

# Which Cross Fields can be Quadrangulated? Global Parameterization from Prescribed Holonomy Signatures

HANXIAO SHEN, LEYI ZHU, RYAN CAPOUELLEZ, DANIELE PANZZO, New York University, USA

MARCEL CAMPEN, Osnabrück University, Germany

DENIS ZORIN, New York University, USA




Fig. 1. Our method is able to construct seamless parametrizations with prescribed *holonomy signature*, i.e. it offers full control over map topology. This topological information can be given as values (holonomy numbers) on a system of certain loops on the surface (a) or, e.g., be derived from a cross-field that the parametrization is supposed to align to. A state-of-the-art method [Campen et al. 2019] can reliably generate seamless parametrizations (b), but it ignores this information at the global level, i.e. it does not offer full control over map topology. Our method adjusts the loop system with associated numbers into a topologically equivalent state (c) of very particular type. Using a cut graph constructed from this loop system, we are then able to reliably generate a seamless parametrization (d) that perfectly matches the prescribed holonomy signature, and for instance allows for lower distortion and better cross-field alignment.

We describe a method for the generation of seamless surface parametrizations with guaranteed local injectivity and full control over holonomy. Previous methods guarantee only one of the two. Local injectivity is required to enable these parametrizations' use in applications such as surface quadrangulation and spline construction. Holonomy control is crucial to enable guidance or prescription of the parametrization's isocurves based on directional information, in particular from cross-fields or feature curves, and more generally to constrain the parametrization topologically. To this end we investigate the relation between cross-field topology and seamless parametrization topology. Leveraging previous results on locally injective parametrization and combining them with insights on this relation in terms of holonomy, we propose an algorithm that meets these requirements. A key component relies on the insight that arbitrary surface cut graphs, as required for global parametrization, can be homeomorphically modified to assume almost any set of turning numbers with respect to a given target cross-field.

CCS Concepts: • Computing methodologies → Computer graphics; Mesh models.

Additional Key Words and Phrases: seamless parametrization, conformal map, quad mesh, cross field, turning number, holonomy

---

Author's addresses: Hanxiao Shen, Leyi Zhu, Ryan Capouillez, Daniele Panzzo, Denis Zorin, Courant Institute of Mathematical Sciences, New York University, New York, NY, USA. Marcel Campen, Institute for Computer Science, Osnabrück University, Germany.

---

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3528223.3530187>.

## ACM Reference Format:

Hanxiao Shen, Leyi Zhu, Ryan Capouillez, Daniele Panzzo, Marcel Campen, and Denis Zorin. 2022. Which Cross Fields can be Quadrangulated? Global Parameterization from Prescribed Holonomy Signatures. *ACM Trans. Graph.* 41, 4, Article 59 (July 2022), 12 pages. <https://doi.org/10.1145/3528223.3530187>

## 1 INTRODUCTION

Seamless surface parametrization is one of the most common approaches to constructing seamless texture atlases, conforming surface quadrangulations, and high-order (spline or subdivision) approximations to surface data. A chart-based parametrization is called *seamless* if it satisfies certain conditions on its transitions between charts or across cuts.

In particular, a seamless parametrization of a discrete surface defines a metric, i.e., an edge length assignment on a mesh, that is intrinsically flat almost everywhere, i.e. angles around vertices sum to  $2\pi$ , except at a (often small) set of *cone vertices* with an angle deficit (or excess) of some multiple of  $\frac{\pi}{2}$ . More generally, the *holonomy angle* for any closed loop on the surface is a multiple of  $\frac{\pi}{2}$ . The holonomy angle is the angle between the first and last edge when laying out a closed chain of mesh triangles in the plane according to the metric (Fig. 2, cf. [Bright et al. 2017; Crane et al. 2010]). Informally, this holonomy condition on a parametrization's metric ensures that parametric lines continue seamlessly across cuts, although, e.g., a  $u$ -parametric line may become a  $v$ -parametric line.

While the set of closed triangle chains on a discrete surface is infinite, all their holonomy angles are actually defined by a set of angles on a finite basis, the *holonomy signature*, which we define

more precisely below. In essence, loops around individual vertices capture all *local* aspects of holonomy, while (in case of non-trivial topology,  $\text{genus} > 0$ ) a system of non-contractible loops captures the additional *global* aspects.

*Holonomy Control.* To clarify the importance of parametrization topology defined by the holonomy signature, consider quad meshes or quad layouts obtained from (constrained classes of) seamless parametrizations by tracing a grid of parametric lines on the surface. The cones become the extraordinary vertices, where  $n \neq 4$  quads meet. The holonomy angles determine how many quads meet at such extraordinary vertices, and more generally, how many turns the edges of the quads make along any closed curve on the surface, e.g., a feature line. As a consequence, controlling the parametrization’s topology in the form of its holonomy angles is critical for obtaining a high-quality parametrization with intended behavior.

In many approaches to seamless parametrization, the target topology is provided as input, e.g., it is derived from a given cross-field or partially or completely specified by the user. At the same time, as we discuss in detail in Section 2, no existing general method guarantees that the target topology is fully respected, although significant progress was made towards this goal.

*Existence.* Moreover, to the best of our knowledge, the answer to the following question is not known:

*For which holonomy signatures, seamless parametrizations with corresponding topology exist?*

Partially, this question was answered in [Jućović and Trenkler 1973], and more specifically in [Campen et al. 2019], where angles at cones, but not complete signatures (including global aspects), were considered. In this paper, we resolve the question of existence for a broad class of signatures, subject to only a mild condition.

Remarkably, it turns out that for surfaces of genus  $\neq 1$ , there is a seamless parametrization for *any* holonomy signature (e.g. implied by a cross-field) under this condition. For genus 1, we show that in this class the one known example of holonomy signatures for which there is no seamless parametrization (signatures with exactly two cones, with angles  $3\pi/2$  and  $5\pi/2$ ) is the only one.

For the condition to be satisfied, it is already sufficient (but not necessary) to have one cone with angle deficit  $+\frac{\pi}{2}$  or  $-\frac{\pi}{2}$  in the signature (corresponding to at least one valence 3 or valence 5 extraordinary quad vertex). This is essentially always satisfied for holonomy signatures implied by cross-fields optimized for smoothness or curvature alignment [Vaxman et al. 2016].

*Contribution.* We describe an algorithm for the construction of seamless parametrizations with full control over holonomy. It extends the construction of [Campen et al. 2019] (referred to as Seamless Padding (SP) in the following) which provides control only over local holonomy aspects (i.e. cone angles). Our contribution includes:

- An existence result for seamless parametrizations with given holonomy signature, indicating a remarkably small topological gap between cross-fields and parametrizations;
- An algorithm for, given a holonomy signature, constructing an alternative system of loops on which the equivalent holonomy signature has arbitrary desired angles;

- A variant of the SP method that, based on the above, builds a valid seamless parametrization with prescribed holonomy.

We note that the topology of cross-fields—which are often used to guide the computation of seamless parametrizations—can be controlled very flexibly and precisely using existing discrete construction algorithms. In fact, one can easily construct a cross-field with any given turning number signature (the field analogue of a holonomy signature, cf. Section 3.1) by solving a linear system of equations [Crane et al. 2010]. The ability to near-universally match this signature, provided by our method, means that this possibility of precise topology control extends to parametrizations.

## 2 RELATED WORK

*Seamless Parametrization.* Seamless surface parametrization with prescribed singularities (locations and indices) is a problem that has received significant attention recently. Results include [Bommes et al. 2009; Bright et al. 2017; Campen et al. 2015, 2019; Chien et al. 2016; Ebke et al. 2016; Fang et al. 2018; Fu et al. 2015; Hefetz et al. 2019; Kälberer et al. 2007; Kovacs et al. 2011; Lyon et al. 2019; Myles et al. 2014; Myles and Zorin 2013; Tong et al. 2006; Zhou et al. 2018, 2020]. Prominent use cases are surface quadrangulation [Bommes et al. 2013b; Campen 2017] and spline conversion [Campen and Zorin 2017; Marinov et al. 2019].

*Cross-Field Guidance.* Most often such parametrizations are generated and optimized guided by a cross-field or frame field on the surface [Vaxman et al. 2016]. Seminal works on cross-field guided parametrization are [Kälberer et al. 2007; Knupp 1995]. Important ideas for cross-field generation are presented by [Bommes et al. 2009; Crane et al. 2010; Li et al. 2006; Ray et al. 2009, 2008]; many of these offer control over the fields’ turning numbers.

*Local Injectivity.* For common use cases, parametrizations are valid only if they are locally injective, i.e. free of fold-overs. Local injectivity constraints required to ensure a valid parametrization are, due to their challenging non-convex nature, not rarely omitted [Bommes et al. 2009; Ebke et al. 2016; Kälberer et al. 2007; Kovacs et al. 2011; Myles and Zorin 2013; Zhou et al. 2018] or convexified in a conservative manner [Bommes et al. 2013a; Bright et al. 2017; Campen et al. 2015; Hefetz et al. 2019; Lipman 2012].

*Guarantees.* In special cases (restricted genus, restricted cone configurations) convex formulations can be used to reliably yield locally injective seamless parametrizations [Aigerman and Lipman 2015; Gortler et al. 2006; Gu and Yau 2003]. Alternatively, additional user input like a surface partition may be exploited to ensure validity [Tong et al. 2006], or more general, non-piecewise-linear forms of parametrization may be employed [Aigerman and Lipman 2016].

Recently, first methods have emerged that provide validity guarantees while supporting arbitrary genus and general cone configurations [Campen et al. 2019; Zhou et al. 2020]. In this sense, they offer control over *local* holonomy aspects. No *global* control over holonomy is provided, though. Therefore, when for instance aiming to generate a cross-field guided parametrization, while locally cones are reproduced, there may be global topological mismatches between the given cross-field and the constructed parametrization, for instance precluding proper alignment, as in Fig. 1 (b).

*Existence.* [Jućović and Trenkler 1973] consider the question of existence of quadrilateral meshes with prescribed irregular vertex valences, [Campen et al. 2019] the very closely related question of existence of seamless parametrizations with prescribed cones. Both consider only local holonomy (irregular vertex valence, cone angles), not global holonomy.

*Distortion Optimization.* The task of parametrization optimization (with respect to varying measures of distortion) while maintaining properties such as local injectivity or seamlessness, is addressed in a number of recent works [Hormann and Greiner 2000; Kovalsky et al. 2016; Mandad and Campen 2020; Rabinovich et al. 2017; Schüller et al. 2013; Shtengel et al. 2017; Zhu et al. 2018]. They are useful as a post-process in the context of our method as we initially focus mainly on validity, holonomy, and seamlessness.

*Cone Choice.* The choice of cones (and more generally guiding cross-fields, holonomy signatures) is an application dependent matter. Various approaches have been proposed for the selection of a cone configuration, for instance curvature-based (e.g. via cross-fields [Vaxman et al. 2016]), distortion-based [Ben-Chen et al. 2008; Kharevych et al. 2006; Soliman et al. 2018], or interactive [Campen and Kobbel 2014; Ebke et al. 2016]. The problem of positioning cones such that conformal maps with these cones become seamless is addressed by [Chen et al. 2019, 2020].

*Holonomy.* Some of the above methods for parametrization construction (such as [Bommes et al. 2009; Bright et al. 2017]) offer full control over the resulting parametrizations’ holonomy, but do not guarantee local injectivity. Those that guarantee local injectivity in a general setting (e.g. [Campen et al. 2019; Myles et al. 2014; Zhou et al. 2020]), in turn, do not offer full control over holonomy. The method of [Campen and Zorin 2017] offers full holonomy control, albeit only for the broader class of seamless *similarity* parametrizations.

### 3 HOLOMONY SIGNATURE

We consider a closed orientable manifold mesh  $M$  of genus  $g$  and a cut graph  $G$  on  $M$  that cuts  $M$  to one or more topological disks. We let  $M^c$  denote the resulting cut mesh, which has a canonical map  $\pi : M^c \rightarrow M$  that is the identity on the interior and maps exactly two boundary edges in  $M^c$  to each edge in  $G \subset M$ . We call two edges  $e, e'$  in the boundary of  $M^c$  *mates* if  $\pi(e) = \pi(e')$ . We also define a *loop* as an oriented closed walk of facets (or dual vertices) of  $M$  and a *simple loop* as an oriented cycle of facets (or dual vertices).

**DEFINITION 1 (SEAMLESS PARAMETRIZATION).** *A discrete seamless parametrization, as in [Myles and Zorin 2013], is a continuous piecewise linear, locally injective map  $F : M^c \rightarrow \mathbb{R}^2$  such that for any boundary edge  $e$  with mate  $e'$ , there is a rigid transformation  $\sigma_e(x) = R_e x + t_e$ , where  $R_e$  is a rotation by an integer multiple of  $\frac{\pi}{2}$ , that maps  $F(e)$  to  $F(e')$ , i.e.  $\sigma_e(F(e)) = F(e')$ .*

A seamless parametrization naturally induces a discrete metric  $E \rightarrow \mathbb{R}^{>0}$  on  $M^c$  by letting the length of an edge  $e$  of  $M^c$  be the length of  $F(e) \subset \mathbb{R}^2$ . Since mated edges  $e, e'$  in the boundary of  $M^c$  are related by a rigid transformation,  $F(e)$  and  $F(e')$  have the same length, so this metric extends to a well-defined metric on




Fig. 2. The holonomy angle  $\kappa_y^F$  (Def. 2) of a dual loop (cyclic triangle strip) under a metric  $F$  is the sum of signed inner angles (yellow and orange). Up to multiples of  $2\pi$  (if the loop makes multiple turns) this corresponds to the angle between first and last edge when laying out the strip in the plane.

$M$ . Moreover, the metric on  $M$  is flat except at isolated vertices  $C = \{v_1, \dots, v_m\}$  in  $G$ , i.e. in the boundary of  $M^c$ , called *cones*.

**DEFINITION 2 (HOLONOMY ANGLE).** *For a loop  $y$ , the holonomy angle (or discrete geodesic curvature [Crane et al. 2010]) of  $y$  under a seamless parametrization  $F$  is*

$$\kappa_y^F = \sum_{f^* \in y} \alpha_y(f^*),$$

where  $f^*$  is the vertex dual to facet  $f$ ,  $\alpha_y(f^*)$  is the signed angle of  $F(f)$  at the vertex of  $f$  incident to the preceding and succeeding facets in the loop. The sign is positive (negative) for vertices on left (right) hand side of the loop. See Fig. 2 for an illustration.

Let  $v^*$  denote the dual facet of vertex  $v$ , and  $\partial v^*$  the cycle of this dual facet’s dual vertices. In other words,  $\partial v^*$  is the loop around the single vertex  $v$ .

**DEFINITION 3 (HOLONOMY NUMBER).** *For a loop  $y$ , we define the holonomy number as  $k_y^F = \kappa_y^F / 2\pi$ . In the case of a vertex-loop,  $y = \partial v^*$ , we additionally define the index of the vertex as  $I_v^F = 1 - k_{\partial v^*}^F$ .*

Since the images of edges in the boundary of  $M^c$  under  $F$  are related by rotation angles that are integer multiples of  $\frac{\pi}{2}$ , the holonomy numbers are always integer multiples of  $\frac{1}{4}$ .

**DEFINITION 4 (HOLONOMY SIGNATURE).** *We call the holonomy numbers for a choice of homology basis loops  $\gamma_1, \dots, \gamma_{2g+m}$  of  $M \setminus C$  the holonomy signature of  $F$ . A natural choice of basis is a homology basis  $\gamma_1, \dots, \gamma_{2g}$  of  $M$  together with cone vertex loops  $\partial v_1^*, \dots, \partial v_m^*$ . The loops are referred to as signature loops.*

A homology basis of  $M$  can, for instance, be chosen as a so-called system of loops [Erickson and Whittlesey 2005].

Importantly, such a finite holonomy signature completely captures the holonomy number of *any* loop: First, since the metric is flat except at cones, any two loops homotopic in  $M \setminus C$  (i.e., loops that can be continuously deformed into each other within  $M$  without crossing a cone vertex) have the same holonomy number (cf. Prop. 1 in [Myles and Zorin 2012]). Second, given an arbitrary loop  $y$  and a homology basis of  $M \setminus C$ , there is (by the nature of a homology basis [Hatcher 2002]) a surface  $\tilde{M}$  so that its boundary is composed of  $y$  and some combination of the basis loops. The Gauss-Bonnet theorem then gives a formula for the holonomy number of this boundary in terms of the Euler characteristic of  $\tilde{M}$ . Thus, the holonomy number of  $y$  is determined by the holonomy numbers of the




Fig. 3. Algorithm overview: (a) Example input signature loops (yellow and green) and cones (red and blue). (b) Loops of an equivalent signature obtained by strategically modifying this input; notice that the yellow loop takes a different path between the cones. (c) Conformal parametrization respecting the prescribed cones and aligned with the cut graph that is formed by the loops; due to this alignment, it has a specific holonomy pattern along the loops. (d) The map is modified by parametric padding to make it seamless while preserving its holonomy properties. (e) Finally, the map can be continuously optimized for low distortion and possibly cross field alignment, naturally within its topological class.

loops of a homology basis (cf. Prop. 2 in [Myles and Zorin 2012]). Fig. 3 (a) shows an example of signature loops: the green and yellow loops form a homology basis of  $M$ , while the small cone vertex loops are visualized as red and blue dots at the respective vertices.

Note that the holonomy signature is not unique, neither its loops nor its numbers (Fig. 4). For a fixed seamless parametrization  $F$ , a different choice of signature loops will lead to different associated holonomy numbers—even though the same parametrization topology is represented. Such holonomy signatures are called *equivalent*.

### 3.1 Relation to Cross-Fields

Seamless parametrizations are often employed in conjunction with cross-fields, most importantly when parametrizations are built and optimized for directional alignment with such a field. In such cases it is important for field topology and parametrization topology to match. In this context, there is a close connection between the holonomy of a seamless parametrization and the turning numbers of a cross-field.

For a smooth surface  $S$ , a cross-field  $\vec{d}$  is a differentiable mapping of four tangent vectors to each point  $p \in S$  (except at isolated singularities) that are invariant to rotation by  $\pi/2$  around the normal  $\hat{n}_p$  to  $S$  at  $p$ . Given such a field and a loop  $\gamma$  on the surface  $S$ , the field will make some number of rotations along this loop, and due to the rotational symmetry of the field these turning numbers  $T_\gamma$  can be integer multiples of  $\frac{1}{4}$ .




Fig. 4. Example of two equivalent holonomy signatures. Red and blue cones have index  $-\frac{1}{4}$  and  $+\frac{1}{4}$ , respectively; the holonomy numbers of the green and yellow loops are indicated. Note that from left to right, the loops are essentially deformed across a cone (the leftmost red cone), and this affects the loops' associated holonomy numbers accordingly.

Moreover, as shown in [Ray et al. 2008], there is a discrete analogue of cross-fields and turning numbers for triangle meshes, and turning numbers along loops satisfy a theorem analogous to the Poincaré-Hopf theorem for vector fields that states  $T_{\partial S} = -\chi(S)$ . This implies that the holonomy number of the boundary of a flat surface, which does not contain any cone, is the same as the turning number of a singularity-free cross-field along that boundary. Hence, if turning numbers of a cross field agree with holonomy numbers on a set of signature loops then they will agree for *any* loop on the surface. Consequently, by taking as our desired holonomy numbers the turning numbers of a given cross-field on a homology basis of  $M \setminus C$ , where  $C$  is the set of the cross-field's singularities, the seamless parametrization is fully constrained to topologically match the input cross-field—in terms of local (cone indices) as well as global behavior.

## 4 APPROACH OVERVIEW

Given a holonomy signature (or a cross-field implying a holonomy signature, cf. Section 3.1), consisting of loops associated with a holonomy number each, on a surface  $M$ , our goal is to construct a valid seamless parametrization  $F$  for  $M$  that respects this signature. In Section 5 we discuss the question for which signatures this is actually feasible.

**Key Idea.** We show in Section 5 that, given a holonomy signature, we can find an equivalent signature (by exchanging or modifying the signature loops) such that the associated holonomy numbers assume almost any desired values. Essentially, we are exploiting the above mentioned non-uniqueness of the signature (cf. Fig. 4). We make use of this algorithmically in Section 6 in the following way: The SP method [Campen et al. 2019] enables constructing a seamless parametrization that has prescribed local holonomy (i.e. cones), but it lacks the ability to prescribe holonomy globally. However, its result has not a random but a fixed holonomy pattern along the cut graph that is used in the construction. We therefore modify the  $2g$  global signature loops such that their union forms a cut graph and such that their corresponding holonomy numbers in an equivalent signature match exactly the fixed pattern that SP will produce. Fig. 3 illustrates the main steps of our construction process.




Fig. 5. A hole-chain cut graph  $G$ , as used in [Campen et al. 2019]. As an example, the contained loop that is highlighted in red, because it makes two left turns (in ccw sense), will have holonomy number  $\frac{2}{4}$  in the parametrization constructed by that method.

**Cut Graph.** The seamless parametrization construction by SP relies on using cut graphs with certain structural restrictions, so-called hole-chains (or modifications thereof). Fig. 5 shows an example, details follow in Section 6.1. Let  $G$  be such a cut graph. Let  $H$  be a system of loops of  $M$ , i.e.,  $H$  is a homology basis and cuts  $M$  into a topological disk. In particular, as  $G$  cuts  $M$  into a topological disk,  $H$  can be chosen such that the (non-disjoint) union of its loops equals  $G$  (Section 6.2).

**Fixed-Holonomy Parametrization.** By construction, SP will yield certain predetermined holonomy numbers along  $G$ , thus on these loops  $H$ , regardless of the loops' geometry. More concretely, each branch point of the hole-chain  $G$  has four incident cut segments, conceptually forming a cross. When following a loop through  $G$ , at such a branch point it may therefore take a left-turn, a right-turn, or continue straight. The generated parametrization's holonomy number along any loop in  $G$  is simply the (signed) difference between its number of right-turns and its number of left-turns (times  $\frac{1}{4}$ ). This is due to the SP-parametrization being aligned to all segments of  $G$ , i.e., they are geodesic under this metric, and the corners between segments at branch points form right angles under this metric.

**Cut Graph Rerouting.** Given target holonomy numbers on the loops of  $H$  (e.g., derived from an input cross-field), by Prop. 2 we can modify the loops of  $H$ , yielding  $H'$ , such that their target holonomy numbers equal their left-right-turn balance. Conceptually, this rerouting of loops is described in Section 5; in Section 6.3 we describe algorithms that practically implement it. As this rerouting does not alter the loops' intersections, i.e., it preserves the branch points of  $G$  and the loops' left-right-turns, the union of loops from  $H'$  still forms a hole-chain, which we can then use as prescribed cut graph for the parametrization construction.

**Holonomy Soft-Guidance.** In order to yield parametrizations that are not just topologically correct but also (already initially, before distortion optimization) of reasonable geometric quality, we aim to reduce the need for cut graph rerouting as much as possible. To this end we construct the individual paths that the initial hole-chain  $G$  is made of in a holonomy-guided (e.g. cross-field guided) manner (Sections 6.1.1 and 6.1.2). This promotes hole-chains that largely have the desired holonomy properties right away.




Fig. 6. Illustration for Prop. 1 concerning quasi-additivity of holonomy numbers on loops.

#### 4.1 Algorithmic Outline

Our method's overall algorithmic pipeline can be outlined as follows:

- (1) Construct Cut Graph
  - Initial field-guided hole-chain  $G$  (Section 6.1)
  - Extract loop basis  $H$  of  $G$  (Section 6.2)
  - Reroute  $H \rightarrow H'$ , yield  $G'$  (Section 6.3)
- (2) Construct Seamless Parametrization
  - Construct  $G'$ -aligned mapping  $f : M^c \rightarrow \Omega$  (Section 7.1)
  - Pad  $f$  to yield seamless map  $f' : M^c \rightarrow \Omega'$  (Section 7.2)
  - Optimize  $f'$  and  $\Omega'$ , maintaining seamlessness (Section 7.3)

## 5 EXISTENCE OF SEAMLESS PARAMETRIZATIONS

Before discussing the algorithmic details, let us settle the question of existence of seamless parametrizations for prescribed holonomy signatures. In particular, this will allow us to guarantee that the above mentioned rerouting can actually be performed as needed.

While any choice of homology basis loops will yield a holonomy signature, our method relies on bases whose loops' holonomy numbers are some specific values. As a first step towards achieving this, the following proposition gives us a simple way to "add" together two loops so that the holonomy number of the new loop is determined by the holonomy numbers of the constituent loops.

**PROPOSITION 1 (QUASI-ADDITION).** Suppose  $\gamma$  and  $\delta$  are two non-intersecting simple oriented loops and  $\alpha$  is a path from the right side of  $\gamma$  to the right side of  $\delta$  that only intersects these loops at its endpoints and does not contain any cones (see Fig. 6). Then (provided the mesh is suitably refined) there is a simple loop  $\gamma_0$ —which can be made arbitrarily close to  $\delta$ ,  $\gamma$ , and  $\alpha$ —such that

$$k_{\gamma_0}^F = k_{\gamma}^F + k_{\delta}^F - 1$$

If  $\alpha$  is a path from the left side of  $\gamma$  to the left side of  $\delta$  we have nearly the same result, but the holonomy number of  $\gamma_0$  is instead given by

$$k_{\gamma_0}^F = k_{\gamma}^F + k_{\delta}^F + 1$$

A proof and an illustration can be found in appendix A.

**Rerouting around a Cone.** In particular, for a cone  $v_i$ , provided there is a path from the right side of  $\gamma$  to  $v_i^*$ , the above proposition tells us there is a loop  $\gamma_0$  such that  $k_{\gamma_0}^F = k_{\gamma}^F - I_{v_i}^F$ . We refer to the construction of the latter loop as *rerouting  $\gamma$  around  $v_0$*  (with a counterclockwise orientation). On the other hand, if there is a path from the left hand side of  $\gamma$  to  $v_i^*$  (the dual facet of vertex  $v_i$ ), the above proposition gives us a loop  $\gamma_0$  such that  $k_{\gamma_0}^F = k_{\gamma}^F + I_{v_i}^F$ , which we refer to as *rerouting  $\gamma$  around  $v_i$*  with a clockwise orientation.




Fig. 7. Rerouting (ccw, twice in a row) of a loop around a cone of index  $\frac{1}{4}$ .

Moreover, it is clear from the construction of these loops that  $\gamma_0$  is homotopic to  $\gamma$  on  $M$ , although the homotopy will necessarily cross the cone as otherwise the holonomy numbers of the two loops would be the same. Fig. 7 shows an example of rerouting a loop around a cone of index  $\frac{1}{4}$  twice, so as to yield a loop whose holonomy number differs by  $\frac{1}{2}$ .

These observations lead to the following key proposition.

**PROPOSITION 2.** *Let  $H = \{\gamma_1, \dots, \gamma_{2g}\}$  a basis of loops for  $M$  that cuts  $M$  into a topological disk, and let  $v_1, \dots, v_m$  be vertices of  $M$ . Also, let  $k_1, \dots, k_{2g}, I_1, \dots, I_m \in \frac{1}{4}\mathbb{Z}$  and assume  $\gcd(I_1, \dots, I_m) = \frac{1}{4}$ . Then there is another basis of loops  $\delta_1, \dots, \delta_{2g}$  that cuts  $M$  into a disk such that  $k_{\delta_i}^F = k_{\gamma_i}^F + k_i$ ,  $i = 1, \dots, 2g$ , for any seamless parametrization  $F$  that has cones with indices  $I_{v_j}^F = I_j$  at the vertices  $v_1, \dots, v_m$ .*

For a constructive proof see appendix B. Conceptually, we can reroute the loops  $\gamma_i$  one-by-one around suitable subsets of the cones in a manner that preserves the topology of  $H$ , such that their holonomy numbers change exactly by the desired values  $k_i$ .

For the purpose of our method, this result means that we can start from a cut graph formed by the union of  $2g$  loops  $\gamma_1, \dots, \gamma_{2g}$ , and modify these loops using an appropriate choice of integers  $k_1, \dots, k_{2g}$  to yield loops  $\delta_1, \dots, \delta_{2g}$  instead, with any holonomy numbers we want, forming an equivalent signature—under the only condition that  $\gcd(I_1, \dots, I_m) = \frac{1}{4}$ . This ability is sufficient for the method presented in the following to construct a seamless parametrization with the desired holonomy, which constructively shows existence, under the above condition.

**GCD-Condition.** This condition is obviously satisfied as soon as there is even just one cone of index  $\pm \frac{1}{4}$  among all prescribed cones. But this (practically very mild assumption) is not even necessary; even if all indices are of higher magnitude, they may have a greatest common divisor of  $\frac{1}{4}$ . If the gcd is indeed larger than  $\frac{1}{4}$  (a potentially realistic scenario is one with indices restricted to multiples of  $\frac{1}{2}$ ), note that while not all holonomy numbers can be achieved by rerouting, it may still be possible to achieve those desired.

## 6 HOLOMONY-CONSTRAINED CUT GRAPH

The cut graph with particular holonomy pattern is built in three steps. We start by constructing a hole-chain  $G$ ; in deviation from the algorithm described for this purpose in [Campen et al. 2019] we employ soft-guidance by a given input cross-field already in this step. Afterwards a holonomy basis of loops  $H$  is extracted from  $G$ , and its associated target holonomy numbers are derived from the

cross-field. Finally, these loops are rerouted where necessary, i.e., where soft-guidance did not yield exactly those holonomy numbers we require for the subsequent stage.

### 6.1 Field-Guided Hole-Chain

A hole-chain  $G$  on  $M$  is built out of  $g$  loops (non-contractible, non-separating, non-homotopic) and  $2g - 1$  connecting paths. Intuitively, cutting the surface by the  $g$  loops yields a topological sphere with  $2g$  holes, and the  $2g - 1$  paths connect these in a chain-like manner, further cutting the surface to a topological sphere with one hole, i.e. a disk. Fig. 5 shows an example with 4 loops (circular, inside the tunnels) and 7 connectors. We here describe how to construct these loops and connectors guided by a given cross-field.

*Remark:* For certain special cases (genus  $\leq 2$ ) a slightly modified hole-chain structure needs to be chosen. This is done exactly as in the SP method. Likewise, an extra connector path possibly needs to be added; this occurs after rerouting (Section 6.3).

**6.1.1 Field-Guided Loops.** We construct  $g$  non-contractible, non-separating, non-homotopic loops on  $M \setminus C$  following the algorithm of [Diaz-Gutierrez et al. 2009]. To promote cross-field alignment (thus turning number zero along the loop) we employ the field-alignment metric of [Campen et al. 2012] in this process. This in particular means that the loop construction is performed on  $M_4$ , a four-sheeted covering of  $M$ , owing to the four different directions a cross-field specifies per point.

A loop resulting from this, while simple (i.e. intersection-free) on  $M_4$  by construction, may in some cases be self-intersecting when projected down onto  $M$ . Conceptually, it may pass “over” itself on a different sheet of  $M_4$ , which corresponds to an actual crossing on  $M$ . In such a case we fall back to a non-guided construction of a replacement loop directly on  $M$ .

**6.1.2 Field-Guided Connectors.** Connector paths between the loops are selected in a Hamiltonian path manner as in the SP method. By contrast, however, we do not build these from simple shortest paths but again in a field-guided manner. As in the loop construction in Section 6.1.1 we use an anisotropic metric and perform the path search on  $M_4$ . As the above loops have an embedding in  $M_4$  by construction, we know on which sheet of  $M_4$  to start and end the search, respectively: one sheet lower or higher than the respective loop, as a connector will be orthogonal (rather than parallel) to its two incident loops under the final parametrization. Again, should a self-intersecting connector path occur, we fall back to a shortest path computed on  $M$ .

*Remark:* Loops constructed by the fallback method (if any) have no native embedding on  $M_4$ . We locally (at the intended connector start or end point) assign them to the sheet on which the field direction best fits the loop’s tangent, so as to have a reasonable setup for the computation of incident connectors. In any case, however, let us remind that the worst possible outcome of a locally suboptimal choice is a hole-chain that requires some more rerouting—structural correctness is not at stake in this soft-guided approach.

The resulting loops and connectors are embedded in edges of  $M$  and together form the discrete cut graph  $G$ .




Fig. 8. Two equivalent holonomy signatures, based on different signature loops; the different associated holonomy numbers are not shown in the figure. Both are the result of rerouting so as to achieve the required holonomy pattern, therefore the resulting optimized seamless parametrizations based on the cut graphs formed by these loop systems are identical (up to seamless transformation, due to a differently located cut graph).

## 6.2 Homology Basis Extraction

We construct a homology basis  $H$  of  $M$  in the form of  $2g$  loops contained in the cut graph  $G$ . To this end we compute a spanning tree  $T$  in  $G$ . The remainder  $G \setminus T$  consists of  $2g$  edges, called bridges [Erickson and Whittlesey 2005]. For each bridge, its union with the two paths from its incident vertices to the root of  $T$  is a loop, and these  $2g$  loops form a homology basis, and more specifically a system of loops.

Note that these loops may coincide partially. Each of the  $2g$  bridge edges, however, is part of exactly one of these loops only. By rerouting the segments of  $G$  that contain these bridge edges (called bridge segments) we are therefore able to *individually* alter the holonomy number or turning number of each of these  $2g$  loops with respect to a given field. Effectively, the bridge segments are the places where the conceptual  $\alpha$ -path from the proof of Prop. 2 can be attached without intersecting any other basis loops. Ultimately, a modified cut graph  $G'$  with the desired holonomy number for each basis loop can be obtained in this way, as detailed in the following.

## 6.3 Segment Rerouting

For each loop of  $H$  we count its number of left turns  $m_l$  and right turns  $m_r$  (in ccw sense). The holonomy number along this loop in the parametrization we will construct will be  $\frac{1}{4}(m_l - m_r)$  (cf. Fig. 5). If its target holonomy number is  $t$  (e.g., the cross-field turning number along this loop), we need to reroute this loop such that this number changes by  $k = \frac{1}{4}(m_l - m_r) - t$ .

This is performed by rerouting the loop's bridge segment, which we tackle in a two-tier manner. We first attempt to find a replacement path for the segment by an efficient field-guided method, detailed in Section 6.3.1. As this method is not guaranteed to yield a *simple* path (which, however, is needed), where necessary a guaranteed (but less geometry aware) fallback strategy is employed, as described in Section 6.3.2. Note that the resulting loops are not unique; many different equivalent signatures exhibit the desired holonomy numbers. Due to equivalence, however, the final parametrization's topology is not affected by this (see Fig. 8).

*Remark:* Optionally, we may perform a pre-rerouting of the field-guided loops from Section 6.1.1 already before moving on to the connector computation. This is possible because these loops' target holonomy is known to be zero, regardless of how the connectors will interact with them. This pre-rerouting is not necessary for correctness, but empirically it reduces the total amount of rerouting required. As the  $g$  loops do not yet form a complete cut graph that cuts the surface to a disk, one needs to take one additional precaution, though, so as to ensure that a loop is rerouted homotopically. Namely, we cut  $M$  minus the loops to a disk by additional temporary cuts (using the method of [Erickson and Whittlesey 2005]) and perform rerouting within this disk.

**6.3.1 Holonomy-Aware Dijkstra.** Given a bridge segment  $\ell$  of  $G$ , supposed to be rerouted such that the holonomy number of its unique containing loop  $\gamma$  from  $H$  changes by  $k$ , we employ a holonomy-constrained Dijkstra's algorithm, as described in [Campen et al. 2019, §5.1]. This entails the following. We compute a spanning tree of  $M^C$  (cut by the current  $G$ ), rooted on  $\ell$ . Indices of cone vertices are propagated through this tree towards the root, and tree edges are marked with the sum of indices propagated through them. By then keeping track of the sum of these values of edges crossed during Dijkstra's shortest path algorithm (applied to the dual mesh), we can read off the index sum of cones enclosed between a Dijkstra path  $\ell'$  and bridge segment  $\ell$ . The algorithm terminates when a path enclosing the desired index sum (which determines the change to the holonomy number of  $\gamma$ ) is found.

Similar to the path construction on  $M_4$  described in Section 6.1, this holonomy-constrained path search effectively occurs on an (in this case infinite) cover of  $M$  (akin to the universal cover of  $M \setminus C$ ). Consequently, a non-simple path (in  $M$ ) can be the result in some cases. The following guaranteed fallback strategy takes care of such cases.

**6.3.2 Fallback Strategy.** By following the rerouting construction used in the proof of Prop. 2, a proper simple replacement path for a bridge segment  $\ell$  can safely be found. Let  $v$  be a vertex on a bridge segment  $\ell$  whose loop's holonomy number needs to be increased by  $k$ .

Assume for a moment that there is a cone vertex  $v^*$  with  $|I(v^*)| \leq k$ . Let  $\alpha$  be the shortest path from cone  $v^*$  to  $v$ —either meeting  $\ell$  from the right if  $I$  and  $k$  have opposite sign, or from the left in the case of equal sign. Among all suitable cones, we choose the one for which the path  $\alpha$  is shortest, so as to reduce the amount of modification.

Let the two vertices on  $\ell$  which are directly adjacent to  $v$  be  $v^-$  and  $v^+$ . Closely following the conceptual rerouting from Fig. 11 we remove the edges  $v^-v$  and  $vv^+$  from  $\ell$  and replace them by the path from  $v^-$  to  $v^+$  tightly along  $\alpha$  and around  $v^*$ . Where necessary we split edges of  $M$  to make room so that this path does not touch any other part of  $G$ .

This changes the loop's holonomy number by  $I(v^*)$  or  $-I(v^*)$ , depending on which side of  $\ell$  the path  $\alpha$  connects to. This procedure can be repeated as long as the remaining difference  $k \leftarrow k \pm I(v^*)$  is not zero yet.

While this strategy proved sufficient in all practical test cases (cf. Section 8) (and indeed is guaranteed to work if there is at least

one cone of index  $\pm \frac{1}{4}$ ), in general a greedy selection of reroute-cones  $v^*$  in this greedy manner is insufficient. Instead, let  $V$  be a multiset of cones such that its index sum is  $k$ . Under the GCD-condition (Section 5),  $V$  exists, as also exploited in the proof of Prop. 2. By selecting the cones from  $V$  as  $v^*$  in the above one after the other (also considering multiplicity), the desired result is achieved. A suitable multiset  $V$ , i.e. a subset of cones and multiplicities, is easily computed using the Extended Euclidean Algorithm. The incorporation of distances also in this general case would enable smaller rerouting modification, but given the practical irrelevance of this multiset-case, the effort would hardly pay off.

*Remark:* In practice we tentatively perform the fallback-rerouting starting from multiple root vertices  $v$  (sampled equidistantly on the bridge segment; we use 10 samples in our experiments) and retain the result of shortest length to reduce the complexity of the final cut graph.

## 7 SEAMLESS PARAMETRIZATION

Once the final, i.e., rerouted and possibly extended (recall the remark in Section 6.1), cut graph  $G'$  is available, the next step is to construct a domain that is compatible with the cut surface  $M^c$  and suitable to serve as parameter domain for a seamless parametrization of  $M$ . The domain shape is derived from a conformal metric computed on  $M^c$  with prescribed cones and prescribed boundary curvature.

### 7.1 Cut Graph aligned Metric

A key role in the SP method that we build on is played by a discrete conformal metric computation on the cut mesh  $M^c$ . While the conformal metric algorithm from [Campen and Zorin 2017] that is used in SP works adequately in most cases, it does not provide strict guarantees of convergence. The cut graph rerouting used in our method can sometimes lead to rather complex cut shapes, thus boundary shapes of  $M^c$ , implying additional metric distortion, making the problem instances particularly challenging.

Very recently, a novel algorithm for discrete metric computation with prescribed (boundary and cone) angles has been proposed [Campen et al. 2021; Gillespie et al. 2021], based on mathematical insights [Gu et al. 2018; Springborn 2019] that guarantee convergence. We employ an implementation based on this work.

Using this algorithm we compute a discrete metric (i.e., edge lengths) for  $M^c$ , prescribing the angles of cone vertices and the geodesic curvature on the boundary, i.e., along the cut  $G'$ . Concretely, the segments of  $G'$  are constrained to be straight under the resulting metric, and the corners (at branch points of  $G'$ ) are constrained to be right.

### 7.2 Padding

Under the computed metric the two boundary segments of  $M^c$  corresponding to a common segment of  $G'$  are straight, their mutual angle is a multiple of  $\pi/2$  (as required for seamlessness), but their lengths may differ. The SP method uses so-called *padding*, i.e., parametrically stretching out strips of the surface under the metric along the boundary segments so that the boundary segments' lengths expand to equalize the lengths of all paired segments, and this provably is always possible.




Fig. 9. Illustration of padding operation (in parameter domain). A thin strip along the top straight cut segment (with no interior vertices) is stretched in vertical direction by its required padding width. Then, vertices are shifted horizontally to match their mates across the cut.

The following steps perform this padding, analogous to the original SP method from [Campen et al. 2019]:

- (1) Add cuts from all cones to the boundary of  $M^c$ . Make sure each boundary segment is reached in at most one point. Around each interior vertex the angles under the metric from Section 7.1 now sum up to  $2\pi$ , i.e. the cut surface is flat.
- (2) Lay out this flat mesh in the plane, i.e., assign  $(u, v)$ -parameters to all vertices of the mesh, e.g., in a breadth-first traversal. The global rotation is chosen such that the straight boundary segments are aligned with  $u$  or  $v$  axis directions. This yields the (non-seamless) domain  $\Omega$ .
- (3) For each straight boundary segment, compute the amount of padding (width  $w_i$ ) required to equalize parametric lengths of the paired segments, using the equation system of the SP method.
- (4) Along each segment, determine a parametrically rectangular strip free of cones, and make the mesh conform to this strip by inserting the strip's boundary line by splitting. Then apply a stretch transformation to the  $(u, v)$ -coordinates inside the strip, so as to shift the segment in perpendicular direction by its padding width  $w_i$  (Fig. 9).
- (5) In cases where the cut graph has cut the surface into more than one disk, glue these together parametrically along pairs of boundary segments by means of rigid transformations applied to the  $(u, v)$ -coordinates to finally obtain a map  $F'$  onto a single connected domain  $\Omega'$ .

### 7.3 Optimization

As a final step, we optimize the established map for reduced distortion. As objective, we employ a local cross-field (orientation and sizing) alignment energy  $E_A$  [Bommes et al. 2009] and add (with a small factor of  $s = 10^{-3}$ ) the symmetric Dirichlet energy  $E_D$  [Rabinovich et al. 2017], which contributes its barrier behavior to prevent parametric inversions in the course of optimization. Linear constraints are added to preserve seamlessness. We use a projected Newton solver and use an explicit triangle inversion check in the line search [Smith and Schaefer 2015], using exact predicates, to reliably maintain local injectivity. We experimentally discovered that using an unconstrained Newton optimizer over the set of independent variables computed using a reduced row echelon form of the constraint matrix is numerically more stable than solving a KKT system at each iteration, leading to faster convergence.

We emphasize that we do not aim to address map optimality here; our focus is on constructing a topologically correct initial map, subject to further improvement geometrically.




Fig. 10. Comparison of seamless parametrizations on surfaces of non-trivial topology, computed by the bare SP method [Campen et al. 2019] (row b, e) and by our method (row d, f). The used cut graphs are shown in red, the initial hole-chain used for SP (row a, e) and the rerouted version used by our method (row c, f). Notice their topologically differing structure (i.e. they wind around some handles or cones differently), as well as the higher distortion of the results by the bare SP method due to being unable to properly align to the underlying smooth cross-field for topological reasons. Notice that this distortion cannot be reduced further by continuous optimization; there are topological obstacles.

Table 1. Statistics about the number of cut segment reroutings performed. It is further split into the numbers of field-guided and fallback reroutings.

Model	Genus	#Reroutings	#Field-Guided	#Fallback
twirl	1	0	0	0
robocatdeci	1	0	0	0
knot1	1	0	0	0
holes3	3	0	0	0
dancer2	1	0	0	0
sculpt	2	0	0	0
fertilitytri	4	0	0	0
rockerarm	1	1	1	0
genus3	3	1	1	0
elk	1	1	1	0
trimstar	1	1	1	0
wrench50K	1	1	1	0
bumpytorus	1	1	1	0
dancer25k	1	1	1	0
camel	1	1	1	0
dragonstandrecon	1	1	1	0
pulley	1	1	1	0
kitten	1	1	1	0
knot	1	1	1	0
masterycylinder	3	1	1	0
eight	2	1	1	0
femur	2	2	2	0
block	3	2	2	0
greeksculpture	4	2	2	0
elephant	3	2	2	0
thaistatue	3	2	2	0
oilpump	4	2	2	0
neptune0	3	2	2	0
carter	7	2	2	0
cup	2	2	2	0
botijo	5	3	3	0
chair	7	3	3	0
rollingstage	7	3	3	0
helmet	3	4	2	2
pegaso	6	4	4	0
chair	7	4	4	0
bozbebozzel	5	5	5	0
dancingchildren	8	5	5	0
grayloc	9	6	6	0
seahorse2	8	10	5	5
raptor50K	10	12	6	6
heptoroid	22	15	14	1
gearbox	78	57	43	14
filigree	65	73	40	33
brain	57	83	70	13
vhskin	79	128	21	107

## 8 EVALUATION

We apply our method to a dataset of 3D models with cross-fields [Myles et al. 2014]. We restrict ourselves to models of genus  $> 0$ , as on topologically trivial surfaces there are no global holonomy aspects to account for. The method succeeds in generating a cut graph with exactly the needed holonomy numbers in all cases. As all cases satisfy the  $\text{gcd}=\frac{1}{4}$  condition, and all crucial operations are combinatorial/discrete, the general success of this step is indeed to be expected. For each model, Table 1 lists the number of rerouting operations that our method performed.

The construction of a seamless parametrization based on this cut then succeeds in most cases; in six, however, the initial metric distortion is very high, causing subsequent steps (padding or the simple distortion optimization approach) to get into numerical trouble. In Figs. 1 and 10 obtained optimized seamless parametrizations for examples from the dataset are shown, matching the input cross-field by construction. Table 2 reports the final distortion of these.

Table 2. Residual energy (normalized by surface area) for the models from Figs. 1 and 10. The columns “without rerouting” correspond to the direct application of SP, without regard for global holonomy. From the last column the advantage in terms of field alignment and distortion becomes clear.

Fig.	Model	with rerouting		ours/SP
		$E_A+sE_D$	$E_A$	
1	cup	0.0125	0.0125	0.3288 0.2882 3.8%
9	block	0.0136	0.0136	0.1115 0.1052 12.2%
9	eight	0.0350	0.0328	0.1524 0.1432 22.9%
9	genus3	0.0221	0.0208	0.1891 0.1747 11.7%
9	oilpump	0.0296	0.0293	0.0450 0.0370 65.7%
9	rollingstage	0.0127	0.0124	0.2666 0.1163 4.8%
9	thaistatue	0.0225	0.0225	0.0272 0.0257 82.7%

## 8.1 Comparison

To demonstrate the importance of our contribution in the context of guaranteed locally injective seamless parametrization construction, we also apply the bare SP method of [Campen et al. 2019] (which takes local holonomy (cones) but not global holonomy into account) to these models.

While SP is able to respect the singularities of the prescribed field by construction, whether or not its resulting map matches the cross-field topologically is essentially a matter of chance. If the cross-field is very smooth (as generally is the case in this data set) and the cut graph for the map is constructed from certain shortest paths, the chance of a match may be higher than that of any particular mismatch. Nevertheless, we encounter a mismatch for a large number of models—in line with the fact that, as can be seen in Table 1, our method had to employ at least one rerouting operation in the majority of cases. In case of a mismatch, the resulting map cannot continuously be optimized to achieve reasonable alignment between map isolines and the field, as there is a topological obstacle. This can be observed in Table 2, where the remaining final distortion is significantly higher when not employing rerouting. The difference is also illustrated in Fig. 10. Our method, in essence by adjusting the cut graph in the described manner, ensures a topological match between the signature induced by the cross-field and the signature of the generated seamless parametrization.

## 9 CONCLUSION AND FUTURE WORK

We have explored the relation between cross-fields and seamless surface parametrizations (and therefore quadrangulations) on a topological level. A key insight is that there are hardly any practically important obstacles to generating a seamless parametrization (or quadrangulation) that topologically matches a given cross-field. We have described a method to generate such a seamless parametrization, given an input cross-field or an abstract topological specification in form of a holonomy signature. It is based on a variation of the SP method [Campen et al. 2019], with the main difference being:

- The initial hole chain cut graph is constructed taking cross-field guidance into account.
- The hole chain is then modified by extracting a loop basis and rerouting of loop segments based on our theory.
- The generation of a cut-aligned parametrization is performed using a different, theoretically sound conformal mapping method.

From the SP method that we employ for the parametrization construction we inherit the restriction to surfaces without boundary. While there are no fundamental obstacles to adding boundary support to our rerouting procedure, padding feasibility requires additional theory in this more general context. The situation regarding support for alignment to feature curves, which is of interest in some use cases of seamless parametrizations, is very similar.

The algorithm stage described in Section 6, in particular the holonomy-constrained cut graph generation using rerouting, relies on discrete operations and therefore is not only sound theoretically, but can be executed without the risk of numerical issues and limits in practice. The algorithm stage described in Section 7 (initial parametrization followed by constrained optimization), by contrast, involves numerical computations, with consequent limits in practice. While for initial parametrization a discrete approach is imaginable [Zhou et al. 2020], at least for the final distortion optimization a numerical approach is inevitable.

While we observe the choice of loops that form the initial cut graph to not affect the final result conceptually (Fig. 8), the distortion of the initial parametrization, and therefore the numerical challenges in the final optimization, can depend strongly on this choice. By testing various random root placements for the loop construction [Diaz-Gutierrez et al. 2009] employed in Section 6.1.1, initial parametrizations of low distortion could be found, but a more direct approach—or a more resilient final distortion optimization technique—is desirable.

The GCD-condition asserts that, for any given signature, there is an equivalent signature whose loops have any desired set of holonomy numbers. It therefore is a sufficient condition for the existence of a seamless parametrization that topologically matches a given signature. It is not necessary, though. While likely of limited practical relevance, the exploration of even tighter conditions may be interesting.

## ACKNOWLEDGMENTS

Supported by NSF CAREER award 1652515, DMS-1821334, OAC-1835712, OIA-1937043, CHS-1908767, CHS-1901091, a Sloan Fellowship, a gift from Adobe Research, a gift from nTopology, and a gift from Advanced Micro Devices, Inc.

## REFERENCES

- Noam Aigerman and Yaron Lipman. 2015. Orbifold Tutte Embeddings. *ACM Trans. Graph.* 34, 6 (2015), 190:1–190:12.
- Noam Aigerman and Yaron Lipman. 2016. Hyperbolic Orbifold Tutte Embeddings. *ACM Trans. Graph.* 35, 6, Article 217 (2016), 14 pages.
- Mirela Ben-Chen, Craig Gotsman, and Guy Bunin. 2008. Conformal Flattening by Curvature Prescription and Metric Scaling. *Computer Graphics Forum* (2008).
- David Bommes, Marcel Campen, Hans-Christian Ebke, Pierre Alliez, and Leif Kobbelt. 2013a. Integer-Grid Maps for Reliable Quad Meshing. *ACM Trans. Graph.* 32, 4 (2013), 98:1–98:12.
- David Bommes, Bruno Lévy, Nico Pietroni, Enrico Puppo, Claudio Silva, Marco Tarini, and Denis Zorin. 2013b. Quad-Mesh Generation and Processing: A Survey. In *Computer Graphics Forum*. Wiley Online Library.
- David Bommes, Henrik Zimmer, and Leif Kobbelt. 2009. Mixed-integer quadrangulation. *ACM Trans. Graph.* 28, 3 (2009), 77.
- Alon Bright, Edward Chien, and Ofir Weber. 2017. Harmonic Global Parametrization with Rational Holonomy. *ACM Trans. Graph.* 36, 4 (2017).
- Marcel Campen. 2017. Partitioning Surfaces Into Quadrilateral Patches: A Survey. *Computer Graphics Forum* 36, 8 (2017), 567–588.
- Marcel Campen, David Bommes, and Leif Kobbelt. 2012. Dual Loops Meshing: Quality Quad Layouts on Manifolds. *ACM Trans. Graph.* 31, 4 (2012).
- Marcel Campen, David Bommes, and Leif Kobbelt. 2015. Quantized global parametrization. *ACM Trans. Graph.* 34, 6 (2015), 192.
- Marcel Campen, Ryan Capouillez, Hanxiao Shen, Leyi Zhu, Daniele Panizzo, and Denis Zorin. 2021. Efficient and Robust Discrete Conformal Equivalence with Boundary. *ACM Trans. Graph.* 40, 6 (2021).
- Marcel Campen and Leif Kobbelt. 2014. Dual Strip Weaving: Interactive Design of Quad Layouts Using Elastica Strips. *ACM Trans. Graph.* 33, 6 (2014), 183:1–183:10.
- Marcel Campen, Hanxiao Shen, Jiaran Zhou, and Denis Zorin. 2019. Seamless Parametrization with Arbitrary Cones for Arbitrary Genus. *ACM Trans. Graph.* 39, 1 (2019).
- Marcel Campen and Denis Zorin. 2017. Similarity Maps and Field-Guided T-Splines: a Perfect Couple. *ACM Trans. Graph.* 36, 4 (2017).
- Wei Chen, Xiaopeng Zheng, Jingyao Ke, Na Lei, Zhongxuan Luo, and Xianfeng Gu. 2019. Quadrilateral mesh generation I: Metric based method. *Comput. Methods Appl. Mech. Engrg.* 356 (2019), 652–668.
- Wei Chen, Xiaopeng Zheng, Jingyao Ke, Na Lei, Zhongxuan Luo, and Xianfeng Gu. 2020. Quadrilateral Mesh Generation II: Meomorphic Quartic Differentials and Abel-Jacobi Condition. *Comput. Methods Appl. Mech. Engrg.* 366 (2020).
- Edward Chien, Zohar Levi, and Ofir Weber. 2016. Bounded Distortion Parametrization in the Space of Metrics. *ACM Trans. Graph.* 35, 6 (2016).
- Keenan Crane, Mathieu Desbrun, and Peter Schröder. 2010. Trivial Connections on Discrete Surfaces. *Computer Graphics Forum* 29, 5 (2010), 1525–1533.
- Pablo Diaz-Gutierrez, David Epstein, and Meenakshi Sundaram Gopi. 2009. Curvature aware fundamental cycles. In *Computer Graphics Forum*, Vol. 28. 2015–2024.
- Hans-Christian Ebke, Patrick Schmidt, Marcel Campen, and Leif Kobbelt. 2016. Interactively Controlled Quad Remeshing of High Resolution 3D Models. *ACM Trans. Graph.* 35, 6 (2016), 218:1–218:13.
- Jeff Erickson and Kim Whittlesey. 2005. Greedy optimal homotopy and homology generators. In *SODA*, Vol. 5. 1038–1046.
- Xianzhong Fang, Hujun Bao, Yiyi Tong, Mathieu Desbrun, and Jin Huang. 2018. Quadrangulation through mose-parameterization hybridization. *ACM Trans. Graph.* 37, 4 (2018), 92.
- Xiao-Ming Fu, Yang Liu, and Baining Guo. 2015. Computing Locally Injective Mappings by Advanced MIPS. *ACM Trans. Graph.* 34, 4 (2015).
- Mark Gillespie, Boris Springborn, and Keenan Crane. 2021. Discrete Conformal Equivalence of Polyhedral Surfaces. *ACM Trans. Graph.* 40, 4 (2021).
- Steven J. Gortler, Craig Gotsman, and Dylan Thurston. 2006. Discrete one-forms on meshes and applications to 3D mesh parameterization. *Computer Aided Geometric Design* 23, 2 (2006), 83 – 112.
- Xianfeng Gu and Shing-Tung Yau. 2003. Global conformal surface parameterization. In *Proc. Symp. Geometry Processing 2003*. 127–137.
- Xianfeng David Gu, Feng Luo, Jian Sun, and Tianqi Wu. 2018. A discrete uniformization theorem for polyhedral surfaces. *Journal of differential geometry* 109, 2 (2018).
- Allen Hatcher. 2002. *Algebraic Topology*. Cambridge University Press.
- Eden Fedida Hefetz, Edward Chien, and Ofir Weber. 2019. A Subspace Method for Fast Locally Injective Harmonic Mapping. In *Computer Graphics Forum*, Vol. 38. 105–119.
- K. Hormann and G. Greiner. 2000. MIPS: An Efficient Global Parametrization Method. In *Curve and Surface Design: Saint-Malo 1999*. Vanderbilt University Press, 153–162.
- Ernest Jucović and Marián Trenkler. 1973. A theorem on the structure of cell-decompositions of orientable 2-manifolds. *Mathematika* 20, 01 (1973), 63–82.
- F. Kälberer, M. Nieser, and K. Polthier. 2007. QuadCover: Surface Parameterization using Branched Coverings. *Computer Graphics Forum* 26, 3 (2007), 375–384.
- Liliya Kharevych, Boris Springborn, and Peter Schröder. 2006. Discrete conformal mappings via circle patterns. *ACM Trans. Graph.* 25 (April 2006), 412–438. Issue 2.
- P. Knupp. 1995. Mesh Generation Using Vector Fields. *J. Comput. Phys.* 119, 1 (1995).
- Denis Kovacs, Ashish Myles, and Denis Zorin. 2011. Anisotropic quadrangulation. *Computer Aided Geometric Design* 28, 8 (2011), 449 – 462.
- Shahar Z. Kovalsky, Meirav Galun, and Yaron Lipman. 2016. Accelerated Quadratic Proxy for Geometric Optimization. *ACM Trans. Graph.* 35, 4 (2016), 134:1–134:11.
- W. Li, B. Vallet, N. Ray, and B. Levy. 2006. Representing Higher-Order Singularities in Vector Fields on Piecewise Linear Surfaces. *IEEE TVCG* 12, 5 (2006), 1315–1322.
- Yaron Lipman. 2012. Bounded Distortion Mapping Spaces for Triangular Meshes. *ACM Trans. Graph.* 31, 4 (2012), 108:1–108:13.
- Max Lyon, Marcel Campen, David Bommes, and Leif Kobbelt. 2019. Parametrization Quantization with Free Boundaries for Trimmed Quad Meshing. *ACM Trans. Graph.* 38, 4 (2019).
- Manish Mandad and Marcel Campen. 2020. Efficient piecewise higher-order parametrization of discrete surfaces with local and global injectivity. *Computer-Aided Design* 127 (2020).
- Martin Marinov, Marco Amagliani, Tristan Barback, Jean Flower, Stephen Barley, Suguru Furuta, Peter Charrot, Iain Henley, Nanda Santhanam, G. Thomas Finnigan, Siavash Meshkat, Justin Hallet, Maciej Sapun, and Paweł Wolski. 2019. Generative Design Conversion to Editable and Watertight Boundary Representation. *Computer-Aided Design* 115 (2019), 194 – 205.
- Ashish Myles, Nico Pietroni, and Denis Zorin. 2014. Robust Field-aligned Global Parametrization. *ACM Trans. Graph.* 33, 4 (2014), 135:1–135:14.

- Ashish Myles and Denis Zorin. 2012. Global parametrization by incremental flattening. *ACM Trans. Graph.* 31, 4 (2012), 109.
- Ashish Myles and Denis Zorin. 2013. Controlled-distortion constrained global parametrization. *ACM Transactions on Graphics* 32, 4 (2013), 105.
- Michael Rabinovich, Roi Poranne, Daniele Panozzo, and Olga Sorkine-Hornung. 2017. Scalable Locally Injective Mappings. *ACM Trans. Graph.* 36, 2 (2017), 16:1–16:16.
- Nicolas Ray, Bruno Vallet, Laurent Alonso, and Bruno Lévy. 2009. Geometry-Aware Direction Field Processing. *ACM Trans. Graph.* 29, 1 (2009).
- Nicolas Ray, Bruno Vallet, Wan Chiu Li, and Bruno Lévy. 2008. N-Symmetry Direction Field Design. *ACM Trans. Graph.* 27, 2 (2008).
- Christian Schüller, Ladislav Kavan, Daniele Panozzo, and Olga Sorkine-Hornung. 2013. Locally Injective Mappings. *Computer Graphics Forum* 32, 5 (2013), 125–135.
- Anna Shtengel, Roi Poranne, Olga Sorkine-Hornung, Shahar Z. Kovvalsky, and Yaron Lipman. 2017. Geometric Optimization via Composite Majorization. *ACM Trans. Graph.* 36, 4 (2017), 38:1–38:11.
- Jason Smith and Scott Schaefer. 2015. Bijective Parameterization with Free Boundaries. *ACM Trans. Graph.* 34, 4, Article 70 (2015), 9 pages.
- Yousuf Soliman, Dejan Slepčev, and Keenan Crane. 2018. Optimal Cone Singularities for Conformal Flattening. *ACM Trans. Graph.* 37, 4 (2018), 105:1–105:17.
- Boris Springborn. 2019. Ideal Hyperbolic Polyhedra and Discrete Uniformization. *Discrete & Computational Geometry* (2019), 1–46.
- Y. Tong, P. Alliez, D. Cohen-Steiner, and M. Desbrun. 2006. Designing quadrangulations with discrete harmonic forms. *Symposium on Geometry Processing* (2006), 201–210.
- Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommes, Klaus Hildebrandt, and Mirela Ben-Chen. 2016. Directional Field Synthesis, Design, and Processing. *Comp. Graph. Forum* 35, 2 (2016).
- Jiaran Zhou, Marcel Campen, Denis Zorin, Changhe Tu, and Claudio T. Silva. 2018. Quadrangulation of non-rigid objects using deformation metrics. *Computer Aided Geometric Design* 62 (2018), 3–15.
- J. Zhou, C. Tu, D. Zorin, and M. Campen. 2020. Combinatorial Construction of Seamless Parameter Domains. *Computer Graphics Forum* 39, 2 (2020), 179–190.
- Yufeng Zhu, Robert Bridson, and Danny M. Kaufman. 2018. Blended Cured Quasi-newton for Distortion Optimization. *ACM Trans. Graph.* 37, 4 (2018), 40:1–40:14.

## A PROOF OF PROPOSITION 1

**PROOF.** We consider the case where  $\alpha$  is a path between the right hand sides of the two loops. If the mesh is suitably refined, there is a topological disk  $D$  in the dual mesh that contains  $\alpha$  and does not contain any cones (Fig. 11 left). Without loss of generality, we may assume that  $\partial D$  intersects  $\gamma$  along a single nontrivial path  $\beta_\gamma$  and intersects  $\delta$  along a similar path  $\beta_\delta$ . We denote the endpoints of  $\beta_\gamma$  as  $f_1^*$  and  $f_2^*$  and the endpoints of  $\beta_\delta$  as  $g_1^*$  and  $g_2^*$ . We now can define  $\gamma_0$  as the simple loop given by traversing  $\gamma \setminus \beta_\gamma$  (with respect to the orientation of this loop) starting at  $f_2^*$ , then traversing the component of  $\partial D$  (with boundary orientation) from  $f_1^*$  to  $g_2^*$ , then traversing  $\delta \setminus \beta_\delta$ , and finally by traversing  $\partial D$  from  $g_1^*$  to  $f_2^*$ . We note that we can choose  $D$  so that the boundary is arbitrarily close to  $\alpha$  and thus so  $\gamma_0$  is arbitrarily close to the original loops and path.

Since  $\partial D$  is the boundary of a topological disk that does not contain any cones, we have that  $\kappa_{\partial D}^F = 2\pi$ . In the computation of  $\kappa_{\gamma_0}^F$ , we have that the signed angles satisfy  $\alpha_{\gamma_0}(f^*) = \alpha_\gamma(f^*)$  for




Fig. 11. Quasi-additivity of holonomy numbers, on the same example as in Fig. 6. The inset on the right is a blow-up of the spot circled on the left.




Fig. 12. Example of iteratively rerouting one loop around two singularities. Left: initial state with given loop  $\gamma_i$  and two paths  $\alpha_j, \alpha'_j$  connecting to the singularity  $v_j$ . Center: reroute around singularity  $v_j$  and find paths  $\alpha_k, \alpha'_k$  for the next singularity  $v_k$ . Right: result after rerouting around  $v_j$  and  $v_k$

$f^* \in \gamma \setminus \beta_\gamma$  and  $\alpha_{\partial D}(f^*) = -\alpha_\gamma(f^*)$  for  $f^* \in \beta_\gamma \setminus \{f_1^*, f_2^*\}$  (Fig. 11 right). Furthermore, since  $\alpha$  intersects  $\gamma$  on the right, we must have that the angle of  $f_1$  that corresponds to  $\alpha_\gamma(f_1^*)$  is on the left hand side of  $\gamma$ , a different angle of  $f_1$  corresponds to  $\alpha_{\partial D}(f_1^*)$  and is on the left hand side of  $\partial D$ , and the final angle of  $f_1$  corresponds to  $\alpha_{\gamma_0}(f_1^*)$  and is on the right hand side of  $\gamma_0$ . Therefore, we have that

$$\alpha_\gamma(f_1^*) + \alpha_{\partial D}(f_1^*) - \alpha_{\gamma_0}(f_1^*) = \pi,$$

and by a similar analysis the same result for  $f_2^*$  is obtained. The situation is similar for  $\delta$  and  $\partial D$ , so we have that

$$\begin{aligned} \kappa_{\gamma_0}^F &= \sum_{f^* \in \gamma_0} \alpha_{\gamma_0}(f^*) \\ &= \sum_{f^* \in \gamma} \alpha_\gamma(f^*) + \sum_{f^* \in \delta} \alpha_\delta(f^*) + \sum_{f^* \in \partial D} \alpha_{\partial D}(f^*) - 4\pi \\ &= \kappa_\gamma^F + \kappa_\delta^F - 2\pi. \end{aligned}$$

Thus, we have that

$$k_{\gamma_0}^F = k_\gamma^F + k_\delta^F - 1.$$

The proof where  $\alpha$  is on the left hand side of the two loops is analogous.  $\square$

## B PROOF OF PROPOSITION 2

**PROOF.** We have that cutting  $M$  along  $\gamma_1, \dots, \gamma_{2g}$  results in a disk, so, if the mesh is sufficiently refined, for any  $\gamma_i$  there is a path  $\alpha_j/\alpha'_j$  (Fig. 12 left) from either side of  $\gamma_i$  to any  $v_j^*$  such that neither path intersects any of the other basis loops or cones. Thus, by the above, we may reroute  $\gamma_i$  around  $v_j$  clockwise or counterclockwise to obtain a new loop  $\gamma'_i$  such that  $H' = (H \setminus \{\gamma_i\}) \cup \{\gamma'_i\}$  also cuts  $M$  into a topological disk and such that, for any seamless parametrization  $F$  satisfying the properties listed in the proposition, we have

$$k_{\gamma'_i}^F = k_{\gamma_i}^F \pm I_{v_j}^F = k_{\gamma_i}^F \pm I_j$$

Since  $H'$  still cuts  $M$  to a disk, we may still reroute any loop around any cone with either orientation, so we may iteratively reroute the basis loops to modify their holonomy number by integer multiples of  $I_j$ . Since  $\frac{1}{4}$  is the greatest common divisor of  $I_1, \dots, I_m$ , we have there are integers  $a_i$  such that

$$\frac{1}{4} = \sum_{i=1}^m a_i I_i$$

Thus, we have that iteratively rerouting each loop  $\gamma_i$  around the cone  $v_j$   $|4k_i a_j|$  times, with orientation determined by the signs of  $k_i$  and  $a_j$ , will give us the desired system of loops.  $\square$