

Harmonic Global Parametrization with Rational Holonomy

ALON BRIGHT*, Bar-Ilan University
EDWARD CHIEN*, Bar-Ilan University
OFIR WEBER, Bar-Ilan University




Fig. 1. Robustness test. We present results of our method (HGP) for four challenging models from the [Myles et al. 2014] benchmark. Many other state-of-the-art methods: [Aigerman et al. 2014], [Levi and Zorin 2014], [Springborn et al. 2008], failed to produce locally injective seamless parametrization maps. The method of [Chien et al. 2016b] managed to produce locally injective maps, but required more than an hour to parametrize the four models above. HGP required just one or two iterations of a second-order cone program (SOCP) per model, and took less than a minute. The heat maps illustrate the conformal distortion k and show that we produce maps with relatively low distortion. The cone points are denoted by the colored dots.

We present a method for locally injective seamless parametrization of triangular mesh surfaces of arbitrary genus, with or without boundaries, given desired cone points and rational holonomy angles (multiples of $2\pi/q$ for some positive integer q). The basis of the method is an elegant generalization of Tutte’s “spring embedding theorem” to this setting. The surface is cut to a disk and a harmonic system with appropriate rotation constraints is solved, resulting in a harmonic global parametrization (HGP) method. We show a remarkable result: that if the triangles adjacent to the cones and boundary are positively oriented, and the correct cone and turning angles are induced, then the resulting map is guaranteed to be locally injective. Guided by this result, we solve the linear system by convex optimization, imposing convexification frames on *only* the boundary and cone triangles, and minimizing a Laplacian energy to achieve harmonicity. We compare HGP to state-of-the-art methods and see that it is the most robust, and is significantly faster than methods with comparable robustness.

CCS Concepts: • **Computing methodologies** → **Mesh models**; **Mesh geometry models**; • **Mathematics of computing** → **Topology**;

*Equal contributions

This work is supported by the Israel Science Foundation, under grants 1869/15 and 2102/15.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2017/7-ART89 \$15.00
DOI: <http://dx.doi.org/10.1145/3072959.3073646>

Additional Key Words and Phrases: mesh parametrization, global parametrization, injective maps, discrete harmonic

ACM Reference format:

Alon Bright, Edward Chien, and Ofir Weber. 2017. Harmonic Global Parametrization with Rational Holonomy. *ACM Trans. Graph.* 36, 4, Article 89 (July 2017), 15 pages.

DOI: <http://dx.doi.org/10.1145/3072959.3073646>

1 INTRODUCTION

Parametrization is a foundational problem in computer graphics and geometry processing, and its solution is relevant to many other processes such as quadrangulation, remeshing, texture mapping, shape correspondence, rotational symmetry fields-and-patterns design, and compression.

A key feature that is desired in mesh parametrizations is that of local injectivity, and harmonic methods have proven useful in achieving such parametrizations from the beginning. For disk meshes, a classic solution involves the use of discrete harmonic maps, based on a seminal result of Tutte [Tutte 1963]. The method is linear, so is quite simple and fast, and ensures an embedding, but the boundary of the disk must be fixed in \mathbb{R}^2 as the boundary of a convex polygon.

More generally, one would like to allow for non-convex boundaries and for more complicated topologies. Theoretical results in this direction were presented in [Gortler et al. 2006]. By considering discrete harmonic 1-forms and their indices, the authors provide an elegant proof of Tutte’s classic result. With the same methods, they show that for a disc, potentially with punctures and non-convex

boundary embeddings (with appropriate turning numbers), that the resulting discrete harmonic map will still be an embedding if the boundary triangles (or n -gons) are non-degenerate and positively oriented.

For more complicated topologies, the mesh must be cut to a disc before mapping, and they suggest integrating two linearly independent harmonic 1-forms over the cut mesh. This works (provably) in the torus case, but runs into significant problems for higher genus, requiring stronger conditions than linear independence on the two harmonic 1-forms. In addition, distortion for the resulting parametrizations is typically quite high, if one does not allow for a careful choice of cut and cone placement, and this method does not allow for any control in that respect.

Furthermore, at such cone points, if the cone angle is not a multiple of 2π , there is no sensible linear harmonicity condition. These problems were tackled in [Tong et al. 2006], which asks the user to design a singularity graph, decomposing the mesh into patches (disks topologically). Harmonic 1-forms are solved for on individual patches with various continuity conditions enforced between patches. Harmonicity constraints are dropped at the meta-vertices of the singularity graph, allowing for any cone angles that are multiples of $\pi/2$.

In the first part of this paper, given cone points and desired rational holonomies, we present a linear harmonic system similar to that of Tong et al. It takes a mesh cut to a disc and enforces rotation conditions on the two copies of any cut edge to achieve holonomy angles that are multiples of $2\pi/q$, for any positive integer q (see Section 4). The resulting parametrization maps are called q -CCMs (for convex combination maps). Utilizing the index-counting arguments of [Gortler et al. 2006], we generalize their results to a mesh S of any topology:

THEOREM 6.1. *Let f denote a q -CCM of S , with specified cone points and holonomy angles determining the rotation constraints. If the cone and boundary triangles are mapped in an orientation-preserving manner, and the induced metric on S achieves the desired cone angles and turning angles, then f is locally injective.*

In the proof, we show that a solution to our system results in a 2-dimensional space of mostly harmonic 1-forms on a q -fold branched cover, and it is the index of these 1-forms that is analyzed to prove local injectivity. Inspiration for this perspective was drawn from [Kälberer et al. 2007], and the Riemann-Hurwitz formula is used to determine the topology of the branched cover. Additionally, we modify the definition of turning angle to allow for boundary vertices with turning angle less than $-\pi$ and generalize the results of [Gortler et al. 2006] in this fashion as well. Lastly, we note that in this work, we restrict to triangle meshes, as these are of the greatest practical interest, but generalization to polygonal meshes is straightforward.

Inspired by these results, we design a method that enforces positive orientation with convexification frames [Lipman 2012], but only requires them on the cone and boundary triangles. As the number of frame conditions needed is much fewer than other frame-based approaches, this reduces feasibility issues and enhances the speed of the optimization. With the frame constraints, our linear harmonic system is then solved by convex optimization. The rotation constraints are kept hard, while the harmonicity constraints

are softened and ultimately enforced by minimization of a Laplacian energy. Additionally, if there are initial feasibility issues with the frame choices, we are able to fix the frames locally, an impossibility with other frame-based methods. We refer to our method with the acronym HGP, for “harmonic global parametrization.”

As a test of HGP, we run it on the benchmark from [Myles et al. 2014]. We compare it to several state-of-the-art methods and see that it is the most robust. Only the method of [Chien et al. 2016b] is comparable in terms of success rate, and our method is 2 to 3 orders of magnitude faster. It usually requires just one iteration of a second-order cone program (SOCP) to find a solution, and is faster than the other methods we compared to (though not by such a large degree). Beyond speed and robustness, we find that HGP also produces maps with relatively low distortion (see Figure 9).

We make a final note on organization of the paper: for readers interested only in the application of the theoretical results, only Section 3.3, Section 4, and Theorem 6.1 need to be read from Sections 3-6. The rest of these sections contain a nearly self-contained exposition leading to the main theoretical result (Theorem 6.1). They include a detailed summary of the index arguments from [Gortler et al. 2006] (Section 3.2) and examples of the branched cover construction and generalized index arguments (Figure 7), which help to build intuition for the main result.




Fig. 2. Proof of concept for $q = 6$. A demonstration of HGP for hexagonal seamless parametrization on two sphere models from the Princeton and AIM@Shape repositories, respectively, with 3 cones each of cone angle $2\pi/3$. This case is of interest and has been considered before [Nieser et al. 2012].

2 PREVIOUS WORK

Parametrization of triangle meshes is a problem with a long history in computer graphics, and there are too many approaches to describe individually. Here, we focus on approaches that are aimed at achieving locally injective seamless parametrizations without remeshing, or that have somehow inspired our method. We refer the reader to some excellent surveys [Floater and Hormann 2005; Hormann et al. 2007] for more complete expositions. As our method has obvious relevance to quad meshing and directional field design, we also recommend the following surveys for more information on these topics: [Bommes et al. 2013; Vaxman et al. 2016].

We begin by noting that the majority of parametrization methods use optimization over spatial variables, which describe the coordinates for the images of the mesh vertices in \mathbb{R}^2 . HGP is amongst

these. To parametrize complicated topologies, the mesh surface must be cut to a mesh disk before mapping; and without additional constraints, artifacts will arise along the cuts.

We use the language of [Myles and Zorin 2012] and describe such parametrizations as *global parametrizations*. Given a cut, Myles and Zorin note a useful correspondence between locally injective maps (up to post-composition by a Euclidean isometry) and discrete metrics on the mesh. In this context, they have a useful discussion on holonomy of the resulting metric, and define maps to be *seamless* if they induce $\pi/2$ holonomy (we of course generalize this definition to include $2\pi/q$ holonomy). This is discussed in greater detail in Section 3.3.

2.1 Harmonic methods

A subfamily of the spatial variable approaches are harmonic methods. The introduction already gave a brief description of these, citing [Gortler et al. 2006; Tong et al. 2006; Tutte 1963], so we just make a few additional points here. First, we also note the work of Floater [Floater 1997], who realized the utility of Tutte’s result, and generalized it to allow for arbitrary convex combination weights in the discrete Laplacian. Since then, much analysis has been done on various choices for such weights, and we refer the reader to [Wardetzky et al. 2007].

As noted, early harmonic methods were limited in scope of applicability, and the work of [Tong et al. 2006] developed a framework allowing for application of such methods to arbitrary topologies, achieving $\pi/2$ holonomy. They ask for user-assisted creation of a singularity graph which decomposes the surface into patches, and solve a harmonic system on each patch individually. Suitable continuity and harmonicity conditions are enforced on the edges between patches, and harmonicity constraints are dropped on the meta-vertices of the singularity graph.

The individual patch maps may be glued together to form a global parametrization, and the system that we solve is nearly equivalent to theirs in the case of a single patch and $\pi/2$ holonomy. In HGP, we generalize to arbitrary rational holonomy, and develop theoretical criteria which ensure local injectivity of the resulting parametrization map. These criteria could also be applied to their approach.

2.1.1 Orbifold methods. Here, we cite the harmonic orbifold-based methods which have partially inspired us: [Aigerman and Lipman 2015, 2016]. The first shows that for sphere and disk meshes with 3 or 4 cone points and specific cone angles, solution of a harmonic system produces a bijective parametrization to a standard domain: a Euclidean orbifold. The limited scope is due to the finite number of Euclidean orbifolds.

HGP is a generalization of the Euclidean orbifold approach to manifolds of arbitrary topology, numbers of cone points, and rational holonomy; along with sufficient conditions for local injectivity in this wider domain of applicability. The fact that these conditions are only applied to cone and boundary triangles was inspired by recent works demonstrating that distortion is maximized on the boundary for planar harmonic maps [Chen and Weber 2015; Chien et al. 2016a; Levi and Weber 2016].




Fig. 3. Results on large models. We show here three large models, from the AIM@Shape database, with approximately one million faces, that HGP succeeded in parametrizing in about 5 minutes. Such models are impossible for slower methods to handle in an efficient manner. The texture allows one to see the quality of these maps.

In [Aigerman and Lipman 2016], the authors again consider mapping to orbifolds with harmonic maps, but this time map to hyperbolic orbifolds. As there are many more hyperbolic orbifolds (an infinite number), this method is much larger in scope. However, the method is primarily useful for surface correspondence, as the seam isometries are hyperbolic isometries, not Euclidean (in particular, they are not Euclidean rotations). Additionally, their cone angles are limited to being $2\pi/q$ for some positive integer q , and may be no greater than π . Lastly, rectangles do not exist in the hyperbolic plane, so there is no square lattice pattern that may be pulled back onto the surface.

2.2 Local injectivity through convexification

Key in our work is the use of convexification in the form of frames, as introduced in [Lipman 2012]. They were also used in [Aigerman et al. 2014], a paper focused on bijections between meshes, but which may also be used for parametrization. Such frames enforce positive orientations of triangles, but must be chosen carefully to ensure feasibility. Our method allows one to only enforce frames on cone and boundary triangles, alleviating some of these feasibility concerns. For our frame choices, we utilize the frame field generated by [Bommes et al. 2009].

Additionally, we note that maps that satisfy frame conditions induce cone angles at interior vertices that are multiples of 2π . Local injectivity only results if these angles are all 2π . HGP allows for a local fix to these problems (see Section 7.1), which is not easily applicable when frames are imposed over the entire mesh.

2.3 Other relevant works

There are several methods that work with variables that determine the metric induced by the parametrization map. The most recent of these is [Chien et al. 2016b], which uses the metric directly in the form of edge lengths squared. Distortion is convex in these variables, but the curvature conditions are not, and the resulting optimization problem must be solved by sequential convex programming. While quite robust, the running time can be quite high, and control

over holonomy for higher genus surfaces requires dense nonlinear constraints.

Many earlier metric-based methods aimed to approximate conformal maps in their parametrizations: [Kharevych et al. 2006; Springborn et al. 2008]. The first uses triangle circumcircle radii as variables, while the second (denoted CETM) uses conformal scale factors at the vertices. These methods tend to lack robustness, as we show with CETM in Section 8. Additionally, limiting to the space of conformal maps is perhaps too restrictive, and harmonic maps are the natural generalization. These methods also do not allow for holonomy control for higher genus surfaces.

A related conformal approach is that of [Gu and Yau 2003] which computes parametrizations of surfaces of arbitrary genus with holomorphic 1-forms. These are just pairs of orthogonal harmonic 1-forms, and similar to the higher-genus approaches of [Gortler et al. 2006], the method usually induces high (isometric) distortion and does not allow for control of cone placement, or cone angles that are not multiples of 2π . The work of [Dong et al. 2005] traces the flow lines of the pair of harmonic 1-forms to produce quad-dominant meshes with user-controlled scaling.

There are many methods which have attacked the problem of seamless parametrization by incorporating cone placement and holonomy determination into their approach. These include: [Ben-Chen et al. 2008; Bommers et al. 2009; Diamanti et al. 2015; Myles and Zorin 2012, 2013]. HGP assumes that the cone points, angles, and holonomy have been given, and aims to achieve these, so we do not compare to these in this paper. Our method may be viewed as being complementary to these methods.

In addition to [Chien et al. 2016b], there have been two other recent works which solve similar problems. [Myles et al. 2014] present a notable method for computing a seamless parametrization guided by a cross field. While the method is fully robust, it does not necessarily preserve the topology of the cross field. Moreover, the produced map is typically not piecewise linear on the original triangulation. [Fu and Liu 2016] use per-triangle affine transformations as variables, striving to make these integrable. A barrier term ensures positive triangle orientation. While the method is quite quick, it is not guaranteed to produce integrable transformations that correspond to a locally injective seamless parametrization with the prescribed cone angles.

Lastly, we must mention QuadCover [Kälberer et al. 2007], and the related HexCover [Nieser et al. 2012] and [Knöppel et al. 2015], which solved for stripe patterns on surfaces. These methods considered q -fold branched covers for $q = 4, 6$, and 2 , respectively, and our branched cover construction is directly inspired by these approaches.

3 PRELIMINARIES

In this section, we summarize some relevant previous works and set some necessary definitions and notation. We begin with a description of Tutte’s classic result [Tutte 1963], and then sketch the proof based on discrete 1-forms given in [Gortler et al. 2006]. In the final subsection, we review the discussion of global parametrization and holonomy from [Myles and Zorin 2012]. For the reader interested primarily in application of HGP, only this last subsection, Section 3.3, need be read.

3.1 “How To Draw A Graph”

In 1963, Tutte proposed a simple method for embedding a planar 3-connected simple graph G in \mathbb{R}^2 as a straight-line drawing [Tutte 1963]. One embeds a peripheral cycle (a “boundary”) as a non-degenerate convex polygon, and replaces the other edges with springs. The resulting equilibrium positions of the interior vertices (those not in the peripheral cycle) and springs gives a valid straight-line embedding of G .

To be more precise, let us set some notation. Let V denote the vertices of G , let B denote the vertices of G in the chosen peripheral cycle (referred to as the boundary from here on), and let $N(v_i)$ denote the neighbors of vertex v_i . Additionally, we denote the positions of vertex v_i by (x_i, y_i) , and the constrained positions of the boundary vertices by (b_i^x, b_i^y) . One solves for the positions of the interior vertices by solving a linear system of the following type:

$$\sum_{v_j \in N(v_i)} w_{ij}(x_i - x_j) = 0 \quad v_i \in V \setminus B \quad (1)$$

$$\sum_{v_j \in N(v_i)} w_{ij}(y_i - y_j) = 0 \quad v_i \in V \setminus B \quad (2)$$

$$\begin{aligned} x_i &= b_i^x & v_i \in B \\ y_i &= b_i^y & v_i \in B \end{aligned}$$

In Tutte’s result, the weights are uniform, $w_{ij} = \frac{1}{|N(v_i)|}$, resulting in an embedding where the position of each interior vertex is the average of the positions of its neighbors (Equations (1), (2)). With this system, Tutte showed that:

THEOREM 3.1 (TUTTE). *If B is embedded as the vertices of a non-degenerate convex polygon in \mathbb{R}^2 , i.e. the polygon has non-zero area and no two vertices are coincident, then the resulting straight-line drawing will be an embedding of G .*

The utility of this linear method for parametrizing mesh discs was recognized in [Floater 1997], and Floater showed that more generally, the weights w_{ij} just need to be positive to ensure the embedding of the graph (and thus the mesh). For such weights, the position of each interior vertex is some general convex combination of the positions of its neighbors. As a result, such maps are called *convex combination maps* (abbreviated as CCMs from here on).

3.1.1 Complex Notation. The linear system above may be more compactly expressed with a complex number $z_i = x_i + iy_i$ describing the position of each vertex v_i :

$$\sum_{v_j \in N(v_i)} w_{ij}(z_i - z_j) = 0 \quad v_i \in V \setminus B \quad (3)$$

$$z_i = b_i \quad v_i \in B \quad (4)$$

In Equation (4) above, $b_i = b_i^x + ib_i^y$ denotes the constrained position of boundary vertex v_i . For brevity, we will utilize such complex notation for the remainder of this paper.

3.2 Proof by Index Argument

As our work extends the results from [Gortler et al. 2006] and generalizes their arguments, we recall the relevant definitions made there, and sketch their basic arguments in this section. Before continuing, we note that the language of discrete differential geometry is used,

and the concepts are described narrowly, only in ways that are relevant to our particular setting. For more complete discussions, the reader is referred to [Crane et al. 2013].

3.2.1 Discrete 1-Forms. We begin with an oriented mesh surface S with vertices V , edges E , and faces F ; and we also require that the graph formed by the vertices and edges is 3-connected. We may consider each edge e_{ij} as consisting of two half-edges h_{ij} (going from v_i to v_j) and h_{ji} (going the other way). Due to the orientability of S , we may consider each face f_k as consisting of a cycle of half-edges consistent with the orientation, and we also get a natural ordering for cycling through edges around vertices. On such an object, we define discrete 1-forms (which we simply refer to as 1-forms, hereafter):

Definition 3.2. A **discrete 1-form** ρ on a mesh surface S is an assignment of a real number ρ_{ij} to each half-edge h_{ij} such that $\rho_{ij} = -\rho_{ji}$.

1-forms may be thought of as discretizations of vector fields. For our purposes, we consider 1-forms that arise from CCMs. Consider the inset example, which illustrates the image of a CCM map (non-uniform weights) of a simple mesh disc and two resulting 1-forms. The 1-forms are $\rho_{ij}^x := \text{Re}(z_j - z_i)$ and $\rho_{ij}^y := \text{Im}(z_j - z_i)$, and values (on the half-edges denoted by the arrows) are shown in blue and red, respectively. These two 1-forms may be thought of as discrete representations of the gradients of the real and imaginary parts of the map.




Fig. 4. An illustration of $\rho^{0.8,0.6}$ for the inset example. The geometric intuition for the value on a particular half-edge is illustrated on the right.

Generally speaking, most 1-forms are not closed, but any of the 1-forms $\rho^{\alpha,\beta}$ defined previously will be closed. This is because they were obtained from an actual function on the vertices (because they are exact).

Given a 1-form ρ , we may apply the boundary operator ∂ (determined by the weights w_{ij}) and obtain a real number $\partial\rho(v)$ for every vertex of the mesh. It is given by the formula:

$$\partial\rho(v_i) := \sum_{v_j \in N(v_i)} w_{ij} \rho_{ji}$$

If $\partial\rho(v_i) = 0$ then we say that ρ is *co-closed* at v_i . If $\partial\rho(v_i) = 0$ for all vertices, we say ρ is *co-closed*. For a 1-form $\rho^{\alpha,\beta}$ from a CCM with weights w_{ij} , we see that it is co-closed at all interior vertices, as the co-closed condition is just a linear combination of Equations (1) and (2). For this reason, we also sometimes describe a 1-form that is co-closed at v_i as being *harmonic* at v_i .

Lastly, we note that a 1-form that is both closed and co-closed is *harmonic*. We do not utilize fully harmonic 1-forms, but 1-forms that are mostly harmonic. For example, note that any $\rho^{\alpha,\beta}$ for our small example is not harmonic at any of the boundary vertices.

The map may be reconstructed from ρ^x and ρ^y , up to translation, by integrating these 1-forms. In particular, if one fixes a vertex v_0 to be mapped to the origin, then the position of any other vertex v_i may be determined by taking any path from v_0 to v_i and summing ρ^x and ρ^y along this path. If considering the inset example above, remember to take the half-edge orientations into account.

More generally, we obtain a 2-dimensional (real) space of 1-forms by taking linear combinations of these 1-forms, $\rho^{\alpha,\beta} := \alpha\rho^x + \beta\rho^y$, for arbitrary $(\alpha, \beta) \in \mathbb{R}^2$. Geometrically, we may see that $\rho_{ij}^{\alpha,\beta}$ is the product of $\|(\alpha, \beta)\|$ and the signed length of the projection of $\rho^x \mathbf{e}_1 + \rho^y \mathbf{e}_2$ (the image of the half-edge h_{ij}) onto the subspace spanned by (α, β) . See Figure 4, where we give an example with $\|(\alpha, \beta)\| = 1$. An analysis of these 1-forms will show that a CCM is always an embedding when the boundary is mapped to a non-degenerate convex polygon.

3.2.2 Closedness and Co-Closedness. Given a 1-form ρ , we may apply the co-boundary operator δ (the discrete exterior derivative) and obtain a real number $\delta\rho(f_k)$ for every face f_k of the mesh:

$$\delta\rho(f_k) := \sum_{h_{ij} \in f_k} \rho_{ij}$$

If $\delta\rho(f_k) = 0$ then we say that ρ is *closed* at f_k . More generally, we say that ρ is *closed* if $\delta\rho(f_k) = 0$ for all faces. For intuition, note that this means ρ may be integrated locally about f_k to obtain a well-defined function, as summation around the half-edges of f_k is zero.

3.2.3 Index of a 1-Form. Before defining index, we make two important notes. The first is that in our index counting arguments, we actually consider 1-forms only on closed (oriented) mesh surfaces. The example and discussions above also apply to a closed mesh of genus 0, which we form by adding another face to the mesh disk whose edges are the boundary edges of the disk. We refer to such a face as an *exterior face*, and continue to use the language of boundary and interior vertices to describe those sets from the original disk mesh. It is the 1-forms ρ^x , ρ^y , and $\rho^{\alpha,\beta}$ on the sphere mesh, which we analyze ultimately. Also, the definition below for index is not meant to apply to true boundary vertices of some non-closed mesh.

The second note is that in the main body of the text, we only analyze 1-forms that are *non-vanishing*: which have nonzero values on all half-edges, and assume that all CCMs do not result in edges of zero length and faces with zero area. These simplify the presentation greatly. However, to be rigorous, we will have to deal with vanishing 1-forms, and the second (related) assumption must be proven, so we defer discussion of these technical details to Appendix B.

For a non-vanishing 1-form, let $\text{scg}_\rho(v)$ or $\text{scg}_\rho(f)$ denote the number of sign changes in ρ_{ij} when the half-edges around v or f are cyclically traversed. For $\text{scg}_\rho(v)$, one may consider either the set of outgoing half-edges or the set of ingoing half-edges. Note

that $\text{scg}_\rho(v)$ and $\text{scg}_\rho(f)$ will necessarily be a non-negative even number.

Definition 3.3. The **index** of ρ about a vertex or face p is:

$$\text{ind}_\rho(p) := \frac{1}{2} \left(2 - \text{scg}_\rho(p) \right).$$

If our 1-form is thought of as being sampled at the midpoint of each edge from a smooth interpolated vector field which is zero at each vertex and in the middle of each face, then $\text{ind}_\rho(p)$ is exactly the index of the smooth vector field at p . As might be expected, we have a theorem analogous to the Poincaré-Hopf index theorem:

THEOREM 3.4 (INDEX FORMULA). A non-vanishing 1-form ρ on a mesh surface S of genus g satisfies the following formula:

$$\sum_{v \in V} \text{ind}(v) + \sum_{f \in F} \text{ind}(f) = \chi(S) = 2 - 2g \quad (5)$$

We refer to Equation (5), as the index formula for the remainder of the paper. A simple and satisfying proof is given in [Gortler et al. 2006]. In addition, there is a connection between closedness and co-closedness and the index at vertices and faces:

LEMMA 3.5. If a non-vanishing 1-form ρ is closed or co-closed at a face or vertex p , respectively, then $\text{ind}_\rho(p) \leq 0$.

This follows simply by noting that if ρ is non-vanishing and closed or co-closed at a vertex or face p , then there must be values of opposite signs (otherwise the relevant sum cannot be zero).

3.2.4 Wheel Vertices & Local Injectivity. In order to achieve an embedding or locally injective map, we would like interior vertices and faces to be well-behaved. In particular, we would like them to be *wheel* and *convex*, as defined in [Gortler et al. 2006]. Given our setting of triangular meshes, and CCMs that produce nonzero area faces, we have automatic convexity of the faces. Thus, we recall only the definition for wheel vertices, and state a key lemma relating the condition to the index of 1-forms $\rho^{\alpha,\beta}$ at the vertices:

Definition 3.6. Let γ_i be the signed angles between adjacent outgoing half-edges in the 1-ring of v (angles are measured by going the “short” way between half-edges, so $0 < |\gamma_i| < \pi$). A vertex v of a CCM is **wheel** if all the γ_i have the same sign and $\sum_i |\gamma_i| = 2\pi$.

LEMMA 3.7. If a vertex v of a CCM is non-wheel, then for some α and β , the 1-form $\rho^{\alpha,\beta}$ is non-vanishing and $\text{ind}_{\rho^{\alpha,\beta}}(v) < 0$.

In particular, Lemma 3.5 tells us that we may show an interior vertex in a CCM to be wheel by showing that the index of $\rho^{\alpha,\beta}$ is zero for all α, β (Lemma 3.5 prohibits positivity). The intuition behind Lemma 3.5 is illustrated and explained in Figure 5. In our setting, local injectivity is equivalent to having all non-cone interior vertices be wheel. This is discussed further in Section 6.

3.2.5 Index Counting. We may now discuss the key index counting argument. Assume the conditions of Theorem 3.1 (with arbitrary convex combination weights). For the resulting CCM, let us see that any non-vanishing $\rho^{\alpha,\beta}$ will have index zero over all interior vertices (and faces, by the appropriate analogue to Lemma 3.7).

We consider the index formula (5). As noted previously, $\rho^{\alpha,\beta}$ is closed and co-closed at all interior vertices and faces. So by Lemma




Fig. 5. A reproduction of half of Figure 3 from [Gortler et al. 2006] (excludes face diagrams). If we consider ρ^y near the illustrated vertex, we see that a sign change occurs each time the line is crossed as we trace the boundary of the 1-ring. Thus, the index is 0 for the wheel vertex, and negative for the non-wheel vertices. More generally, if any vertex is non-wheel, then a similar line may be found and we may consider $\rho^{\alpha,\beta}$ for any nonzero (α, β) perpendicular to the line. For this particular illustration $(\alpha, \beta) = (0, 1)$.

3.5, all of these elements make nonpositive index contributions and we let $s \leq 0$ denote the sum of these contributions. It remains to consider the contribution of the exterior face and boundary vertices. An example schematic (inspired by Figure 2 from [Gortler et al. 2006]) has been inset and shows the convex polygon boundary and the potential index contributions for $\rho^{\alpha,\beta}$ at the boundary vertices. Arrows on the edges denote the half-edges on which $\rho^{\alpha,\beta}$ is positive. As the polygon is convex and $\rho^{\alpha,\beta}$ is non-vanishing, the exterior face is closed and has index 0. All but two boundary vertices have two outgoing boundary half-edges on which $\rho^{\alpha,\beta}$ has opposite signs, so make nonpositive index contribution. The remaining two vertices make contribution no more than 1 as this the maximum possible index value. Let t denote the sum of the exterior face and boundary vertex index contributions.


Thus the index formula gives us: $s + t = \chi(\mathbb{S}^2) = 2$. As $t \leq 2$, we see that $s = 0$ and that indeed all interior vertices are wheel (and faces are convex). This demonstrates local injectivity, and to prove that the result is an embedding (as stated in Theorem 3.1), just a bit of additional work is needed. As it is not immediately relevant to us here, we refer the reader to [Gortler et al. 2006] for these details.

3.2.6 Non-Convex Generalization. Before moving on, we merely state a generalization of Tutte’s theorem (3.1) to non-convex boundaries, from [Gortler et al. 2006], and discuss it informally. The precise definitions of turning angle and reflex boundary vertices are given in Appendix C.

THEOREM 3.8. Consider a 3-connected mesh disc, with its boundary mapped to a simple polygonal curve with positive edge lengths and turning angle 2π . The resulting CCM is an embedding if all the reflex vertices of the boundary are in the convex hull of their neighbors.

The closed polygonal curve bounds a compact domain, and reflex vertices are those for which this domain is locally non-convex. The condition at a reflex vertex v is implied if all interior triangles having v as a vertex are positively oriented. This mirrors the condition in our generalization, Theorem 6.1.

The result is proven by a more sophisticated counting argument, and shows that for a more complicated boundary curve, the exterior



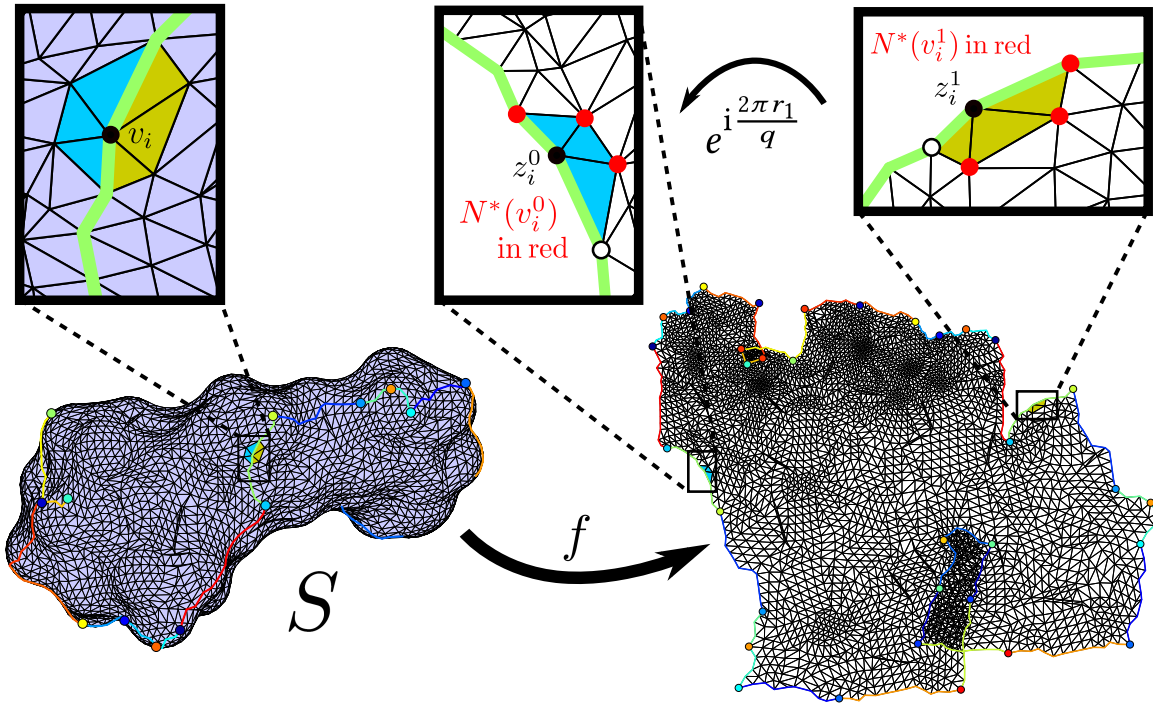


Fig. 6. An example parametrization, and the rotated harmonicity condition. A locally injective parametrization is shown, with the seam edges and cones colored to illustrate duplication after cutting. Insets set notation for Equation (8). In the leftmost inset, we have the 1-ring of a degree 2 (in G_s) seam vertex v_i in S . In the other insets, we have the 1-rings of v_i^0 and v_i^1 in the image $f(S_c)$ (z_i^0 and z_i^1 are $f(v_i^0)$ and $f(v_i^1)$, respectively). The sets $N^*(v_i^0)$ and $N^*(v_i^1)$ are the vertices labelled with red dots. The rotation $e^{2\pi r_1 i/q}$ denotes the rotation constraint for the seam edges neighboring v_i . See text for further commentary.

face and boundary vertices still make a total index contribution of 2 (when the condition on reflex vertices holds). Lastly, we note that arguments like this lead to similar results in [Gortler et al. 2006] for multiply-connected mesh discs, which may be found there.

3.3 Global Parametrization

The above ideas are useful for parametrizing mesh discs, but one also encounters meshes of more complicated topologies in graphics (and in other fields of application). Because of this, we briefly review some concepts from Section 3 of [Myles and Zorin 2012] and set some notation. As noted in the introduction, and to match the discussion in [Myles and Zorin 2012], we assume our mesh S is a triangular mesh. It is straightforward to generalize our theoretical result to polygonal meshes, but the conditions are more complicated to state, and triangular meshes are the most relevant domain by far.

Let g and m denote the genus and number of boundary components of S . Unless $g = 0$ and $m > 0$, topological considerations imply that there is no locally injective map of S to \mathbb{R}^2 . Thus, one searches in these cases for a *global parametrization* of S , which is a selection of cuts along edges transforming S into a mesh disc S_c , and a piecewise linear parametrization map $f : S_c \rightarrow \mathbb{R}^2$. Let E_s denote the edges in the cuts, and V_s denote the vertices at their endpoints. They form a subgraph of the mesh, which we denote with G_s and refer to it as the *seam graph*. For a seam edge e_{ij} , we use e_{ij}^0 and

e_{ij}^1 to denote the two copies of the edge in S_c that are created upon cutting.

If $f(e_{ij}^0)$ and $f(e_{ij}^1)$ are of equal length for all $e_{ij} \in E_s$, then we may pullback the standard metric on \mathbb{R}^2 to obtain a cone metric g_f on S . This metric is flat everywhere except for a discrete set of cone points, which will be a subset of V_s if f is locally injective. If a seam graph is chosen arbitrarily, then the resulting map will likely have high distortion, as the original metric (the restriction of the standard metric on \mathbb{R}^3) is likely to differ greatly from g_f . Thus, it is important to choose G_s well, and to place cone points judiciously.

One may also go the other direction and begin with a cone metric g which is flat on each triangle of S , and produce a global parametrization. In particular, one cuts the mesh S into a disc S_c along a seam graph G_s which includes all the cone points of g and intersects all the boundary components of S . Then, one may lay down the triangles one-by-one following a spanning tree on the dual graph of S_c . Flatness of g at non-cone vertices ensures that there is no dependence on the choice of spanning tree.

The parametrization map f is a *seamless* parametrization if $f(e_{ij}^0)$ and $f(e_{ij}^1)$ may be mapped onto each other by an orientation-preserving isometry of \mathbb{R}^2 for which the rotational part is a rotation by some multiple of $\pi/2$. This is useful for quadrangulation, but we will more generally allow this angle to be any rational multiple of 2π . In particular, suppose all angles are multiples of $2\pi/q$, where q is a positive integer.

Correspondingly, a cone metric g is defined to be *seamless* if the holonomies are multiples of $\pi/2$. The holonomy angles are generated by those for a homology basis of $S \setminus C$, where C denotes the cone points. As a brief reminder, recall that the holonomy of g about a loop in the dual graph is the total angular difference between a vector and a copy of it parallel transported around the loop. Parallel transport is done here with the discrete Levi-Civita connection [Crane et al. 2013]. An alternate and equivalent description is also given in [Myles and Zorin 2012], and we refer readers to these two sources, if further explanation is desired.

As noted in Proposition 2 of [Myles and Zorin 2012], a cone metric is seamless if and only if any global parametrization inducing this metric is seamless. Furthermore, the rotation angles between copies of seam edges are linear functions of the holonomy angles, and may be obtained by solving the system described in Appendix B (of [Myles and Zorin 2012]). However, because rotation angles are indistinguishable up to multiples of 2π , they only determine the holonomy angles up to multiples of 2π . This is a slight clarification of Proposition 2, as stated in [Myles and Zorin 2012], but is quite relevant here as it explains why cone angles and turning angles are verified in our method to check for success (see Section 7.1).

4 THE HARMONIC SYSTEM

Before describing our linear harmonic system, we recall our object to be parametrized: an oriented triangular mesh surface S of arbitrary topology, whose underlying graph of vertices and edges is 3-connected (almost always true in application). The variables in our system are $z_i := f(v_i)$ for the vertices v_i of S_c , the mesh disc resulting from cutting of S . This is well-defined if $v_i \notin V_s$, but for vertices in the seam, cutting causes them to be duplicated a number of times equal to their degree in G_s . Let us denote these copies with superscript indices, e.g., vertex v_i of degree d will have its copies labelled as $v_i^0, v_i^1, \dots, v_i^{d-1}$. Their positions will be denoted by $z_i^0, z_i^1, \dots, z_i^{d-1}$, as might be expected.

The first set of conditions are the rotation constraints for each seam edge e_{ij} , as referenced in Section 3.3. Let the endpoints of e_{ij}^0 be v_i^a and v_j^a and the endpoints of e_{ij}^1 be v_i^b and v_j^b . The constraints are of the form:

$$z_i^a - z_i^b = e^{i \frac{2\pi r_{ij}}{q}} (z_j^b - z_j^a), \quad e_{ij} \in G_s \quad (6)$$

where $r_{ij} \in \{0, 1, \dots, q-1\}$ is used to give the rotation angle, obtained from the desired holonomy angles.

The second set of conditions are the harmonicity conditions for non-seam vertices:

$$\sum_{v_j \in N(v_i)} w_{ij}(z_i - z_j) = 0, \quad v_i \in V \setminus G_s \quad (7)$$

Note that these are the same as Equation (3), and the w_{ij} again denote arbitrary positive convex combination weights.

The third and last set of conditions are the harmonicity conditions on the seam vertices that are not cones. These are quite convoluted to state precisely, but are conceptually simple, so we present here the condition for such seam vertices that are of degree 2 in G_s . This should be most, if not all non-cone seam vertices (and the cut seam may always be modified to ensure all non-cone seam vertices are of

degree 2). Figure 6 establishes notation, and the condition follows for such a vertex v_i :

$$\sum_{v_j \in N^*(v_i^0)} w_{ij}(z_i^0 - z_j) + \sum_{v_j \in N^*(v_i^1)} w_{ij} e^{i \frac{2\pi r_1}{q}} (z_i^1 - z_j) = 0. \quad (8)$$

The basic conceptual idea is easily seen in the figure. The 1-rings of seam vertices have been split in the cutting process and the various sectors have been rotated by known rotation angles. These rotations must be accounted for, and the positional differences (in the sums) must be rotated back to create the proper harmonic condition. The general condition for non-cone seam vertices of higher degree is given in Appendix A as Equation (14).

Our harmonic system consists of Equations (6), (7), and (14). We refer to solutions of this system as *q-convex combination maps* (or q-CCMs, for short). Lastly, we note that our system is related to that of [Tong et al. 2006]: it consists of just one patch, S_c , as opposed to many; and allows for arbitrary rational cone angles and holonomies, instead of just multiples of $\pi/2$.

5 A Q-FOLD BRANCHED COVER

With a q-CCM, we will see in this section that we may construct from it a 2-dimensional space of mostly harmonic 1-forms on a q -fold branched cover. For the reader interested only in the application of HGP, this section may be skipped. No explicit construction of the branched cover (or the 1-forms) is needed for implementation. Lastly, throughout this section and the next, we refer to the examples from Figure 7, as examples A, B, and C.

Given a q-CCM f , consider the q maps denoted f_0, f_1, \dots, f_{q-1} , where $f_n := \left(e^{2\pi ni/q} \right) f$. These are just rotated copies of f . Examples are shown in Figure 7. To construct the branched cover, we will take q corresponding copies of the cut mesh, denoted $S_c^0, S_c^1, \dots, S_c^{q-1}$ and identify seam edges of these cut meshes according to the f_n maps.

In particular, for each seam edge duplicate e_{ij}^0 (there should be $q-1$ of them), we identify it to the other duplicate e_{ij}^1 for which their images under the f_n maps are parallel and their component half-edges are pointing in the same direction. This duplicate is likely on a different copy of S_c , and must exist by the rotation constraints (6) and the fact that we have q rotated copies e_{ij}^1 . Again, Figure 7 provides useful examples, with edge identifications being color-coded.

In these, f happens to be locally injective, and we may think of the images of the f_n as representations of the cut meshes S_c^n . f will not be locally injective in general, but the process will remain the same, as the construction is topological. We will denote the resulting branched cover by \tilde{S}_q , and it is an oriented triangular mesh, as the edge identifications are made in a way that keeps the orientations of the copies of S_c consistent.

The covering map $P_q : \tilde{S}_q \rightarrow S$ is implicit in the construction. Let $Q : \bigcup_{n=0}^{q-1} S_c^n \rightarrow \tilde{S}_q$ denote the quotient map which describes the edge identification described above. For any non-seam vertex or edge, or triangular face in S , its pre-image is the set obtained by applying Q to the q copies of it in $\bigcup_{n=0}^{q-1} S_c^n$. An analogous statement holds for seam edges e_{ij} and vertices v , except there are $2q$ and

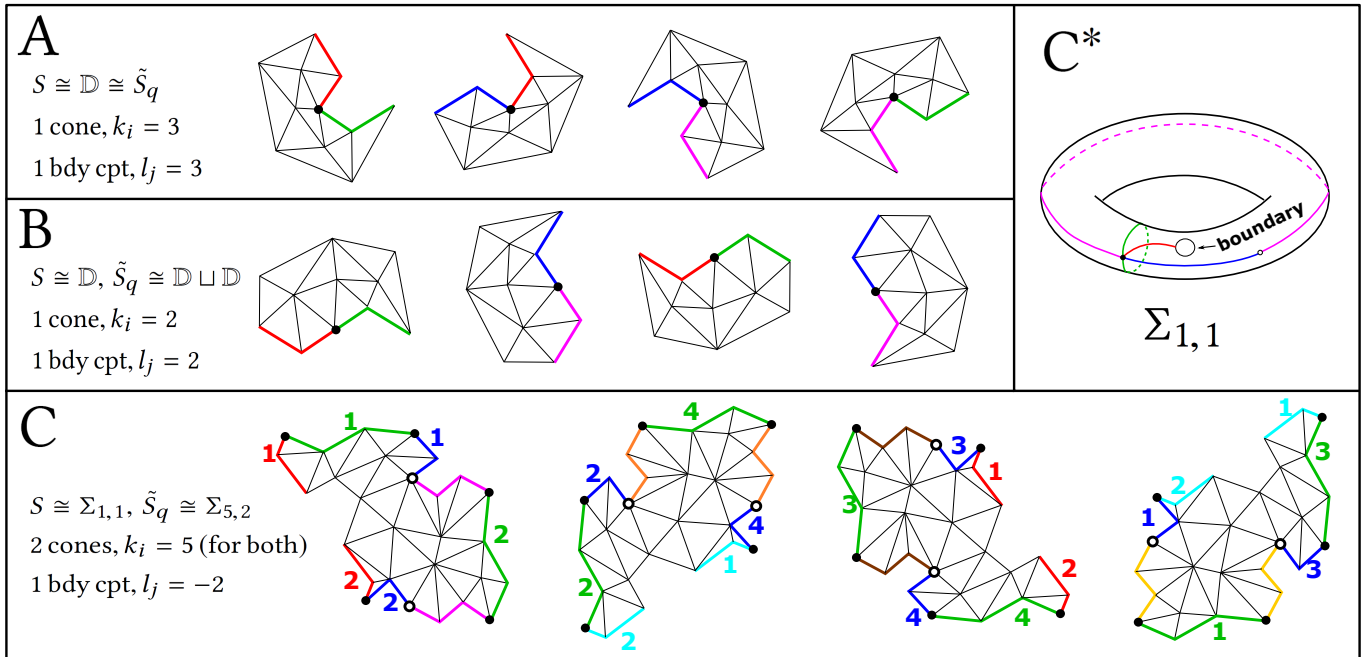


Fig. 7. Branched covering construction examples ($q = 4$). Illustrated are $f_n(S_c)$ for three examples, labelled A, B, and C. These may also be thought of as the mesh copies S_c^n . Edge identifications are color-coded. In example C, 16 identifications are needed, so numbers are used as well, with an identification being made if both color and number match. Next to the examples, there is information about the topology of S and \tilde{S}_q , the cone points and angles of S , and the boundary components of S and their turning angles (given in the notation of Sections 5 and 6). In examples A and B, the mesh S is a disk mesh with a single cone point, of cone angle $3\pi/2$ and π , respectively. In window C, the mesh S is a torus with two cone points (black and white vertices) of cone angles $5\pi/2$ and one boundary component, with turning angle $-\pi$. A topological illustration of S for example C is in window C*. If cut along the colored edges, we obtain a disc that is topologically equivalent to the leftmost mesh in example C. Further discussion of these examples is in Sections 5 and 6.

$q \deg(v)$ copies, respectively, in $\bigcup_{n=0}^{q-1} S_c^n$. Note that the cone points, denoted by C as a set, are the branch points, with other structures having exactly q preimages. Further discussion is in Section 5.2.

5.1 Some 1-forms on \tilde{S}_q

On \tilde{S}_q , we define a 2-dimensional space of 1-forms. As done in Section 3.2.1, we first define two 1-forms: $\tilde{\rho}^x$ and $\tilde{\rho}^y$. For any edge \tilde{e}_{ij} of \tilde{S}_q not resulting from edge identification, it corresponds to an edge e_{ij}^n from some cut mesh S_c^n . If we let v_i^n and v_j^n denote the endpoints of e_{ij}^n , then we define $\tilde{\rho}_{ij}^x := \operatorname{Re}(f_n(v_j^n) - f_n(v_i^n))$ and $\tilde{\rho}_{ij}^y := \operatorname{Im}(f_n(v_j^n) - f_n(v_i^n))$.

For any edge \tilde{e}_{ij} resulting from an edge identification, we just pick one of the identified edges, and define $\tilde{\rho}_{ij}^x$ and $\tilde{\rho}_{ij}^y$ analogously. It will be well-defined, since the identified edges had parallel images that are isometric, so it does not matter which edge we take. Finally, we define $\tilde{\rho}^{\alpha,\beta} := \alpha\tilde{\rho}^x + \beta\tilde{\rho}^y$ for any $(\alpha, \beta) \in \mathbb{R}^2$. These 1-forms will be analyzed to prove Theorem 6.1.


5.2 The topology of \tilde{S}_q

To perform this analysis, we must determine the topology of \tilde{S}_q . This is determined by the Riemann-Hurwitz formula:

THEOREM 5.1. *If \tilde{M} is an N -fold branched cover of a base surface M , and \mathcal{R} denotes the ramification points of \tilde{M} and e_P denotes the ramification index at P , then the following formula holds:*

$$\chi(\tilde{M}) = N\chi(M) - \sum_{P \in \mathcal{R}} (e_P - 1).$$

Recall that a ramification point P is a pre-image of a branch point, and near P the covering map looks like $z \mapsto z^{e_P}$, defining the ramification index. The inset (from Figure 5 of [Kälberer et al. 2007]) illustrates a self-intersecting immersion of the neighborhood of a ramification point of index 4, with points in the branched cover sharing horizontal coordinates if they map to the same point under the covering map. Note that this inset may also serve as a topological picture of \tilde{S}_q for example A (though the geometry induced by the immersion does not match).



It is also quite easy to deduce the topology of \tilde{S}_q for example B: that of two disjoint discs. This example helps to show that \tilde{S}_q need not be connected (though it usually is). For example C, the topology is hard to deduce by eye, so we will use the Riemann-Hurwitz formula.

For this (and to apply the formula generally), we must determine how many ramification points lie above each branch point (cone point) and what their ramification indices are. We shall see that this

depends only on the desired cone angle at the branch point, so let us define an *angle index* k_i at branch (cone) point v_i , so that this desired angle is $2\pi k_i/q$.

We need to consider the 1-ring $R_1(v_i)$ about a cone point v_i and $P_q^{-1}(R_1(v_i))$. Note that $R_1(v_i)$ is cut into sectors by the seam graph G_s , and the q copies of these in $\bigcup_{n=0}^{q-1} S_c^n$ are identified under Q to form $P_q^{-1}(R_1(v_i))$. There is a ramification point for each connected component of $P_q^{-1}(R_1(v_i))$, and the number of copies of $R_1(v_i)$ required to form the component gives the ramification index.

Let us first consider the case where v_i is of degree 1 in G_s , as is true for the cone points in examples A and B. Here, there is only one sector and duplicate of v_i in each copy of S_c , and we may begin at the duplicate of v_i in S_c^0 and move around it in a counterclockwise manner (according to the orientation on S) making identifications. The first identification is to the sector in $S_c^{k_i \pmod{q}}$, by the rotation constraints (6) (derived from the desired cone angle).

Similarly, each subsequent identification will take you to a different copy of S_c where the superscript is incremented by $k_i \pmod{q}$ each time. Once a return has been made to S_c^0 , a connected component of $P_q^{-1}(R_1(v_i))$ has been formed. The formation of the other components is an analogous and parallel process. For intuition, the reader is invited to follow this argument for examples A and B, where $k_i = 3$ and 2 , respectively (and $q = 4$, of course).

We have reduced the problem to one in basic modular arithmetic. Each ramification point has ramification index equal to the order of k_i in $\mathbb{Z}/q\mathbb{Z}$, and the number of ramification points is equal to the index of $\langle k_i \rangle$ (the subgroup generated by k_i) in $\mathbb{Z}/q\mathbb{Z}$. We have argued for the following result in the degree 1 case.

LEMMA 5.2. *The number of ramification points above a branch point v_i with angle index k_i is $\gcd(k_i, q)$. The ramification index at each of these points is $q/\gcd(k_i, q)$.*

For cone points v_i of higher degree in G_s , it is not hard to see that the same equivalence to the modular arithmetic problem holds. The fact that $R_1(v_i)$ is split into multiple sectors just means that additional identifications are needed to create a copy of $R_1(v_i)$. As the identifications are parallel, we still end up incrementing the superscript by $k_i \pmod{q}$ after each copy of $R_1(v_i)$ has been added. Example C provides two examples of these higher degree cases with cones of degrees 2 and 5 (and $k_i = 5$ and $q = 4$ for both).

5.2.1 Boundary components. In order to use the Riemann-Hurwitz formula and prove Theorem 6.1, we must also consider boundary components B_j of S and their preimage $P_q^{-1}(B_j) \subset \tilde{S}_q$. As can be seen in example C, $P_q^{-1}(B_j)$ is not merely q copies of B_j in many cases. Additionally, we must define turning angle and prove some relevant lemmas about index contributions. To not further delay the main index counting argument, we summarize the main lemmas below and refer the reader to Appendix C for further details.

First, we state an analogue of Lemma 5.2 which follows by exactly the same line of reasoning. Under the conditions of Theorem 6.1, let us use l_j to denote the *turning index*, where the turning angle of B_j is $2\pi l_j/q$.

LEMMA 5.3. *The number of connected components of $P_q^{-1}(B_j)$ above a boundary B_j with turning index l_j is $\gcd(l_j, q)$. The turning angle of each of these components is $2\pi l_j/\gcd(l_j, q)$.*

With a way to count boundary components in \tilde{S}_q , we may now apply the Riemann-Hurwitz formula to example C. It has one boundary component with turning index -2 , so \tilde{S}_q has $\gcd(-2, 4) = 2$ boundary components. $\chi(S) = -1$, as a once-punctured torus, and there are two ramification points of index $4/\gcd(5, 4) = 4$. Thus, $\chi(\tilde{S}_q) = -10 = 2 - 2\tilde{g} - 2$, where $\tilde{g} = 5$ denotes the genus of \tilde{S}_q . Thus, \tilde{S}_q is a twice-punctured genus 5 surface in example C.

Secondly, we present a formula for the index contribution of a boundary component for any non-vanishing $\tilde{\rho}^{\alpha, \beta}$, which is the main aim of Appendix C.

LEMMA 5.4. *If a boundary component B_j has turning angle $2\pi\phi$ for some integer ϕ , then the total index contribution (for a non-vanishing $\tilde{\rho}^{\alpha, \beta}$) of its exterior face and vertices is $\phi + 1$.*

Lastly, we again note that our constructed \tilde{S}_q is a generalization of the branched covers described in QuadCover [Kälberer et al. 2007], HexCover [Nieser et al. 2012], and the stripe patterns paper [Knöppel et al. 2015].

6 INDEX COUNTING ON THE BRANCHED COVER

We may now finally state and prove the main theoretical result of the paper. The reader only interested in application of HGP, only needs to understand the statement of Theorem 6.1. We note that analogous to the definition of cone triangles, *boundary triangles* refer to triangles of S containing a boundary vertex (or two). The definition for turning angle is given in Appendix C.

THEOREM 6.1. *Let f denote a q -CCM of S , with specified cone points and holonomy angles determining the rotation constraints. If the cone and boundary triangles are mapped in an orientation-preserving manner, and the induced metric on S achieves the desired cone angles and turning angles, then f is locally injective.*

To show that f is locally injective it suffices to show that any non-vanishing 1-form $\tilde{\rho}^{\alpha, \beta}$ is of index zero at all interior vertices of \tilde{S}_q . This ensures that all interior vertices of $f(S_c)$ are wheel, and that the seam vertices will also be wheel under reconstruction of the 1-ring by isometries of the component sectors. We state first a useful lemma which stems directly from the discrete Gauss-Bonnet equation.

LEMMA 6.2 (GAUSS-BONNET).

$$q|C| - \sum_{v_i \in C} k_i + \sum_{j=1}^m l_j = q(2 - 2g - m).$$

As noted before, we consider the index formula (5). As with the proof of the disc case, we need to consider our 1-forms $\tilde{\rho}^{\alpha, \beta}$ on a closed mesh, so we add exterior faces to each of the boundary components of \tilde{S}_q . The Euler characteristic of the cover goes up by 1 for each exterior face we add. With the previous discussions, we have determined the right-hand side (RHS) of formula (5). Below, we utilize the Riemann-Hurwitz formula and Lemmas 5.2, 5.3, and

6.2.

$$\begin{aligned}
 \text{RHS} &= q(2 - 2g - m) - \sum_{P \in \mathcal{R}} (e_P - 1) + \sum_{j=1}^m \gcd(l_j, q) \\
 &= q(2 - 2g - m) - q|C| + \sum_{v_i \in C} \gcd(k_i, q) + \sum_{j=1}^m \gcd(l_j, q) \quad (9) \\
 &= \sum_{v_i \in C} \gcd(k_i, q) + \sum_{j=1}^m \gcd(l_j, q) - \sum_{v_i \in C} k_i + \sum_{j=1}^m l_j
 \end{aligned}$$

Now we consider the left-hand side (LHS) which counts the indices of the 1-forms $\tilde{\rho}^{\alpha, \beta}$ on \tilde{S}_q . Again, let s denote the contribution of interior vertices (in this case, non-cone and boundary points) and interior faces (all but the exterior faces). All that remains is the contributions of the ramification points and the boundary components, whose neighboring triangles are positively-oriented and achieve the desired cone and turning angles.

For the ramification points, note first that a cone point of angle $2\pi N$ has index contribution $1 - N$ (a simple analogue to Lemma 5.4). Let us group the ramification points into sets $P_q^{-1}(v_i)$ for branch (cone) points v_i . By Lemma 5.2, there are $\gcd(k_i, q)$ ramification points in $P_q^{-1}(v_i)$ all of cone angle $2\pi k_i / \gcd(k_i, q)$. Thus, we have that the index contribution for the ramification points is:

$$\sum_{v_i \in C} \gcd(k_i, q) \left(1 - \frac{k_i}{\gcd(k_i, q)} \right) = \sum_{v_i \in C} \gcd(k_i, q) - k_i.$$

Analogous arguments with Lemmas 5.3 and 5.4 show that the index contribution for the boundary components of \tilde{S}_q is:

$$\sum_{j=1}^m \gcd(l_j, q) + l_j.$$

Thus, we get that:

$$\text{LHS} = s + \sum_{v_i \in C} \gcd(k_i, q) - k_i + \sum_{j=1}^m \gcd(l_j, q) + l_j \quad (10)$$

and we may only have $\text{RHS} = \text{LHS}$ if $s = 0$, indicating that all interior vertices are wheel. QED.

6.1 Example A Calculation

In this brief subsection, we summarize the above argument for example A, to provide some intuition. Afterwards, the reader is invited to do the same with examples B and C with the starting information from Figure 7 (left of examples). With this information, we see that Lemmas 5.2 and 5.3 tell us that for \tilde{S}_q from example A, we find 1 ramification point of ramification index 4, with cone angle 6π ; and 1 boundary component with turning angle 6π . Recalling that $\chi(\mathbb{D}) = 1$, we have Equation (9): $\text{RHS} = 4(1) - 3 + 1 = 2$. The cone angle and turning angle give us Equation (10): $\text{LHS} = s + (-2) + 4$. Thus, we find $s = 0$, as desired.

7 THE METHOD (HGP)

The above generalization of [Gortler et al. 2006] demonstrates that with positive orientation of the cone and boundary triangles, and the correct induced cone angles and turning angles, we have a guarantee on local injectivity of any resulting q-CCM. This inspires the use of

convexification frames [Lipman 2012] for these triangles. However, the choice of frames must be made carefully, in a way which does not exclude the linear space of q-CCMs and lead to infeasibility.

The choice of frames is tied to the desired holonomy, so for the results here, both were determined with the use of the frame field from [Bommes et al. 2009]. Experimentally, we found that for models with low topological complexity (low genus), the use of these frames worked well, and did not lead to feasibility problems. Unfortunately, for higher genus models, these frames were not as robust. Our solution to this problem was to soften the harmonicity constraints (7), (14) while maintaining the rotation constraints (6) and the convexification frames.

To state the optimization problem that we solve, recall that our variables are the $z_i := f(v_i)$ for the vertices v_i of S_c . Let us also set some useful notation to make the representation more compact. Let L denote the Laplacian-like matrix obtained when the harmonicity conditions (7), (14) are gathered into a matrix equation $Lz = 0$ (where z denotes the vector of the z_i 's). Let us also inherit the notation of Section 4 for the rotation constraints (6). Finally, let $f_z^j(z)$ and $\tilde{f}_z^j(z)$ denote the similarity and antisimilarity parts of f_j restricted to a triangular face t_j . These are simple linear functions in terms of the z_i for the vertices of t_j . The set of cone and boundary triangles in S_c will be denoted F_{cb} . We solve the following optimization:

$$\begin{aligned}
 &\text{minimize}_{z_i} && \|Lz\|^2 && (11)
 \end{aligned}$$

$$\begin{aligned}
 &\text{subject to} && z_j^a - z_i^a = e^{i\frac{2\pi r_{ij}}{q}} (z_j^b - z_i^b), && e_{ij} \in G_s && (12)
 \end{aligned}$$

$$\begin{aligned}
 &&& \text{Re} \left(f_z \frac{\tilde{f}_z^j}{\tilde{f}_z^j} \right) - |f_z| \geq \epsilon, && t_j \in F_{cb} && (13)
 \end{aligned}$$

Note first that (12) is just the rotation conditions (6) repeated. Conditions (13) are just the frame conditions, applied to the boundary and cone triangles. ϵ is a small positive constant to enforce positivity. The form of this condition was taken from [Chen and Weber 2015], and $\tilde{f}_z^j / |\tilde{f}_z^j|$ is used to denote the frame for t_j . Lastly, we call the energy in (11) the Laplacian energy, following [Jacobson et al. 2011]. If it is zero, the result will be a q-CCM.

The resulting problem is an SOCP (second-order cone program) with quadratic energy, and the solver we use is Mosek [ApS 2015]. If the resulting parametrization map is not locally injective, we use this result to set new frames and attempt to solve again with new weights, obtained from the edge lengths induced by the non-injective map. This iteration is similar to the frame iterations performed in other frame-based methods [Aigerman et al. 2014; Lipman 2012].

A few additional details: we use mean value coordinates for our harmonicity conditions in this paper [Floater 2003], as they are guaranteed to be positive; and $\epsilon = 0.01$ for the results in this paper.

7.1 Frame Fixing

As noted before, the satisfaction of frame conditions may not ensure the correct cone angle at a vertex, as it merely ensures they are mapped in an orientation-preserving manner. The induced cone angle may differ by multiples of 2π . For Theorem 6.1 to apply, the cone angles must be achieved exactly.

If this occurs (it is rare), we can locally embed the 1-ring and just re-assign frames based on this embedding as we *only* have frames at the cone points. The resulting map is guaranteed to have the correct cone angles if the SOCP is feasible. This follows from an observation of [Aigerman et al. 2014]: that the angle sum at any vertex is a continuous function of the vertex positions. Thus, any other map in the convex space cut out by the frames will have the same angle sums.

For the results in this paper, we implemented a simple algorithm for local embedding. In particular, to determine the interior angles (those based at the cone point) we solved an ABF-like system [Sheffer and de Sturler 2001] which minimized the sum squared angle deviations while achieving the desired cone angle. The geometry of the rest of the 1-ring is determined by fixing the lengths (from the initial incorrect embedding) of the edges with the cone as a vertex. Lastly, note that any other local embedding procedure achieving the correct cone angle may also be used to generate new frames.

For other frame-based methods, there is no equivalent local fix, because there are frames on every single triangle. If one simply modifies the frames about a troublesome point, but does not modify other nearby frames in some systematic manner, feasibility issues are likely to occur. Figure 8 shows the results of such a fix.




Fig. 8. Frame fixing demonstration. This femur model (from the AIM@Shape repository) was parametrized with both HGP and the BD method [Lipman 2012], and did not produce the desired cone angles initially in both cases. We highlight the troublesome region in the blowups, where the desired cone points are denoted with pink dots and the yellow dot denotes a regular vertex. The cone point nearest the regular vertex had cone angle 2π instead of 4π and the regular vertex inherited this angle deficit and became a cone with angle 4π . The top illustrates our result after frame fixing, while the bottom illustrates the BD result, which cannot be fixed in a similar fashion.

8 RESULTS

In this section, we discuss the robustness and running time of our method, HGP, and present several figures that demonstrate its strengths. The models used are from the well-known benchmark data set from [Myles et al. 2014], drawn from the AIM@Shape, Stanford, and Princeton shape repositories.

8.1 Robustness & Running Times

On the benchmark, we compare to several state-of-the-art methods. There were 114 models in the data set, and 8 of these are impossible to parametrize without modifying mesh connectivity (fligree, helmet, raptor, seahorse2, vh_skin, brain, pegaso, robocat). Thus, we omit these and are left with 106 models.

In [Chien et al. 2016b], the same database was used, and we utilize the statistics obtained there for comparison. In addition to the metric-based method suggested there, they also consider implementations of the constrained optimization methods of [Aigerman et al. 2014] and the ARAP- L_∞ approach of [Levi and Zorin 2014], as well as the conformal CETM method [Springborn et al. 2008]. We add to these statistics a comparison to an implementation of [Lipman 2012] (denoted BD), another method which uses a full set of frames on every triangle in the mesh. The LSCM energy [Lévy et al. 2002] was optimized for, and the frames are also obtained from the cross-field calculated in [Bommes et al. 2009]. Only a single SOCP iteration was performed in each application of the method (no frame iteration), as distortion was not a measure of success for the statistics.

Ultimately, we find that HGP succeeded in producing locally injective seamless parametrizations for 104 models, failing on only 2. The resulting successful parametrization maps are included as supplementary material. Of these, 96 required just 1 SOCP iteration, 6 required 2 iterations, and 2 required 3-4 iterations. For the other methods under consideration: [Chien et al. 2016b] succeeded on 102, [Aigerman et al. 2014] succeeded on 97, ARAP- L_∞ succeeded on 93, BD succeeded on 90, and CETM succeeded on just 18.

While the success rate of [Chien et al. 2016b] is nearly identical, HGP is much faster. On larger models in the benchmark, it is 2 or 3 orders of magnitude quicker. For example, on the “armchair” model, with 100K faces, our method took just 10.5 seconds, while the metric-based method took about 40 minutes. HGP uses just one iteration of an SOCP problem for almost all of the models in the benchmark, while [Chien et al. 2016b] requires many. Also, we note that the method of [Chien et al. 2016b] did not implement additional constraints to control for the holonomy angles (though possible). Thus on models with positive genus, the parametrizations obtained there may be merely global and not fully seamless.

We may also consider more closely the comparison to the other frame-based methods: BD and [Aigerman et al. 2014]. As reflected in the results above, it is advantageous that we only need to apply frames on the boundary and cone triangles. This not only leads to greater feasibility, but also a slight increase in speed. HGP tended to be about twice as fast on average (BD took 26 seconds on the armchair model). Finally, it is important to recognize that feasibility of the underlying SOCP does not automatically imply success for such frame-based methods. Feasibility only ensures the angle sums about vertices up to some 2π multiple (see Figure 8). If we run into these problems, we are able to apply a local frame fixing approach, while such local fixes are not possible for BD and [Aigerman et al. 2014]. On our benchmark results, this frame fixing was applied to 4 out of 5 models successfully, the last one being one of our only 2 failures.

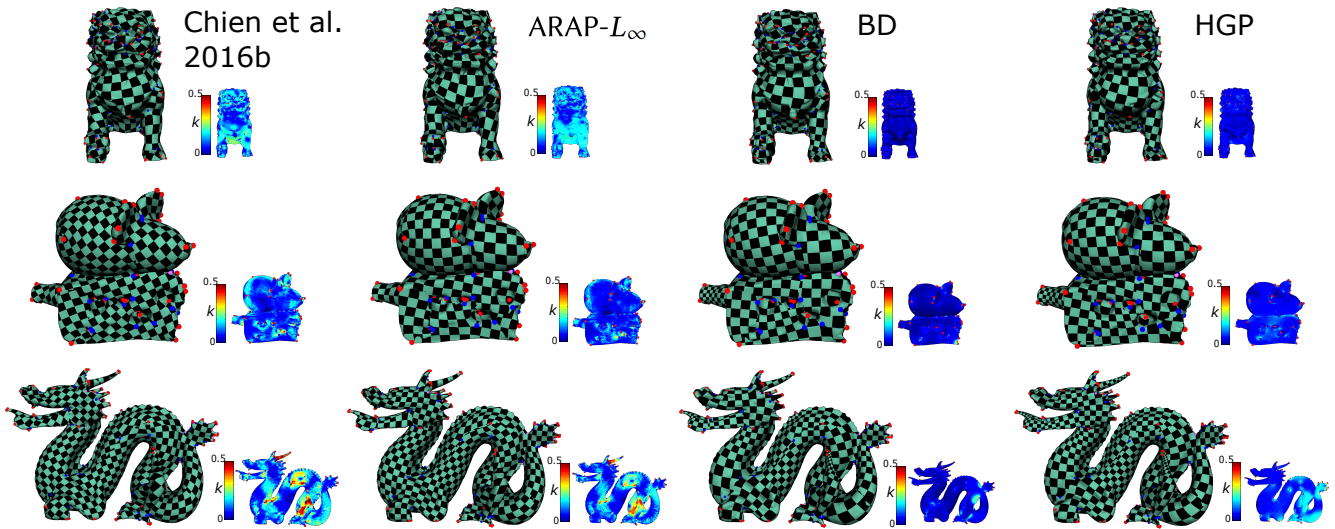


Fig. 9. Map quality comparison on select models from the benchmark [Myles et al. 2014]. First column: Variant V1 of [Chien et al. 2016b]. Second column: ARAP- L_∞ [Levi and Zorin 2014]. Third column: BD [Lipman 2012]. Fourth column: HGP. The heat maps illustrate conformal distortion k , which serves as a measure of map quality. As discussed further in Section 8, HGP produces high-quality maps even without explicit attempts at distortion control.

8.2 Figures

Embedded within the paper, we have several images that illustrate various strengths and features of HGP. First, we have Figure 1. As noted in the caption, it shows four difficult models for which many other methods failed to produce locally injective seamless parametrizations. The method of [Aigerman et al. 2014] fails due to infeasibility of the convex problem cut out by the frames. ARAP- L_∞ produces a parametrization map, but it contains foldovers. CETM takes the original metric and flows out of the region of metrics due to violation of the triangle inequality. The method of [Chien et al. 2016b] managed to produce locally injective maps, but took over an hour to complete on the four models, while HGP took less than a minute. Heat maps are provided in the image for the conformal distortion k , and show that we not only succeed in producing locally injective seamless results, but that the distortion for these models is relatively low (see Figure 9 for some results for other models).

The second image, Figure 2, demonstrates HGP for $q = 6$. In [Nieser et al. 2012], such holonomies were considered for the purposes of hexagonal seamless parametrizations. Thus, we take two models with 3 cones each (cone angles $2\pi/3$) and use our method to create such parametrizations. The result is pleasing and the hexagons agree in orientation across the seam (though not exactly, as we did not place the images of the cone points at Eisenstein integers).

Figure 3 shows the success of HGP in parametrizing three very large models ($\approx 10^6$ faces). The running time for each was approximately 5 minutes. The method of [Chien et al. 2016b] would take significantly longer. This is again due to the fact that we only require one SOCP solve for these models.

Figure 6 is a schematic which illustrates the rotated harmonicity condition, but also serves to give an example parametrization. The model shown is the “retinal” model from the [Myles et al. 2014]

benchmark. Note that the parametrization is locally injective and seamless (seam duplicate edges are labelled with the same color).

Figure 8 shows the successful application of our local frame fixing method. The “femur” model was parametrized with both BD and HGP, with the frames from [Bommes et al. 2009]. The results had incorrect cone angles, with a cone that was supposed to be 4π becoming a regular vertex with cone angle 2π and a nearby regular vertex became a cone vertex of angle 4π . The local frame fix works and produces a valid parametrization.

Last, but not least, we have Figure 9, which shows that HGP produces high quality maps. We show results on a few models alongside the results for a few other methods: variant V1 of [Chien et al. 2016b], ARAP- L_∞ [Levi and Zorin 2014], and BD [Lipman 2012]. As can be seen, our maps have better overall conformal distortion than all methods, except for BD, which is comparable. This is despite the fact that HGP merely minimizes the Laplacian energy without any additional distortion terms. If distortion tends to be poor on any particular model, we may also feed our map to any of these methods, which should improve the result while keeping the map locally injective.

9 CONCLUSIONS & FUTURE WORK

In this paper, we have demonstrated that HGP is a fast and robust method for locally injective seamless parametrization. The desired holonomy angles are allowed to be any rational multiple of $2\pi/q$. The basis for our method is a satisfying generalization of Tutte’s classic “spring embedding theorem” and the work in [Gortler et al. 2006]. The solution to our harmonic system (Section 4) produces a parametrization map, which may be used to construct a 2-dimensional space of mostly harmonic 1-forms on a q -fold branched cover. As done in [Gortler et al. 2006], we utilize an index counting argument on these 1-forms to show our main result (Theorem

6.1). This result states that if the cone and boundary triangles are positively oriented and achieve the correct cone and turning angles, then the map is guaranteed to be locally injective. We solve the linear system by optimizing a Laplacian energy with convexification frames applied only on the cone and boundary triangles. This boosts speed and robustness, as well as allows us to fix frames locally, if cone angles differ from the desired values by some multiple of 2π .

There are many avenues for future work. In Section 7, we noted that with higher genus models it was sometimes difficult to find good frames, and the frames from [Bommes et al. 2009] would sometimes lead to infeasibility. It would be of interest to find a more robust way to choose frames. This might allow us to directly utilize the linear space of q-CCMs, and would allow for more direct distortion control. Lastly, we suspect that with a couple technical arguments, the main result should hold for arbitrary non-rational holonomy (by continuity).

ACKNOWLEDGEMENTS

This research was partially funded by the Israel Science Foundation (grants No. 1869/15 and 2102/15). Additional help was provided by several individuals: Eden Fedida, for assistance with figures; Zohar Levi, for formatting of benchmark input data; and Alec Jacobson, for useful discussions on null spaces of sparse matrices.


REFERENCES

- Noam Aigerman and Yaron Lipman. 2015. Orbifold Tutte Embeddings. *ACM Trans. Graph.* 34, 6, Article 190 (Oct. 2015), 12 pages. DOI: <https://doi.org/10.1145/2816795.2818099>
- Noam Aigerman and Yaron Lipman. 2016. Hyperbolic Orbifold Tutte Embeddings. *ACM Trans. Graph.* 35, 6, Article 217 (Nov. 2016), 14 pages. DOI: <https://doi.org/10.1145/2980179.2982412>
- Noam Aigerman, Roi Poranne, and Yaron Lipman. 2014. Lifted bijections for low distortion surface mappings. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 69.
- MOSEK ApS. 2015. *The MOSEK optimization toolbox for MATLAB manual. Version 7.1 (Revision 28)*. <http://docs.mosek.com/7.1/toolbox/index.html>
- Mirela Ben-Chen, Craig Gotsman, and Guy Bunin. 2008. Conformal flattening by curvature prescription and metric scaling. *Computer Graphics Forum* 27, 2 (2008), 449–458.
- David Bommes, Bruno Lévy, Nico Pietroni, Enrico Puppo, Claudio Silva, Marco Tarini, and Denis Zorin. 2013. Quad-Mesh Generation and Processing: A Survey. *Computer Graphics Forum* 32, 6 (2013), 51–76.
- D. Bommes, H. Zimmer, and L. Kobbelt. 2009. Mixed-integer quadrangulation. *ACM Trans. Graph.* 28, 3 (2009), 77.
- Renjie Chen and Ofir Weber. 2015. Bounded distortion harmonic mappings in the plane. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 73.
- Edward Chien, Renjie Chen, and Ofir Weber. 2016a. Bounded Distortion Harmonic Shape Interpolation. *ACM TOG* 35, 4 (2016).
- Edward Chien, Zohar Levi, and Ofir Weber. 2016b. Bounded Distortion Parametrization in the Space of Metrics. *ACM Trans. Graph.* 35, 6, Article 215 (Nov. 2016), 16 pages. DOI: <https://doi.org/10.1145/2980179.2982426>
- Keenan Crane, Fernando De Goes, Mathieu Desbrun, and Peter Schröder. 2013. Digital Geometry Processing with Discrete Exterior Calculus. In *ACM SIGGRAPH 2013 courses (SIGGRAPH '13)*. ACM, New York, NY, USA, 126.
- Olga Diamanti, Amir Vaxman, Daniele Panozzo, and Olga Sorkine-Hornung. 2015. Integrable PolyVector Fields. *ACM Trans. Graph.* 34, 4, Article 38 (July 2015), 12 pages. DOI: <https://doi.org/10.1145/2766906>
- S. Dong, S. Kircher, and M. Garland. 2005. Harmonic functions for quadrilateral remeshing of arbitrary manifolds. *Computer Aided Geometric Design* 22, 5 (2005), 392 – 423. DOI: <https://doi.org/10.1016/j.cagd.2005.04.004>
- M.S. Floater. 1997. Parametrization and smooth approximation of surface triangulations* 1. *Computer Aided Geometric Design* 14, 3 (1997), 231–250.
- M.S. Floater and K. Hormann. 2005. Surface Parameterization: a Tutorial and Survey. *Advances In Multiresolution For Geometric Modelling* (2005).
- Michael S Floater. 2003. Mean value coordinates. *Computer Aided Geometric Design* 20, 1 (2003), 19–27.
- Xiao-Ming Fu and Yang Liu. 2016. Computing Inversion-free Mappings by Simplex Assembly. *ACM Trans. Graph.* 35, 6, Article 216 (Nov. 2016), 12 pages. DOI: <https://doi.org/10.1145/2980179.2980231>

- 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.
- X. Gu and S.T. Yau. 2003. Global conformal surface parameterization. *Symposium on Geometry Processing* (2003), 127–137.
- K. Hormann, B. Lévy, and A. Sheffer. 2007. Mesh parameterization: Theory and practice. *SIGGRAPH Course Notes* (2007).
- Alec Jacobson, Ilya Baran, Jovan Popović, and Olga Sorkine. 2011. Bounded Biharmonic Weights for Real-time Deformation. *ACM Trans. Graph.* 30, 4, Article 78 (July 2011), 8 pages. DOI: <https://doi.org/10.1145/2010324.1964973>
- F. Kälberer, M. Nieser, and K. Polthier. 2007. QuadCover: Surface Parameterization using Branched Coverings. *Computer Graphics Forum* 26, 3 (2007), 375–384.
- L. Kharevych, B. Springborn, and P. Schröder. 2006. Discrete conformal mappings via circle patterns. *ACM Trans. Graph.* 25, 2 (2006), 412–438.
- Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2015. Stripe Patterns on Surfaces. *ACM Trans. Graph.* 34, 4, Article 39 (July 2015), 11 pages. DOI: <https://doi.org/10.1145/2767000>
- Zohar Levi and Ofir Weber. 2016. On the Convexity and Feasibility of the Bounded Distortion Harmonic Mapping Problem. *ACM TOG* 35, 4 (2016).
- Zohar Levi and Denis Zorin. 2014. Strict minimizers for geometric optimization. *ACM Transactions on Graphics (TOG)* 33, 6 (2014), 185.
- Bruno Lévy, Sylvain Petitjean, Nicolas Ray, and Jérôme Maillot. 2002. Least squares conformal maps for automatic texture atlas generation. *ACM Transactions on Graphics (TOG)* 21, 3 (2002), 362–371.
- Yaron Lipman. 2012. Bounded distortion mapping spaces for triangular meshes. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 108.
- Ashish Myles, Nico Pietroni, and Denis Zorin. 2014. Robust Field-aligned Global Parameterization. *ACM Trans. Graph.* 33, 4, Article 135 (July 2014), 14 pages. DOI: <https://doi.org/10.1145/2601097.2601154>
- Ashish Myles and Denis Zorin. 2012. Global parametrization by incremental flattening. *ACM Trans. Graph.* 31, 4, Article 109 (July 2012), 11 pages. DOI: <https://doi.org/10.1145/2185520.2185605>
- Ashish Myles and Denis Zorin. 2013. Controlled-distortion constrained global parametrization. *ACM Trans. Graph.* 32, 4, Article 105 (July 2013), 14 pages. DOI: <https://doi.org/10.1145/2461912.2461970>
- M. Nieser, J. Palacios, K. Polthier, and E. Zhang. 2012. Hexagonal Global Parameterization of Arbitrary Surfaces. *IEEE Transactions on Visualization and Computer Graphics* 18, 6 (June 2012), 865–878. DOI: <https://doi.org/10.1109/TVCG.2011.118>
- Alla Sheffer and Eric de Sturler. 2001. Parameterization of faceted surfaces for meshing using angle-based flattening. *Engineering with Computers* 17, 3 (2001), 326–337.
- Boris Springborn, Peter Schröder, and Ulrich Pinkall. 2008. Conformal equivalence of triangle meshes. *ACM Transactions on Graphics (TOG)* 27, 3 (2008), 77.
- Y. Tong, P. Alliez, D. Cohen-Steiner, and M. Desbrun. 2006. Designing Quadrangulations with Discrete Harmonic Forms. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing (SGP '06)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 201–210. <http://dl.acm.org/citation.cfm?id=1281957.1281983>
- WT. Tutte. 1963. How to draw a graph. *Proc. London Math. Soc* 13, 3 (1963), 743–768.
- Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommes, Klaus Hildebrandt, and Mirela Ben-Chen. 2016. Directional Field Synthesis, Design, and Processing – State of the Art Report. *Computer Graphics Forum* 35, 2 (2016).
- Max Wardetzky, Saurabh Mathur, Felix Kälberer, and Eitan Grinspun. 2007. Discrete Laplace operators: no free lunch. In *Symposium on Geometry processing*. 33–37.

A HIGHER-DEGREE ROTATED HARMONICITY

To state the general rotated harmonicity condition precisely, we need to set up some notation, and we refer the reader to the inset schematic for added clarity. It illustrates the 1-ring of a degree 3 (in G_s) vertex v_i after cutting (though the following discussion is for any degree d). After cutting, the 1-ring of v_i is split into d sectors by the edges of G_s . These sectors may be cycled through in a fashion that agrees with the orientation on S , and we ensure that the superscript indices labelling the copies of v_i follow this order. Let u_0, u_1, \dots, u_{d-1} denote the neighbors of v_i in G_s and let u_j be the vertex separating sector j from sector $j + 1 \pmod{d}$.



Additionally, let u_j^a and u_j^b denote the copies of u_j obtained after cutting, which are adjacent to v_i^j and $v_i^{j+1 \pmod{d}}$, respectively. Then we may define $N^*(v_i^s) := N(v_i^s) - \left\{ u_{s-1 \pmod{d}}^b \right\}$, where $N(v_i^s)$ is the neighborhood in S_c . Lastly, we define r_s to be a correcting rotation condition which is determined from the rotation conditions on the seam edges including v_i . In particular, $r_0 = 0$ and $r_s = \sum_{j=1}^s \theta_j$ where $2\pi\theta_j/q$ is the rotation condition between the edge from v_i^j to u_{j-1}^b to the edge from v_i^{j-1} to u_{j-1}^a . The general condition is then:

$$\sum_{s=0}^{d-1} \left[\sum_{v_j \in N^*(v_i^s)} w_{ij} e^{i \frac{2\pi r_s}{q}} (z_i^s - z_j) \right] = 0, \quad v_i \in V_s \setminus C. \quad (14)$$

B VANISHING 1-FORMS

In this section, we briefly summarize the basic approach in [Gortler et al. 2006] for dealing with vanishing 1-forms. The details are contained within Appendices A and B in that paper, and the arguments made there generalize easily to our case.

The first thing to note is that we may easily extend the definitions for scg and ind to vanishing 1-forms. If a half-edge is 0, it does not qualify as a sign change. Their basic approach is to note that if a vanishing 1-form from a CCM is encountered, it may be perturbed to a non-vanishing 1-form which does not add any positive index vertices or faces and does not remove any negative index vertices or faces. The standard analysis applies to this perturbed 1-form, and we see that all vertices and faces are well-behaved with respect to the vanishing 1-form.

In Appendix B, they extend these results to show that for a CCM, there will be no singular structures (no edges of zero length, faces of zero area, or angles of 0 or π). In fact, this is the only result we really need. Consider Figure 5 and note that if a vertex is wheel there is an interval of directions that may be used to obtain negative index. As the meshes are finite, there must be a direction within that is not parallel to any of the segments in the parametrization (which is the only way to get vanishing 1-forms with no singular structures).

C WELL-BEHAVED BOUNDARIES

The aim of this section is to establish some necessary definitions of turning angle and reflex vertices and to prove Lemma 5.4. This is a direct generalization of Lemma 4.14 and arguments from Lemma 4.15 in [Gortler et al. 2006], so we are brief with the presentation (and recommend the reader consider the arguments therein first).

LEMMA 5.4. *If a boundary component B_j has turning angle $2\pi\phi$ for some integer ϕ , then the total index contribution (for a non-vanishing 1-form $\tilde{p}^{\alpha,\beta}$) of its exterior face and vertices is $\phi + 1$.*

We are in the setting of Theorem 6.1, so we may assume that all boundary triangles (those having a boundary vertex) are positively oriented. Given such a map, we consider the metric near boundary B_j induced by the map.

Definition C.1. The **turning angle** of B_j is given by the sum of the individual turning angles at each vertex:

$$\sum_{v \in B_j} \pi - \theta_v$$

where θ_v denotes the angle sum at vertex v .

Note that the turning angle at a vertex may now be arbitrarily negative, unlike the definition given in [Gortler et al. 2006]. The following definitions for convex and reflex boundary vertices and extrema correspond to the definitions in [Gortler et al. 2006].

Definition C.2. A boundary vertex is **convex** if $\theta_v \leq \pi$.

Definition C.3. A boundary vertex is **reflex** if $\theta_v > \pi$.

Definition C.4. A boundary vertex is an **extremum** relative to a direction $d = (\alpha, \beta)$ in the plane if the two boundary edges meeting at v both project positively or negatively onto d .

The above are restatements of their definitions in terms of angle sums (though not equivalent due to the possibility that $\theta_v > 2\pi$). With all the boundary triangles positively-oriented, we find that the index of a boundary vertex with respect to a direction $d = (\alpha, \beta)$ is fully determined by whether or not the vertex is an extremum or not, and the angle sum at the vertex.

LEMMA C.5. *A non-extremal boundary vertex has index k if $\theta_v \in (-2k\pi, 2(1-k)\pi)$.*

LEMMA C.6. *An extremal boundary vertex has index k if $\theta_v \in ((-2k+1)\pi, (-2k+3)\pi)$.*

The above lemmas may be obtained with a direct count of sign changes with respect to direction d , noting that the existence of these changes is ensured by the fact that any angle of a triangle is bounded above by π .

For a boundary B_j , let N_k and E_k denote the non-extremal and extremal vertices with index k . The following holds by the same argument used in the proof of Lemma 4.14.

LEMMA C.7. *For a boundary B_j with turning angle $2\pi\phi$, we have:*

$$\begin{aligned} & (|E_1| - |E_0| - 3|E_{-1}| - \dots - (2k+1)|E_{-k}| - \dots) + \\ & (-2|N_{-1}| - 4|N_{-2}| - \dots - 2k|N_{-k}| - \dots) = \phi \end{aligned}$$

In particular, scaling of the boundary perpendicular to $d = (\alpha, \beta)$ gives the above equality in the limit. We may now argue for Lemma 13. We note that the index of the exterior face is:

$$1 - \frac{|E_1| + |E_0| + |E_{-1}| + \dots}{2}$$

When summed with the indices of the boundary vertices, we get the result:

$$\begin{aligned} & 1 - \frac{|E_1| + |E_0| + |E_{-1}| + \dots}{2} \\ & \quad + \text{ind}(N_0) + \text{ind}(N_{-1}) + \text{ind}(N_{-2}) + \dots \\ & \quad + \text{ind}(E_1) + \text{ind}(E_0) + \text{ind}(E_{-1}) + \dots \\ & = 1 - \frac{|E_1| + |E_0| + |E_{-1}| + \dots}{2} + s \\ & \quad + 0 - |N_{-1}| - 2|N_{-2}| + \dots \\ & \quad + |E_1| + 0 - |E_{-1}| + \dots \\ & = 1 + \frac{1}{2} [(|E_1| - |E_0| - 3|E_{-1}| - \dots) + (-2|N_{-1}| - 4|N_{-2}| - \dots)] \\ & = 1 + \phi \end{aligned}$$