myles and Zorin 2012], [Ray et al. 2008], [Bunin 2008b] and [Crane et al. 2010]. We show how considering cross-fields and conformal maps in a unified framework based on 1-forms [Crane et al. 2010] (referred to as field curvature in [Ray et al. 2008]), leads to a unified formulation for the problem of placing cones for constrained parametrization. Our methods include both the greedy approach of [Bommes et al. 2009] for fields and [Myles and Zorin 2012] for conformal maps as special cases. We discuss the relation of our work to these papers in detail in the next sections.

While the details of our work is most closely related to [Bommes et al. 2009] in terms of cone placement, similar field-based approaches were developed in [Ray et al. 2006; Kälberer et al. 2007]; curl reduction in [Ray et al. 2006] is particularly relevant.

Various ways of reducing distortion of conformal parametrization are considered in [Jin et al. 2004; Ben-Chen et al. 2008; Springborn et al. 2008]; these works do not consider alignment constraints.

The method of [Ray et al. 2009], based on [Palacios and Zhang 2007], generates smooth fields with fewer cones on geometrically intricate or noisy data by smoothing surface curvature. In contrast, our method trades off between distortion minimization and field smoothness, but could also, in principle, use curvature smoothing to reduce cones further.

Techniques based on the construction of a base domain using simplification offer good control over the number of cones, and typically result in a relatively even distribution [Eck et al. 1995; Lee et al. 1998; Khodakovsky et al. 2003; Marinov and Kobbelt 2005] as well as [Daniels II et al. 2009; Daniels et al. 2009; Pietroni et al. 2009; Tarini et al. 2010]. These methods can only handle parametrization alignment by including the edges to align to in the boundaries of the triangles/quads of the base domain. The distortion is often suboptimal for the number of cones, and difficult to control. A number of methods [Gu and Yau 2003; Dong et al. 2006; Tong et al. 2006; Ben-Chen et al. 2008; Springborn et al. 2008] use global harmonic or conformal parametrizations with cones; among these, [Dong et al. 2006] and [Ben-Chen et al. 2008] and [Springborn et al. 2008] present automatic cone placement. [Jin et al. 2008; Lai et al. 2009] parameterize a surface by flattening the metric by uniformization using *Ricci flow*, and concentrating curvature at user-specified cone locations. These conformal parametrization methods can only handle boundary constraints for simply connected domains, not internal features or multiply-connected boundaries.

3 Background

We start with summarizing basic concepts we use in the paper. As our formulations make significant use of differential geometry, we refer the reader unfamiliar with these concepts to [O’Neill 2006].

Aligned seamless parametrizations. A *global parametrization* f of a mesh M is obtained by cutting it to a topological disk M_c , and mapping it to the plane. The points on the cut (*seam*) consisting of a number of seam curves, get two or more distinct parametric values. A global parametrization f defines a metric $\nabla f^T \nabla f$ on

the mesh, which is required to be flat at all points of the seam, excluding isolated points (*cones*). In this case, the metric does not depend on the choice of seam curves, as long as the cones remain the same. Conversely, a *cone metric* g on a surface M is a metric with zero discrete Gaussian curvature everywhere except at the set of cones $C = \{c_j\}$, $j = 1 \dots M$, with cone angles $\alpha_j > 0$, corresponding to curvatures $K_j = (2\pi - \alpha_j)\delta_{c_j}$. A cone metric defines a global parametrization uniquely, up to the choice of the seam and a global rigid transformation.

A *seamless* parametrization f has the following property: if p_1 and p_2 are the points on the cut mesh M_c corresponding to the same point p on the seam, then the Jacobians of the parametrization at p_1 and p_2 differ by a rotation r by a rotation angle $k\pi/2$, where k is an integer (cf. Figure 2). Intuitively, a parametrization is seamless, if the parametric lines can be continued across seams smoothly (although a u -line may become a v line or the other way around).

Notation for metric. We use *two* metrics on the surface, often together: the Euclidean metric inherited from R^3 , and the parametrization metric g . All quantities related to metric g have subscript g , and those related to the Euclidean metric have no subscript. K denotes the Gaussian curvature, κ the geodesic curvature of a curve.




Figure 2: Seamless parametrization

Parallel transport and holonomy. Formulating the alignment conditions in differential form requires a way to compare vectors at different points on the surface without introducing global coordinates. In a metric g , a *parallel transport* of a vector v along a curve γ can be defined as follows (cf. Figure 3). Suppose the signed angle (in metric g) from the vector to the curve tangent along a segment of length ds of the curve changes by $d\theta$. The tangent vector itself rotates by $\kappa_g ds$, where κ_g is the curve geodesic curvature. The total rotation of the vector in metric g can be computed as $d\theta - \kappa_g ds$, which should be zero for parallel transport. This yields a direct formula for the total change of the angle from the parallel-transported vector to the tangent t_{end} at the endpoint of γ :

$$\theta_{\text{end}} - \theta_{\text{start}} = \int_{\gamma} \kappa_g(s) ds = \kappa_g^{\text{tot}}[\gamma] \quad (1)$$

where the angle can include an integer number of full rotations. We introduce the notation $\kappa_g^{\text{tot}}[\gamma]$ for the total geodesic curvature along a curve in a metric g , and refer to it as the *rotation angle* of a curve with respect to metric g , or parametrization rotation angle.




Figure 3: Parallel transport and holonomy

A seamless parametrization can be defined in terms of *holonomy angles* without relying on the choice of the seam. The integral in

the formula (1) for a positively-oriented contractible loop in a *flat* metric is 2π , and the holonomy angle is the *deviation* from this value (cf. Figure 3):

$$A_g(\gamma) = 2\pi - \kappa_g^{\text{tot}}[\gamma]. \quad (2)$$

For a cone metric, any two homotopic curves in $M \setminus C$ have the same holonomy angle, and for any counterclockwise curve γ encircling a single cone c_j , $A_g(\gamma) = K_j$. The holonomy angle for any curve is completely determined by the holonomy angles on a set of loops γ_i including a set of non-contractible loops γ_i^h in M , and a single loop γ_j^c around each cone c_j ; when we need not distinguish between two types of curves, we drop the superscripts.

A cone metric g is *seamless*, if for all loops γ_i , $A_g(\gamma_i)$ is a multiple of $\pi/2$, in other words,

$$\kappa_g^{\text{tot}}[\gamma_i] = \frac{k_i \pi}{2}, \quad (3)$$

where k_i is an integer (cf. [Lai et al. 2009]).

As-rigid-as-possible energy. One of the primary measures of parametrization quality is distortion: it is desirable for a parametrization to be close to isometric in a metric. We use the *as-rigid-as-possible* (ARAP) energy, to capture deviation from isometry:

$$E^{\text{ARAP}} = \frac{1}{2} \int_M \min_{R \in SO(2)} \|\nabla f - R\|^2 \quad (4)$$

where the integration is over the surface. As it is shown in [Liu et al. 2008], the integrand can be expressed directly in terms of *signed* singular values σ_1 and σ_2 of ∇f , as $(\sigma_1 - 1)^2 + (\sigma_2 - 1)^2$. For a conformal metric with metric tensor $g = e^{2\phi} I$, the energy $E^{\text{ARAP}}(g)$ can be approximated by a *quadratic energy*:

$$\frac{1}{2} ((\sigma_1 - 1)^2 + (\sigma_2 - 1)^2) = (e^\phi - 1)^2 \approx \phi^2$$

4 Parametrization alignment using connection 1-forms

We describe a unified view of aligned conformal metric and *variation-minimizing* cross-fields (that is, cross-fields minimizing an energy of the type defined in [Ray et al. 2008] and [Bommes et al. 2009]). Our formulation is based on [Bunin 2008b], describing alignment constraints on conformal metric and presenting the connection form view of cross-fields [Crane et al. 2010; de Goes and Crane 2010]. Formulating the aligned-metric problem in connection form allows us to handle all types of constraints in a unified manner, and naturally extend, the greedy approach of [Myles and Zorin 2012] to the aligned metric setting.

We give a basic description in the case of smooth surfaces, as in this case the relation between cross-fields and conformal metric is most transparent, and then describe our algorithm along with the discretization.

Parametrization and metric alignment. Suppose a set of smooth *feature curves* Γ_m , $m = 1 \dots M$, is marked on the surface, including the boundary curves (a single boundary loop with corners can be partitioned into multiple curves). The curves intersect at the endpoints only, and they form nonzero angles at intersection points. We say that the parametrization is *feature-aligned* if each curve Γ_i is an isoparametric curve along the u or v direction (note that if a seam crosses the curve, its parametric direction may change). We refer to the part of the boundary for which feature alignment constraints are not specified as *free*; this may include isolated points in the interior of the mesh.

We say that a cone metric is *aligned* with a set of feature curves, if the corresponding parametrization, for a choice of rigid transformation is aligned. We will show how the alignment constraints

can be expressed in terms of the metric directly, without using the parametrization f .

In our discussion, unless otherwise noted, we cut the mesh along feature curves converting each curve into two boundary curves, so we do not make a distinction between features and boundaries.

Boundary and feature alignment in differential form. Consistently with [Bunin 2008b], all boundary curves are assumed to be positively (counter-clockwise) oriented, with $\mathbf{t}$ and $\mathbf{n} = \mathbf{t}_\perp$ denoting, respectively, the unit tangent and normal for a curve (boundary or otherwise). Consequently, boundary normals point inward.

Observe that parametric lines are geodesic in metric g , i.e. have zero curvature. In addition, if two lines intersect, the angle between them in g has to be a multiple of $k\pi/2$. If all feature curves form a connected set, these two conditions would be sufficient to ensure that each of them is aligned with one of two parametric directions, i.e., that the parametrization is aligned.

For multiply connected sets we need to find the parametric-domain angle between two disjoint isoparametric curves Γ_1 and Γ_2 . Consider a curve γ connecting these curves, starting at p_1 and ending at p_2 . Let θ_i denote the angle from the tangent vector of γ to the tangent vector of Γ_i at p_i . The condition that ensures correct relative orientation of two curves Γ_1 and Γ_2 is [Bunin 2008b]

$$\kappa_g^{\text{tot}}[\gamma] = \frac{k_\gamma \pi}{2} + \theta_2 - \theta_1 \quad (5)$$

where again the curvature and length parametrization are with respect to metric g . This leads to the following definition of alignment in terms of metric:

- **Boundary curvature.** On a feature curve Γ_i , $\kappa_g = 0$, except at the endpoints. At an endpoint p where two curves meet, the curvature of the joint curve is $\kappa_g(p) = k\pi\delta_p/2$, for an integer k ; in a more concise form, on the boundary $\kappa_g = \sum_i k_i \pi \delta_{p_i}$, where $k_i \pi/2$ are angles at curve joints.
- **Component alignment.** For any path γ connecting boundary curves, (5) holds. The integer k_γ may depend on the choice of the curve γ , but as long as all homology basis loops in M have homology angles $k\pi/2$, it is sufficient to enforce this condition for a single curve: the rest will automatically be of the same form.

Figure 4 illustrates boundary alignment paths for a planar shape.




Figure 4: Boundary alignment paths

We note the similarity between the holonomy loop condition (3) and boundary alignment constraints; the only difference is the presence of the $\theta_2 - \theta_1$ term. For brevity, we will treat these constraints as a single group, and define $\Delta\theta = \theta_2 - \theta_1$ for homology loops to be zero.

If the genus of the surface is g , there are n cones and we have m boundary loops, then the total number of *integral* constraints is $n + 2g + m - 1$ (holonomy loops and alignment).

4.1 Conformal metrics and variation-minimizing cross-fields

Aligned conformal metric. The case of particular importance to us is the case of *conformal* cone metrics. For these metrics, $g = e^{2\phi}I$, where ϕ is the *log scale factor*. In this case the angle change expressions for parallel transport are *linear* in ϕ . More specifically, the geodesic curvature is given by

$$\kappa_g = e^{-\phi}(\kappa - \langle d\phi, \mathbf{n} \rangle),$$

where we use $\langle \cdot, \cdot \rangle$ for evaluation of a 1-form on a vector, and length element $ds_g = e^\phi ds$.

This leads to the following set of conditions derived in [Bunin 2008b]:

- (C1) **Gaussian curvature.** $\nabla^2 \phi = K - K_g$, where K_g is zero except at cones, where it is $\frac{k_j \pi}{2} \delta_{c_j}$.
- (C2) **Boundary curvature.** $\langle d\phi, \mathbf{n} \rangle = \kappa - \kappa_g$ on all boundaries Γ_i . κ_g is zero everywhere on the boundary except joints p_j , where it is $\frac{k_j \pi}{2} \delta_{p_j}$.
- (C3) **Homology loops and component alignment.** $\int_{\gamma_i} \langle d\phi, \mathbf{n} \rangle ds = k_i \pi / 2 + \kappa^{\text{tot}}[\gamma_i] + \Delta \theta_i$,

Proposition 1. *In general, no conformal map satisfies seamless, boundary, and feature-alignment constraints.*

Proof. In the absence of free boundary points, the scale factor ϕ is determined by the Poisson equation (C1) $\nabla^2 \phi = K - K_g$ on M , with the Neumann condition (C2) $\langle d\phi, \mathbf{n} \rangle = \kappa - \kappa_g$. Under mild assumptions on the surface and boundary, this problem has a unique solution defined up to a constant [Dindoš 2008], and the remaining constraints on holonomy angles of homology loops and alignment constraints (C3) cannot be satisfied. $\square$

Cross-fields and frames. A *cross-field* is an assignment of quadruples $(\mathbf{e}, \mathbf{e}_\perp, -\mathbf{e}, -\mathbf{e}_\perp)$ of unit-length vectors, separated by right angles, to all points on the surface excluding isolated points (singularities). Cross-fields are commonly used as a guidance field for a global parametrization, with cones located at field singularities and gradients of parametric coordinates aligned in least squares sense with the field vectors. For a cross-field, the alignment constraints are straightforward: we say that a cross-field is *aligned* with a set of feature curves, if the tangents of the feature curve are parallel to one of the vectors of the cross-field at each point.

While in general, it is impossible to separate a cross-field into four continuous unit-vector fields, on a simply connected domain it is possible; in particular it can be done on the cut surface M_c . We assume that $\mathbf{e}$ is a smooth field on M_c ; $(\mathbf{e}, \mathbf{e}_\perp)$ defines an orthonormal *frame* on the surface. At two points p_1 and p_2 corresponding to the same point on a seam curve the frames are related by a $k\pi/2$ rotation, with k dependent only on the curve.

To relate to the metric formulation, we can similarly define the alignment conditions for a cross-field using a local differential quantity, the (Euclidean metric) *connection 1-form*. While this makes the constraints more complicated, in this form, the close relation between cross-field and conformal maps becomes clear.

The *connection 1-form* for the frame is identical to the one-form defined directly for the discrete case in [Crane et al. 2010]. Specifically, for a global frame $(\mathbf{e}, \mathbf{e}_\perp)$ on $M \setminus C$, we define the 1-form ω by

$$\langle \omega, \mathbf{v} \rangle = \mathbf{e}_\perp \cdot \nabla_{\mathbf{v}} \mathbf{e}, \quad (6)$$

where $\nabla_{\mathbf{v}}$ is the surface derivative of $\mathbf{e}$ in the direction $\mathbf{v}$. Evaluated on a tangent vector $\mathbf{v}$, the form produces the rate of rotation of the frame in direction $\mathbf{v}$. Then the relative orientations of the frames at points p_1 and p_2 can be obtained by integrating this rotation along a path γ connecting these points. Observe that for a displacement ds

in a direction $\mathbf{v}$ along a curve γ , the change $d\theta$ in the angle between $\mathbf{e}$ and the tangent to the curve can be represented as $\omega(\mathbf{v})ds - \kappa ds$, the sum of the rotation of the frame, and the rotation of the tangent vector due to the geodesic curvature of γ . Not any 1-form corresponds to a cross field, which leads to additional conditions on the integrals of the 1-form on loops. For a closed loop, the full change in this angle should be of the form $k\pi/2$, as after a full circle the cross-field should be mapped to itself. For loops γ_ϵ of length ϵ around cones, in the limit the condition on the loop leads to one of Cartan's structure equations for frames, which can be written in the form $d\omega = (K - K_g)\beta_1 \wedge \beta_2$, where 1-forms β_i correspond to the frame: $\langle \beta_1, \mathbf{e} \rangle = 1$, $\langle \beta_1, \mathbf{e}_\perp \rangle = 0$, $\langle \beta_2, \mathbf{e}_\perp \rangle = 1$, $\langle \beta_2, \mathbf{e} \rangle = 0$.

- (F1) **Gaussian curvature.** $\star d\omega = K_g - K$.
- (F2) **Boundary alignment.** $\langle \omega, \mathbf{t} \rangle = \kappa - \kappa_g$ along all Γ_i .
- (F3) **Holonomy loop and boundary component alignment.** $\int_{\gamma_i} \langle \omega, \mathbf{t} \rangle ds = k_i \pi / 2 + \kappa^{\text{tot}}[\gamma] + \Delta \theta$.

We note the similarity of the alignment constraints and constraints for the conformal maps, which we will clarify shortly, once we introduce a class of *variation-minimizing* cross-fields.

Variation-minimizing cross-fields. Clearly, there are many cross-fields satisfying the constraints above, as these restrict the field orientation on the boundary curves only. We can choose a unique field among these by picking the one for which $\|\omega\|_2^2 = \int \omega \wedge \star \omega$ is minimal, i.e. the integral of the rotation rate over all directions and all points is as small as possible. This is exactly the functional used in [Crane et al. 2010] as well as in [Ray et al. 2008; Bommers et al. 2009] in different variables.

To make the condition on ω explicit, we use the Hodge decomposition of ω . As we are considering a domain with boundary, the spaces used in the decomposition require specification of boundary conditions, and several decomposition choices are possible [Capell et al. 2006]. If a form ω satisfies Dirichlet (tangent) conditions on the boundary $\langle \omega, \mathbf{t} \rangle = b$,

$$\omega = d\xi + \delta\psi + h \quad (7)$$

where ξ is a function vanishing on the boundary, ψ is a 2-form such that $\delta\psi$ satisfies boundary conditions, and h is a harmonic 1-form, i.e. it satisfies both $\delta h = 0$ and $dh = 0$ and vanishes on the boundary.

We observe the following:

Proposition 2. *The minimal-norm form ω satisfying alignment condition has zero closed part $d\xi$:*

$$\omega = \delta\psi + h; \quad (8)$$

The co-closed and harmonic parts are defined uniquely by the boundary constraints.

The proof of this proposition can be found in the appendix. In particular, we observe that the variation-minimizing field has a co-closed 1-form:

$$\delta\omega = 0 \quad (9)$$

Conformal metric alignment constraints and cross-fields. Comparing conformal metric constraints and constraints for the field expressed in terms of a 1-form, we observe that they are very similar. This is not that surprising: observe that a conformal parametrization defines a frame $(\mathbf{e}, \mathbf{e}_\perp)$ by setting $\mathbf{e} = f_u / \|f_u\|$, the unit parametric line tangent for u direction.

Proposition 3. *For a conformal metric, the connection 1-form of the frame is $\omega = -\star d\phi$ (i.e. $\star\omega = d\phi$), where $\star$ is the Hodge star; i.e., with standard identification of 1-forms and vectors, ω is the gradient of ϕ rotated -90 degrees (see appendix).*

Comparing this formula to the Equation 8, we observe that the difference between the connection 1-form for a general variation-minimizing cross-field and the field corresponding to a conformal

map is the harmonic part. In other words, the following proposition holds:

Proposition 4. *An aligned variation minimizing cross-field corresponds to an aligned conformal map if and only if its (co-closed) connection 1-form is co-exact.*

The relationship between conformal maps and cross-field is the one between closed and exact forms: the difference is the space of harmonic forms. Moreover, the closest in the 1-form norm conformal map for a given cross-field can be obtained by solving a Poisson equation. Setting $\phi = -\star\psi$, and using $dh = 0$ and the fact that h vanishes on the boundary, we get

$$\nabla^2 \phi = \star d\star d\phi = -\star d\omega$$

with Neumann conditions $\langle d\phi, \mathbf{n} \rangle = \langle \omega, \mathbf{t} \rangle$ on the boundary.

There are several important practical consequences of our observations about the relation between 1-forms and conformal metrics.

1. ϕ^2 for ϕ computed from ω can be used as a measure of metric distortion;
2. $\|\omega\|_2$ yields direct control over smoothness of the cross-field and, indirectly, the number of cones;
3. The harmonic part $h = \omega + \star d\phi$ measures how close is a field to a conformal metric.

In the next section, we discuss how these observations can be used to derive an algorithm following the ideas of [Bommes et al. 2009] and [Myles and Zorin 2012].

Remark. A useful interpretation of the co-closed connection 1-form in terms of *similarity* structure on the surface is shown in Figure 5. On the cut mesh M_c which is simply connected, $\star\omega$ can be integrated to yield a globally defined ϕ^c , but with mismatched values on the cut. Due to co-closedness condition, the jump in the resulting value of ϕ^c is *constant* along the cut. In other words, one can think about 1-form as defining a flat conformal metric on any simply connected chart on $M \setminus C$, with transitions between charts given by similarity transforms, instead of isometries as for well-defined global parametrization. In our case, transition from M_c to itself across the cut is the similarity transform, with scaling given by e^s , where $s = \phi_1^c - \phi_2^c$, the difference in values of ϕ^c across the cut.




Figure 5: Interpretation of a closed, but non-exact 1-form $\star\omega$ in terms of a similarity structure. The difference in scale of the parametrization on M_c computed from ω is constant across the cut.

5 Algorithm

5.1 Overview

The idea of the algorithm is to evolve the connection 1-form ω , gradually expanding the domain Ω^f on which the constraints (F1–F3) are enforced, and resolving the energy minimization problem each time. The central element of the algorithm is the definition of the domain Ω^f and how it is expanded.

Our discretization, described below, necessarily introduces an approximation in the meaning of ϕ , which is used only to measure

distortion and cannot be used to compute exact parametric lengths in contrast to, e.g. [Springborn et al. 2008]. On the other hand, our discretization of (F1–F3) and ω induces an explicitly constructable feature-aligned cross-field (cf. Appendix B) from which we compute the final parametrization.

We use *dual paths* (chains of triangles) to discretize all paths (homology loops, alignment paths, and boundary loops). In the discrete case, all constraints can be thought of as loop or path constraints (the Gaussian curvature constraints corresponds to a loop around a vertex, and the Neumann boundary condition corresponds to a path connecting boundary edges on the two sides of a boundary vertex), with loops defined as chains of triangles (Figure 6).




Figure 6: The discretization of interior or boundary curvature condition at a vertex v can be written in terms of the corresponding loop or path γ_v around the vertex.

The domain Ω^f thus can be naturally defined as a collection of paths for which rotation or holonomy angles are fixed. The triangles of these paths form Ω^f . More specifically, we consider the basis of paths on the *whole* mesh M , consisting of $2g$ homology loops, $\max\{m-1, 0\}$ alignment paths, and V loops/paths around vertices shown in (6). The domain Ω^f is characterized by a vector I^f of indices of loops and paths for which a constraint was fixed. Fixing a constraint means choosing an integer k_i for a loop or a path (F3) or a curvature (F1). As the initial domain, we start with a set of loops around all boundary components, i.e. all paths corresponding to Neumann constraints are included in the beginning.

Loops and paths can be added to a domain in different ways; we use a variation of [Bommes et al. 2009] described in greater detail below.

At each update step after new paths are added to I^f , we compute the variation-minimizing 1-forms ω , as close as possible to a conformal metric, i.e. minimizing h , and satisfying all constraints.

Among these 1-forms, we choose the one that minimizes an energy

$$\begin{aligned} E^{\text{reg}} &= \alpha_{\text{smooth}} \|\omega\|_2^2 + \frac{1}{A_M} \|\phi\|_2^2 \\ &= \alpha_{\text{smooth}} (\|d\phi\|_2^2 + \|h\|_2^2) + \frac{1}{A_M} \|\phi\|_2^2 \end{aligned} \quad (10)$$

where A_M is the surface area, which we use to make the parameter non-dimensional. The parameter α_{smooth} is the only parameter of the algorithm. In Section 6 we show that the parameter is necessary, in a sense that with $\alpha_{\text{smooth}} = 0$ a very low (if the mesh can be refined, arbitrarily low) distortion can be achieved at the expense of introducing *cone chains*. Setting α_{smooth} high reduces the number of cones while permitting greater distortion. As a result of recomputing ω , loop holonomies and path rotations are modified, and new paths are selected for inclusion in I^f .

5.2 Discretization

The core computational part of the algorithm consists of solving an energy-minimization problem with constraints. We need to discretize the log scale factor ϕ , the harmonic 1-form h , and the constraints (F1–F3) on the form, including discretization of integrals on loops and alignment paths.

Although the 1-form ω is defined on Ω^f only, we assign variables ϕ and h for the whole mesh as detailed below; outside Ω^f , the constraints are not imposed.




Figure 7: The Hodge dual ω of a 1-form ψ is given via (11).

Discretizing 1-forms and the conformal factor. We discretize the conformal factor ϕ as a vector $\bar{\phi}$ of values at the V vertices of the mesh. The standard 1-form discretization on edges is used for $d\phi$, with $\phi_j - \phi_i$ assigned to the edge e_{ij} . The connection 1-form $-\star d\phi$, requires defining a Hodge star. Recall that on 1-forms identified with vectors, the Hodge star acts as a 90-degree rotation. Similarly to the non-conforming FEM approach of [Pinkall and Polthier 1993; Polthier 2000], we define the Hodge dual form by values on triangle midlines (Figure 7). We note that in the absence of obtuse triangles, this can be viewed a different representation of commonly used DEC dual forms defined on dual edges connecting circumcenters. Explicitly, the relation between values of a 1-form ψ defined on edges of the mesh, and its Hodge dual ω , defined on midline segments, is given by

$$\omega_i = \frac{1}{2} (\psi_k \cot \alpha_k - \psi_j \cot \alpha_j), \quad (11)$$

where index of the midline segments coincide with parallel edges, angles α_i are indexed by vertex, (i, j, k) is a circular permutation of $(1, 2, 3)$, and a counterclockwise orientation is chosen on edges.

It is more efficient to represent ω directly as a co-closed and harmonic part using its Hodge decomposition, which also applies in the discrete case: $\omega = -\star d\phi + h$. A low-dimensional representation of h can be obtained as follows. Consider a set γ_i of $n_h := 2g + \max\{m - 1, 0\}$ non-contractible basis loops, where m denotes the number of boundary components and g the genus. The harmonic 1-form h can be written with a coefficient vector $\bar{h}$:

$$h = \sum_{\gamma_i \in n_h} \bar{h}_i \eta_i, \quad (12)$$

where each closed harmonic form η_i is an exact form on the cut surface $M \setminus \gamma_i$, but not across γ_i . η_i can be determined by computing a harmonic function u_{η_i} that has zero tangent conditions along the boundary of M and a discontinuity on γ_i with matching derivatives on both sides so that Δu_{η_i} is still well-defined on γ_i . We discretize u_{η_i} with non-conforming linear elements, with degrees of freedom on edge midpoints, and two different values for midpoints on γ_i related by a constant shift. This ensures the continuity of $\eta_i = du_{\eta_i}$.

Discretizing integrals on paths. For the purposes of geodesic curvature computation, we view these paths as consisting of midline segments: a path γ is formed by segments $\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_r$ (cf. Figure 4). This path representation naturally matches the 1-form ω , as its values are defined on the same segments $\mathbf{m}_j$. Thus, the integral of ω along a discrete path γ is precisely the sum of the values of ω along this path (cf. [Crane et al. 2010]), with sign adjusted for direction:

$$\int_{\gamma} \langle \omega, \mathbf{t} \rangle ds = \sum_{j=1}^r \pm \omega_j \quad (13)$$

The total geodesic curvature along a path γ is computed as the sum of angles between sequential midline segments, computed after the corresponding triangles are rotated about the shared edge to be coplanar:

$$\kappa^{tot}[\gamma] = \sum_{j=1}^{r-1} \beta_j \quad (14)$$

Discretizing constraints (F1-F3).

We discretize the constraints on the whole mesh M at once, on a complete basis of loops and paths, instead of adding constraints one-by-one. Constraints for loops not included in I^f are disabled by introducing a free variable for the corresponding turning/holonomy angle.

The Gaussian curvature (F1) $-\star d\omega = \nabla^2 \phi = K - K_g$ and the boundary alignment conditions (F2) $\langle \omega, \mathbf{t} \rangle = \langle d\phi, \mathbf{n} \rangle = \kappa - \kappa_g$ together form a Poisson system for ϕ with Neumann boundary conditions, which are discretized with well-known cotangent weights.

For a vertex v , let $\text{VA}(v)$ denote the sum of angles at the vertex on all neighboring triangles. If v is an interior vertex, its Gaussian curvature is $2\pi - \text{VA}(v)$; if it is on the boundary, its geodesic curvature is $\pi - \text{VA}(v)$. The cone metric geodesic and Gaussian curvatures are of the form $k_i^{\text{vert}}\pi/2$. Denoting by b the vector of V geodesic and Gaussian curvatures defined in this way, and by $\bar{k}^{\text{vert}}$ the vector of integers k_i^{vert} , the linear system for (F1-F2) in matrix form is

$$L\bar{\phi} - \frac{\pi}{2}\bar{k}^{\text{vert}} = b$$

where L is the standard $V \times V$ Laplace matrix discretizing $-\nabla^2$.

The n_h constraints (F3) along homology loops and boundary component alignment paths require $\kappa^{tot}[\gamma]$ computed using (14), and $\Delta\theta$ as defined in Section 4. This yields the right-hand-side of the equations in matrix form:

$$B_{\phi}\bar{\phi} + B_h\bar{h} - \frac{\pi}{2}\bar{k}^{\text{path}} = c.$$

where $\frac{\pi}{2}\bar{k}^{\text{path}}$, of length n_h , are the homology loop holonomy angles and relative alignment turning angles, and c is $\Delta\theta - \kappa^{tot}[\gamma_i]$ for the loop or path γ_i .

The (i, j) entry of the $n_h \times n_h$ matrix B_h is $\int_{\gamma_i} \langle \eta_j, \mathbf{t} \rangle ds$ computed via (13). The $n_h \times V$ matrix B_{ϕ} is computed as $-SR^*D^d$, where D^d is the discrete exterior derivative matrix, R^* computes the Hodge dual with entries according to (11), and S is a matrix of ± 1 for path integration via (13).

We collect these systems together in a more compact form for a complete set of $V + n_h$ constraints for the energy minimization problem.

$$C_{\phi}\bar{\phi} + C_h\bar{h} - \frac{\pi}{2}\bar{k} = d, \quad (15)$$

where $C_{\phi} := \begin{bmatrix} B_{\phi} \\ L \end{bmatrix}$, $C_h := \begin{bmatrix} B_h \\ 0 \end{bmatrix}$, $\bar{k} := \begin{bmatrix} \bar{k}^{\text{path}} \\ \bar{k}^{\text{vert}} \end{bmatrix}$, and $d := \begin{bmatrix} c \\ b \end{bmatrix}$.

The variables in this system are $\bar{\phi}$, $\bar{h}$ and $\bar{k}$. $\bar{k}$ can be partitioned into subsets $\bar{k}^{\text{free}}$ and $\bar{k}^{\text{fixed}}$, the latter consisting of elements with indices included in I^f . $\bar{k}^{\text{free}}$ are free variables disabling constraints outside Ω^f . While adding constraints one-by-one is completely equivalent to adding and removing free variables, this way of keeping track of Ω_f allows to keep the constraint matrices fixed.

Change of basis of paths. Rather than using $\bar{k}$ directly to define I_f , we choose a different basis in the space of paths, with coefficients $\bar{p}$. The reason for this change of basis is empirical: we have observed that a greedy algorithm for fixing constraints yields somewhat better results if the basis of loops, rather than including vertex loops directly, consists of loops “growing” from boundary in a way similar to the definition of [Bommes et al. 2009].

To define a new basis, we construct a forest $\mathcal{F}$ of dual trees; the root of each tree is a boundary triangle, and the tree contains a chain of triangles along the boundary component. The forest spans the whole surface (cf. Figure 8). A dual edge *not* in the forest closes a loop or connects two trees in the forest by an alignment path. All




Figure 8: The dual forest $\mathcal{F}$ is in brown. The two blue edges correspond to loops whose holonomies are set by boundary curvature.

such loops and paths are independent; we denote this set of paths G . We denote by $\frac{\pi}{2}\bar{p}$ the vector of holonomies and turning angles of these paths.

There is a unique linear transformation P mapping $\bar{p}$ to $\bar{k}$: $\bar{k} = P\bar{p}$. We can greatly simplify this transformation by ensuring that alignment paths and holonomy loops in both bases are the same. More specifically, we identify a set of $\max\{m-1, 0\}$ curves $G^{\text{comp}} \subset G$ for the relative alignment of boundary components. Adding the associated dual edges outside $\mathcal{F}$ to $\mathcal{F}$ creates a spanning tree from which we identify a basis $G^{\text{hom}} \subset G$ of homology loops using the technique of [Eppstein 2003]. We use $G^{\text{path}} = G^{\text{comp}} \cup G^{\text{hom}}$ of homology loops and boundary component alignment paths for constraints. If $\bar{p}^{\text{path}}$ is the subvector of $\bar{p}$ consisting of entries for homology loops and paths, then $\bar{p}^{\text{path}} = \bar{k}^{\text{path}}$. The transformation between $\bar{p}$ and $\bar{k}^{\text{vert}}$ is easily established using the algorithm described in [Myles and Zorin 2012].

This yields the final form of the constraint

$$C_\phi \bar{\phi} + C_h \bar{h} - \frac{\pi}{2} P \bar{p} = d, \quad (16)$$

The turning angles for the boundary alignment paths corresponding to geodesic curvature constraints are set from the beginning, fixing the loops denoted by the blue dual edges in Figure 8.

5.3 Algorithm

Discrete system for ω . We now evolve $\omega = \star d\phi + h$ by minimizing the quadratic energy (10), subject to a set of linear constraints that consists of two parts: (1) alignment constraints (F1-F3) collected in (16) and (2) harmonic part minimality. The latter constraint, defined in detail below, is obtained as the standard constrained optimality condition.

To enable only the constraints in I^f , we split the vector $\bar{p}$ as $[\bar{p}^{\text{fr}}, \bar{p}^{\text{fix}}]$, where the first component consists of free variables for unfixed constraints, and the second from constraints with already defined angles.

The as-exact-as-possible forms ω minimize the norm of $\|h\|^2 = \bar{h}^T M_1 \bar{h}$, where the (i, j) entry of M_1 is $\langle \eta_i, \eta_j \rangle_2$ computed using the discrete L^2 1-form inner product. The usual constrained minimization optimality condition leads to the constraint on h and ϕ :

$$\begin{aligned} M_1 \bar{h} + C_h^T \lambda &= 0, & C_\phi^T \lambda &= 0, \\ C_\phi \bar{\phi} + C_h \bar{h} - \frac{\pi}{2} P^{\text{fr}} \bar{p}^{\text{fr}} &= d + \frac{\pi}{2} P^{\text{fix}} \bar{p}^{\text{fix}} \end{aligned} \quad (17)$$

where λ is the Lagrange multiplier for the constraint, and $P = [P^{\text{fr}} P^{\text{fix}}]$.

As our constraint is derived from minimizing the norm of the harmonic part, this term can be dropped from (10), and subject to the

constraints above, we minimize the discretized energy

$$E^{\text{reg}, d} = \bar{\phi}^T \left(\alpha_{\text{smooth}} L + \frac{1}{A_M} M_\phi \right) \bar{\phi} \quad (18)$$

where M_ϕ is the diagonal matrix of vertex areas which sum to A_M . The final form of the problem that we solve is

Minimize (18), with respect to $\bar{\phi}$ and $\bar{p}^{\text{fr}}$ subject to constraints (17).

An analysis of the constraints shows that as long as $\bar{\phi}$ has a sufficient number of degrees of freedom not fixed by (15), the harmonic part minimization (17) is solved by $h = 0$; in other words, an exact 1-form ω satisfying all constraints exists.

Once the remaining number of degrees of freedom not fixed by (15) is less than needed to have $h = 0$, (17) determine h and $\bar{\phi}$ uniquely, so solving the optimization problem is unnecessary; the conformal part is determined by the harmonic part minimality. We found it convenient not to separate these stages and retain a unified formulation.

Complete algorithm.

Algorithm 1 Outline of the loop-fixing algorithm

- 1: Construct a spanning forest of dual trees $\mathcal{F}$, and from it, collection of variables $\bar{p}$ corresponding to path basis G .
 - 2: Initialize I^f to the set of boundary loops
 - 3: **while** $I^f \neq G$ **do**
 - 4: Add to I_f a set of unfixed paths E^{fix} in G with rotation angles close to integer multiples of $\pi/2$ (see the solver section below).
 - 5: Define free variables $\bar{p}^{\text{fr}}$ for paths not in I^f .
 - 6: Determine integer $\bar{p}_i$ for turning angles for paths in E^{fix} .
 - 7: Solve for $\omega = -\star d\phi + h$ and $\bar{p}^{\text{fr}}$ satisfying alignment and minimal harmonic part constraints, and minimizing the energy (18).
 - 8: **end while**
-

Once the cones, homology loop holonomies, and alignment angles are determined by this algorithm we proceed exactly as in [Myles and Zorin 2012], computing a seamless global ARAP parametrization satisfying all constraints. We defer the details to Appendix B, which also explicitly relates the discrete connection ω to its natural cross-field representation.

Using a Mixed Integer solver. The description of the evolution above is equivalent to using the Mixed Integer solver of [Bommes et al. 2009] on our system. However, that solver requires an explicit elimination of constraints to yield a symmetric positive definite linear system. The constraints (F1-F2) on ϕ , in particular, are not amenable to elimination as the inverse of the Laplace matrix is dense. Thus, we implemented a slower stand-in mixed integer solver in Matlab using a similar algorithm as [Bommes et al. 2009] with Lagrange multipliers (yielding a symmetric indefinite system) and without local iterative updates. Specifically, Steps 4 and 6 in Algorithm 1 collect E^{fix} so that the total rounding is less than 0.5 and round these variables at once.

Cone location optimization. After all holonomies and turning angles are determined, cone positions on the surface are optimized in the spirit of [Bommes et al. 2009] to reduce distortion further. In this procedure, each cone is moved to a neighboring vertex if it decreases the distortion energy $\|\phi\|^2$ of the resulting connection form ω . This can be iterated on all cones until the energy can no longer be reduced. This stage can be followed by an optional clean-up procedure as in [Myles and Zorin 2012] for collapsing cones not on feature curves that are 1 or 2 edges apart.

Remark. To reproduce results similar to [Bommes et al. 2009], one may remove the constraint minimizing the norm of h and choose a large α_{smooth} weight in E^{reg} . See Figure (15) for an example.

Extension to general guiding feature fields. So far, we have considered the problem with alignment constraints specified on curves. However, in practice it may be useful to define pointwise alignment directions on faces or edges of the mesh, without structuring them into curves. We observe that if there is an area of the surface where the alignment directions are fixed, in general there is no parametrization that matches the field directions exactly; thus we can only hope to compute the field in least-squares sense there. Field singularities determine the cones in these areas.

Thus, for the purposes of cone determination we take the following approach: we partition the surface into connected domains, with field directions fixed on the boundary of each domain only. The areas with the field everywhere in the interior are not considered and cones on these areas are set by the field. We proceed to compute to determine cones on each domain independently, and finally combine the cones together and compute an ARAP parametrization. The only modification of the algorithm needed for each subdomain is changing the curvature of the boundary, used in the boundary constraint (F2) to the rate of rotation of the field on the boundary, and measure $\Delta\theta$ in constraint (F3) with respect to the field directions on the connected boundaries.

The final parametrization is again obtained as in [Myles and Zorin 2012], by computing an ARAP parametrization with fixed cross-field directions as constraints.

6 Tradeoff between the number of cones and distortion

A fundamental question in choosing cone locations is how low one can make the distortion, and what additional conditions need to be imposed on the parametrization if the lowest-distortion parametrization has an excessive number of cones.

For example, if there are no restrictions on cone and holonomy angles, there is always an isometric parametrization of a mesh, obtained by cutting it along a dual spanning tree and unfolding isometrically along the tree. Effectively, each vertex becomes a cone. With $k\pi/2$ constraints on the degrees of freedom, the answer is less obvious: clearly, if all vertices are cones, the distortion is excessive for a smooth mesh with low discrete Gaussian curvature at vertices, as the 1-ring around a nearly flat vertex mapped to an angle $3\pi/2$ or $5\pi/2$ instead of 2π , inevitably has high distortion. Thus there is an optimal number of cones somewhere between topologically required and all vertices.

Somewhat surprisingly, the answer to this question is that for distortion-minimizing parametrization cones form *cone chains*; in the limit case of smooth surfaces, cones converge to lines on the surface (Figure 9).

Moreover, the following proposition holds:

Proposition 5. *For a closed smooth surface with bounded total curvature, and any $\epsilon > 0$, there is a seamless parametrization which has total metric distortion $E^{\text{ARAP}} < \epsilon$.*

A detailed outline of a proof is in the appendix, but the main observation is that one can convert a non-seamless global parametrization (which can have arbitrarily small distortion) into a seamless ones by a local perturbation near the seam. The construction in the proof also indicates that for a sequence of meshes with average edge length decreasing as h , one can expect the distortion of the distortion-minimizing solution with chains of cones to decrease as $\sqrt{h}$ which is confirmed in our experiments on simple shapes.

These observations show that simply minimizing metric distortion is not a practical option for most applications: while the resulting set of cones is not dense on the whole surface it may be dense along curves, if the least possible distortion is desired. The *form smoothness* parameter, α_{smooth} allows the user to choose a preferable tradeoff between the number of cones and distortion.




Figure 9: Chains of cones for low-distortion parametrization.

7 Results and Discussion

Comparison to previous work. As our method can be regarded as an extension of [Bommes et al. 2009] and [Myles and Zorin 2012], our primary comparisons are to these two methods. As previously mentioned the former (which we label *MIQ*) can be thought of as a special case of our method with no harmonic part minimization and no metric term in the energy. While [Bommes et al. 2009] does not use an isometry measure directly in its functional, we note that it is closely related to $\|d\phi\|_2^2$, i.e. a measure of the *gradient* of metric distortion. As a result, to an extent the distortion is factored in. The field-aligned parametrization itself can be regarded as a single step of as-rigid-as-possible parametrization [Myles and Zorin 2012]. Its degree of distortion is similar to the related methods of [Ray et al. 2006] and [Kälberer et al. 2007], so we focus on comparisons with this method. In all figures, distortion is measured by the following average measure

$$\sqrt{\frac{1}{A_M} \int_M \min_{R \in SO(2)} \|\nabla f - R\|} \quad (19)$$

As it is discussed in Section 6, the method of [Myles and Zorin 2012] does not have a smoothness term, but has implicit regularization by a different mechanism: the flattening stage. The downside of this approach is that there is no way to reduce distortion below the one produced by the method if desired. As shown in Figure 14, on models without alignment constraints, our method produces similar results for a moderate choice of the coefficient of the regularization term $\|\omega\|_2$. On the other hand, for high values of the parameter the results are similar to [Bommes et al. 2009] Figure 15.

At the same time, for a number of cases, even for a large value of smoothing parameter, the difference is quite significant. An important simple example is shown in Figure 10. In this case *boundary alignment constraints force parametrization collapse in the absence of cones*. The reason for this can be easily understood: suppose on the ring part of the strip is h_1 , and on the protruding flap it is h_2 ; the opposite sides of each part of the shape have to be parallel in the parametric domain if alignment conditions are imposed, so the parametric width, defined as the length of a parametric isoline connecting opposite sides of a strip, remains constant around the ring, and along the flap. Then at the place on the ring where the flap separates, the width has to be $h_1 + h_2$ on one hand and h_1 on the other, i.e. $h_2 = 0$, and the flap parametrization is collapsed.

In our method, the metric component of the energy prevents this collapse from happening as it leads to high metric distortion, and introduces two cones there.

Of particular interest is a comparison with [Bunin 2008a]. Our continuum constraints for conformal maps (C1-C3) are essentially identical to [Bunin 2008b]; however, [Bunin 2008a] uses a completely different formulation for computing a conformal metric on a connected planar domain, using rational function approximation for ϕ , and inverse Poisson formulation for determining cone positions;




Figure 10: For the shape shown in this figure, a singularity-free cross-field exists, but a singularity-free boundary aligned parametrization does not.

the resulting conformal parametrization is aligned in a least-squares sense. We observe that that method also results in a configuration similar to ours – effectively, for close approximation of constraints the inner boundary loop is separated from the outer by a ring of cones (Figure 11).




Figure 11: Example similar to [Bunin 2008a]; a sequence with increasing smoothness factor.

Dependence on parameters. The method has one main parameter, the weight of the smoothing term. Figure 15 shows the effect of varying this term for two examples. For low smoothness, as expected, we get chains of cones described in Section 6. For high smoothness, with no harmonic part minimization, it behaves qualitatively similar to [Bommes et al. 2009].

Figure 13 shows the distributions of distortion on a number of models, and the effect of taking the metric distortion into account compared to optimizing smoothness only via the approach of [Bommes et al. 2009]. We observe that we obtain an improvement in all instances. Further details, including timings, on the models are listed in Table 1. We expect the timings to be dramatically reduced by replacing our Matlab solver with an improved solver described in Section 8.

Figures 1 and 12 compare a quadrangulation generated from ARAP parametrizations on Mixed Integer fields and ours, obtained after rounding seam translations and cone parametric coordinates.




Figure 12: Quadrangulation results for cone locations found our method and field optimization of [Bommes et al. 2009].

Model (Figure)	faces	α_{smooth}	N_{our}	dist_{our}	T_{main}	T_{opt}	N_{MIQ}	dist_{MIQ}
Fandisk (1)	14454	0.01	42	0.226	15.4	2.64	30	0.353
Bulge (12)	2072	0.05	16	0.296	1.85	0.217	8	0.462
Screwdriver (13)	15354	0.01	65	0.228	240	78.0	44	0.349
Screw (13)	11616	0.01	31	0.205	43.4	1.79	23	0.711
Lion head (13)	16834	0.01	50	0.186	302	55.2	62	0.222
Sculpt (13)	7342	0.01	30	0.237	15.9	0.643	8	0.596
Beetle (13)	38656	0.01	72	0.173	108	18.7	50	0.266
Casting (13)	36828	0.01	119	0.211	202	16.0	106	0.303
Mannequin (15)	14472	0.001	90	0.192	240	52.0		
		0.01	66	0.218	198	51.9		
		1.0	42	0.239	184	45.2		
		100.0	44	0.244	193	41.6		
		100.0*	32	0.296	192	41.4	32	0.274
Aircraft (15)	4656	0.001	74	0.175	14.7	5.80		
		0.01	52	0.183	12.9	3.56		
		1.0	43	0.261	13.0	3.03		
		100.0	43	0.288	13.6	5.53		
		100.0*	41	0.271	12.3	3.21	43	0.263

Table 1: The models are listed in the order in which they appear in the figure in parentheses. Column titles: N_{our} and N_{MI} are, respectively, the number of cones in our parametrization and the MIQ field; dist_{our} and dist_{MI} denote, respectively, the average triangle distortions of seamless parametrizations produced by our result and the MIQ procedure; T_{main} and T_{opt} are the times in seconds for holonomy-rounding and the cone location optimization. The entries indicated with 100.0* for α_{smooth} are without the “exact-as-possible” constraints (17) to emulate [Bommes et al. 2009].

8 Conclusion

Using a formulation unifying a number of previous formulations for conformal maps and cross-fields, we constructed an algorithm for global parametrization providing control over metric distortion and, through smoothness of the associated cross-field, over the number of cones. A single parameter makes it possible to span parametrizations from highly distorted but with a minimal (in the sense of field smoothness) number of cones, and nearly-isometric with cone chains. We demonstrate that this trade-off is effectively unavoidable, as by increasing the number of cones until the cone chains are formed we can always decrease metric distortion further.

Limitations. Just as the method of [Myles and Zorin 2012], we do not consider additional constraints needed for quadrangulation. In addition to the holonomy constraints we describe, it is essential to ensure that cone parametric coordinates and seam translations are integer multiples of the parametric quad size. Our distortion minimization procedure does not use this constraint. However, if the desired quadrangulation density is sufficiently fine this constraints does not have a strong effect on distortion. For rectangular texture atlases, multi-chart geometry images [Carr et al. 2006], and T-mesh surfaces [Myles et al. 2010], rounding is not essential or rounding by small distances is sufficient.

For applications like generating coarse subdivision control meshes, very coarse rounding of cones positions and seam translations is clearly desired and additional effort is required to integrate this rounding into the process.

In our current implementation we do not use the efficient greedy mixed-integer solver based on elimination of constraints and reduction to a positive-definite system; because the system we solve has per-vertex constraints, using Gaussian elimination to reduce the number of variables leads to a dense system. At the same time, constructing a similar efficient solver for symmetric indefinite matrices is possible, e.g., by replacing a Cholesky-based direct solver with an LDL^T solver, and replacing CG with a modified CG with constraint projection for iterative updates. Like [Bommes et al. 2009], our metric-based system has the structure of a Laplacian system, and we conjecture that our mixed integer system has similar locality properties to benefit from the use of local iterative updates in the




Figure 13: Comparison with [Bommes et al. 2009] on seamless parametrizations. The method of [Bommes et al. 2009] on the screw on the top row results in a collapsed parametrization due to reasons similar to those shown in Figure 10. The percentage accompanying the histogram indicates the fraction of triangles with distortion less than 0.3. See Table 1 for the number of cones.




Figure 14: Comparison with [Myles and Zorin 2012] without feature-alignment: For $\alpha_{smooth} = 0.001$, our method produces results with similar numbers of cones and similar distortion distributions to that of [Myles and Zorin 2012] when no feature alignment constraints are specified. The percentage accompanying the histogram indicates the fraction of triangles with distortion less than 0.2.




Figure 15: Varying smoothness parameter α_{smooth} : Distortions of seamless parametrizations of the Mannequin and Aircraft models, respectively, with statistics tabulated in Table 1. Increasing α_{smooth} trades distortion for fewer cones. The entries indicated with 100.0* for α_{smooth} are without the “exact-as-possible” constraints (17) to generate a number of cones and histogram similar to that of [Bommes et al. 2009].

early rounding stages.

Finally, we observe that a greedy algorithm based on local criteria for choosing integer variable values, although behaves well in practice does not guarantee that the distortion is minimized and depends on the choice of the spanning forest. While solving the problem exactly is likely to be intractable due to the presence of integer variables, better and more global approximation methods may yield higher-quality solutions.

Acknowledgements

We are grateful to the authors of [Bommes et al. 2009] for distributing their mixed integer solver. This work was partially supported by NSF awards IIS-0905502, OCI-1047932, DMS-0602235.

A Proofs

Proof of Proposition 2. Indeed, as $\xi = 0$ on the boundary, $d\xi(t) = 0$ on the boundary, so ξ has no effect on boundary alignment. As for any closed loop $\int d\xi(t)ds = 0$ by co-closedness of the form, it has no impact on the holonomy constraint. Finally, $\int_{\gamma_i} d\xi(t)ds = \xi(p_j) - \xi(p_i) = 0$, as both of these values are zero. The precise proof of existence and uniqueness requires a more exact specification of regularity assumptions; we simply observe that the 2-form ψ can be identified with a scalar function, and frame rotation constraints around infinitesimal loops yield a Poisson equation for this function with Dirichlet boundary conditions (harmonic form does not contribute as its Laplacian vanishes), hence it is determined uniquely by contractible loop constraints and boundary alignment. Finally, the space of harmonic forms has exactly the right dimension to satisfy loop constraints for homology basis loops and for alignment between boundary components. As ξ does not contribute to satisfying the constraints, minimality implies that it is 0. $\square$

Proof of Proposition 3. If we consider any geodesic in metric g through a point in a direction $\mathbf{v}$, then it satisfies $d\phi(\mathbf{n}) = \kappa$. Recall that for a connection 1-form ω , $\omega(\mathbf{v})ds - \kappa ds$ is the change in angle between the tangent of a curve with curvature κ and tangent $\mathbf{v}$, and the frame vector $\mathbf{e}$. Consider a geodesic (in metric g) in direction $\mathbf{v}$. As the conformal map preserves angles, and frame vectors $\mathbf{e}$ are parallel in metric g at different points, the angle between the frame and this geodesic has to remain constant. We conclude that $\omega(\mathbf{v}) = \kappa$. Comparing to the formula $d\phi(\mathbf{n}) = \kappa$, and using the

fact that $\alpha(\mathbf{v}) = \star\alpha(\mathbf{v}^\perp)$, for any vector $\mathbf{v}$ (Hodge star acts as a 90-degree rotation), we get the result of the lemma. $\square$

Proof of Proposition 5. We start by reducing the problem to the case of meshes, by considering a triangle mesh approximating the surface both in position and normal with accuracy $O(\epsilon)$ and with triangles of size $O(\epsilon)$. Then one can bound the total metric distortion on each triangle by a $O(\epsilon^2)$ function d_T . The area of a triangle is $A_T O(\epsilon^2)$, and the number of triangles is $O(\epsilon^{-2})$, leading to $E^{\text{ARAP}} \sqrt{\sum_T d_T^2 A_T}$, bounded by $O(\epsilon)$.

The mesh can be mapped to the plane isometrically, by unfolding it along a spanning tree of all triangles, so the composition of the map from the surface to the mesh, and from the mesh to the plane has distortion at most $O(\epsilon)$.

Next, we subdivide the mesh so that all edges have length at most δ ; we will determine the relationship between δ and ϵ later.

Consider two sides of a cut in an isometric parametrization. We are going to make the parametrization seamless along the cut adding $k\pi/2$ cones ($k = 3$ and $k = 5$) along the boundary.




Figure 16: Constructing a cone chain. Left: geometric construction of a pair of cones. Right: a cone chain

Suppose on two sides of a seam curve in the parametric domain we have edges $\mathbf{v}_1$ and $\mathbf{v}_2$ that correspond to the same edge e on the surface. We are going to split each triangle along the edge into several (typically two). These edges have to be of equal length. We are going to put one cone at the beginning of the edge, and a second cone somewhere along the edge, so that the edge is mapped to two two-segment lines in the parametric domain. We denote the two segments corresponding to $\mathbf{v}_i$ by $\mathbf{w}_{i1}$ and $\mathbf{w}_{i2}$. For the parametrization to be seamless, the vectors corresponding to $\mathbf{v}_2$, should be either parallel to vectors corresponding to $\mathbf{v}_1$, or differ




Figure 17: Cross-field representations. The primary direction of each cross (red) is represented by its angle from the reference vector $\mathbf{b}_{ijk}$. (a) [Ray et al. 2008; Bommes et al. 2009]. (b) Our choice of ω induces a representation using non-conforming finite elements.

by a 90-degree rotation. Clearly, they cannot be parallel to each other, because then the angle between $\mathbf{v}_1$ and $\mathbf{v}_2$ has to be zero or 90 degrees. We choose $\mathbf{w}_{21} = \mathbf{w}_{11}^\perp$, and $\mathbf{w}_{22} = \mathbf{w}_{12}$. Representing vectors by complex numbers, we can express the relations between different vectors as

$$\mathbf{v}_1 = \mathbf{w}_{11} + \mathbf{w}_{12}, \quad \mathbf{v}_2 = i\mathbf{w}_{11} + \mathbf{w}_{12}. \quad (20)$$

This system always has a solution that can be constructed geometrically (Figure 16).

Consider the strip S of triangles along the seam. We observe that the angle between $\mathbf{w}_{11}$ and $\mathbf{w}_{12}$ is always at least $\pi/4$, so the triangle T' formed by $\mathbf{v}_i$, $\mathbf{w}_{11}$ and $\mathbf{w}_{2i}$ is necessarily contained in a circle of radius $\sqrt{2}d$, for any directions of $\mathbf{v}_1$ and $\mathbf{v}_2$. Consider the strip of triangles along the seam, and cut it into three parts S_1 , S_2 and S_3 , connecting points splitting all non-seam edges incident at the seam vertices into three. Let the boundary between S_2 and S_3 be L . By subdividing seam edges if necessary we can ensure that all triangles T' are contained in S_1 , closest to the seam. Consider the strip bounded by L and the modified path along the seam, consisting of vectors $\mathbf{w}_{ij}$ for all original segments.

Changing the seam lines to follow $\mathbf{w}_{i1}$ and $\mathbf{w}_{i2}$ instead of $\mathbf{v}_i$, can be achieved by a local modification, entirely within $S_1 \cup S_2$. As the distortion in our norm is always bounded, the total distortion is at most $O(\sqrt{|L|\delta})$, the area of the strip.

Note that the length of the seam cutting the mesh to a tree spanning all faces is proportional to the number faces in the mesh times the bound on the face size, i.e. $O(\epsilon^{-1})$, giving us an estimate of $O(\epsilon^{-1}\sqrt{\delta})$ for the total distortion. We conclude that δ needs to decrease faster than $O(\epsilon)$ to observe a decrease in distortion with refinement. To balance distortion due to meshing and distortion due to cones, we can choose $\delta \sim \epsilon^2$, yielding distortion $O(\epsilon)$ for both. $\square$

B Parametrization from the connection

Once the cones, homology loop holonomies, and alignment angles are determined, one may proceed as in [Myles and Zorin 2012] and

1. cut the surface M along edges to a topological disk M_c so that the seam passes through all cones [Eppstein 2003];
2. compute a cross-field satisfying the holonomy conditions either via [Ray et al. 2008; Bommes et al. 2009; Crane et al. 2010] or from the connection ω directly; and
3. compute a parametrization f minimizing (4), where the initial target gradient R is computed from the field.

The following sections detail steps 2 and 3.

B.1 Cross-fields from connections

Here, we show how to compute the field directly from the connection ω . While [Ray et al. 2008; Bommes et al. 2009] represent the cross-field with angles defined per triangle, the ω we compute relates cross-field degrees of freedom defined on edges interpolated with non-conforming linear elements [Pinkall and Polthier 1993;

Polthier 2000]. Figure 17 contrasts these representations. In both cases, cross-fields are represented in each triangle ijk as angles with respect to a reference vector $\mathbf{b}_{ijk}$, with $\kappa_{ij} = -\kappa_{ji} \in (-\pi, \pi]$ denoting the angle of rotation between adjacent reference vectors $\mathbf{b}_{ijk}$ and $\mathbf{b}_{jil}$. An integer matching variable $m_{ij} = -m_{ji}$ accounts for the $\pi/2$ ambiguity between adjacent triangles so that $\theta_{ijk} + \kappa_{ij} + \frac{\pi}{2}m_{ij}$ is the representation of the cross-field from facet ijk parallel-transported to the representation in triangle jil . Thus, in our non-conforming representation,

$$\beta_{ji} = \beta_{ij} + \kappa_{ij} + m_{ij} \frac{\pi}{2}. \quad (21)$$

Since the connection computes the total rotation of the cross-field along the midline,

$$\omega_j = \beta_{jk} - \beta_{ij}. \quad (22)$$

Like [Bommes et al. 2009], we begin by assigning zero matchings m_{ij} on the edges of a dual spanning tree of M_c . Then, with β fixed at one feature edge, (21) and (22) define the values of β_{ij} on the rest of the surface by propagation along the dual spanning tree. The matchings m_{ij} on the rest of M are then computed via (22).

This representation can be converted to a per-facet angle $\theta_{ijk} = \frac{1}{3}(\beta_{ij} + \beta_{jk} + \beta_{ki})$ to compute the primary vector $\mathbf{e}_{ijk} = \frac{\mathbf{b}_{ijk}}{|\mathbf{b}_{ijk}|} \cos \theta_{ijk} + \frac{\mathbf{b}_{ijk}^\perp}{|\mathbf{b}_{ijk}|} \sin \theta_{ijk}$ of the cross-field, from which we initialize the target Jacobian matrix

$$R_{ijk} = [\mathbf{e}_{ijk} \quad \mathbf{e}_{ijk}^\perp]^\top$$

on each facet for the following section.

B.2 Computing an ARAP parametrization

As in [Bommes et al. 2009; Myles and Zorin 2012], the parametrization f is discretized using linear finite elements with (u, v) degrees of freedom on the vertices of M_c . Recall from Section 3 and Figure 2 that each point $p_1 = p_2$ on the seam is assigned two (or more) parametric values related by $f(p_2) = rf(p_1) + t$. On each seam edge ij , the seam rotation $r_{ij} = \text{Rot}(m_{ij}\pi/2)$ is defined by the matching m_{ij} , and the parametrization f is computed from the prescribed gradients R by minimizing

$$\frac{1}{2} \int_M \|\nabla f - R\|^2 \quad (23)$$

under the seam constraints above, which are linear in f . For quadrangulations, grid lines are made continuous across seams (cf. Figure 2) by rounding all t and cone (u, v) values to integers using the mixed integer solver of [Bommes et al. 2009]. Scaling R globally allows for control over the grid resolution.

[Myles and Zorin 2012] further reduced the isometric distortion by minimizing the ARAP energy (4):

$$E^{\text{ARAP}} = \frac{1}{2} \int_M \min_{R \in \text{SO}(2)} \|\nabla f - R\|^2$$

using the local/global alternating optimization of [Liu et al. 2008] which iterates the following two steps until convergence.

1. **Local step:** Fix f . In each triangle, use a local frame to write ∇f as a 2×2 matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$, and compute the closest rotation

$$R = \frac{B}{\sqrt{\det(B)}}, \quad \text{where} \quad B = \frac{1}{2} \begin{bmatrix} a+d & b-c \\ c-b & a+d \end{bmatrix},$$

taking care to leave the R unchanged on triangles adjacent to features.

2. **Global step:** Fix R and solve (23) for f as before.

References

- BEN-CHEN, M., GOTSCHAN, C., AND BUNIN, G. 2008. Conformal flattening by curvature prescription and metric scaling. *Computer Graphics Forum* 27, 2, 449–458.
- BOMMES, D., ZIMMER, H., AND KOBELT, L. 2009. Mixed-integer quadrangulation. *ACM Trans. Graph.* 28, 3, 77.
- BOMMES, D., LVY, B., PIETRONI, N., PUPPO, E., SILVA, C., TARINI, M., AND ZORIN, D. 2012. Quad Meshing. Eurographics Association, Cagliari, Sardinia, Italy, M.-P. Cani and F. Ganovelli, Eds., 159–182.
- BUNIN, G. 2008. Towards unstructured mesh generation using the inverse poisson problem. *arXiv preprint arXiv:0802.2399*.
- BUNIN, G. 2008. A continuum theory for unstructured mesh generation in two dimensions. *CAGD* 25, 14–40.
- CAPPELL, S., DETURCK, D., GLUCK, H., AND MILLER, E. 2006. Cohomology of harmonic forms on riemannian manifolds with boundary. In *Forum Mathematicum*, vol. 18, 923–932.
- CARR, N., HOBEROCK, J., CRANE, K., AND HART, J. 2006. Rectangular multi-chart geometry images. In *Symposium on Geometry Processing*, Eurographics Association, 190.
- CRANE, K., DESBRUN, M., AND SCHRÖDER, P. 2010. Trivial connections on discrete surfaces. *Computer Graphics Forum* 29, 5 (July), 1525–1533.
- DANIELS, J., SILVA, C. T., AND COHEN, E. 2009. Localized quadrilateral coarsening. *Computer Graphics Forum* 28, 5, 1437–1444.
- DANIELS II, J., SILVA, C. T., AND COHEN, E. 2009. Semiregular quadrilaterally remeshing from simplified base domains. *Computer Graphics Forum* 28, 5 (July), 1427–1435.
- DE GOES, F., AND CRANE, K., 2010. Trivial connections on discrete surfaces revisited: A simplified algorithm for simply-connected surfaces.
- DINDOŠ, M. 2008. *Hardy Spaces and Potential Theory on C1 in Riemannian Manifolds*. American Mathematical Soc.
- DONG, S., BREMER, P., GARLAND, M., PASCUCCHI, V., AND HART, J. 2006. Spectral surface quadrangulation. *ACM Trans. Graph.* 25, 3, 1057–1066.
- ECK, M., DEROSE, T., DUCHAMP, T., HOPPE, H., LOUNSBERY, M., AND STUETZLE, W. 1995. Multiresolution analysis of arbitrary meshes. *SIGGRAPH* 1995, 173–182.
- EPPSTEIN, D. 2003. Dynamic generators of topologically embedded graphs. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, SODA '03, 599–608.
- GU, X., AND YAU, S.-T. 2003. Global conformal surface parameterization. In *Proc. 2003 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, SGP '03, 127–137.
- JIN, M., WANG, Y., YAU, S., AND GU, X. 2004. Optimal global conformal surface parameterization. In *Proc. IEEE Visualization'04*, 267–274.
- JIN, M., KIM, J., LUO, F., AND GU, X. 2008. Discrete surface ricci flow. *IEEE Trans. Visualization and Computer Graphics* 14, 1030–1043.
- KÄLBERER, F., NIESER, M., AND POLTHIER, K. 2007. Quad-Cover: Surface Parameterization using Branched Coverings. *Computer Graphics Forum* 26, 3, 375–384.
- KHODAKOVSKY, A., LITKE, N., AND SCHRÖDER, P. 2003. Globally smooth parameterizations with low distortion. *ACM Trans. Graph.* 22, 3, 350–357.
- LAI, Y., JIN, M., XIE, X., HE, Y., PALACIOS, J., ZHANG, E., HU, S., AND GU, X. 2009. Metric-driven rosy field design and remeshing. *IEEE Trans. Visualization and Computer Graphics*, 95–108.
- LEE, A., SWELDENS, W., SCHRÖDER, P., COWSAR, L., AND DOBKIN, D. 1998. MAPS: multiresolution adaptive parameterization of surfaces. In *SIGGRAPH* 1998, 95–104.
- LIU, L., ZHANG, L., XU, Y., GOTSCHAN, C., AND GORTLER, S. J. 2008. A Local/Global approach to mesh parameterization. *Computer Graphics Forum* 27, 5 (July), 1495–1504.
- MARINOV, M., AND KOBELT, L. 2005. Automatic generation of structure preserving multiresolution models. *Computer Graphics Forum* 24, 3 (Sept.), 479–486.
- MYLES, A., AND ZORIN, D. 2012. Global parametrization by incremental flattening. *ACM Transactions on Graphics (TOG)* 31, 4, 109.
- MYLES, A., PIETRONI, N., KOVACS, D., AND ZORIN, D. 2010. Feature-aligned T-meshes. *ACM Trans. Graph.* 29, 4, 1–11.
- O'NEILL, B. 2006. *Elementary Differential Geometry, Revised 2nd Edition*. Elementary Differential Geometry Series. Elsevier Science.
- PALACIOS, J., AND ZHANG, E. 2007. Rotational symmetry field design on surfaces. *ACM Trans. Graph.* 26, 3 (July).
- PIETRONI, N., TARINI, M., AND CIGNONI, P. 2009. Almost isometric mesh parameterization through abstract domains. *IEEE Trans. Visualization and Computer Graphics* 99, RapidPosts.
- PINKALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics* 2, 1, 15–36.
- POLTHIER, K. 2000. Conjugate harmonic maps and minimal surfaces. Preprint No. 446, TU-Berlin, SFB 288, 2000.
- RAY, N., LI, W., LÉVY, B., SHEFFER, A., AND ALLIEZ, P. 2006. Periodic global parameterization. *ACM Trans. Graph.* 25, 4, 1460–1485.
- RAY, N., VALLET, B., LI, W., AND LÉVY, B. 2008. N-Symmetry direction field design. *ACM Trans. Graph.* 27, 2.
- RAY, N., VALLET, B., ALONSO, L., AND LEVY, B. 2009. Geometry-aware direction field processing. *ACM Trans. Graph.* 29, 1, 1–11.
- SPRINGBORN, B., SCHRÖDER, P., AND PINKALL, U. 2008. Conformal equivalence of triangle meshes. *ACM Trans. Graph.* 27 (August), 77:1–77:11.
- TARINI, M., PIETRONI, N., CIGNONI, P., PANOZZO, D., AND PUPPO, E. 2010. Practical quad mesh simplification. *Computer Graphics Forum* 29, 2.
- TONG, Y., ALLIEZ, P., COHEN-STEINER, D., AND DESBRUN, M. 2006. Designing quadrangulations with discrete harmonic forms. *Symposium on Geometry Processing*, 201–210.