

Robust Field-aligned Global Parametrization: Supplement 1, Proofs and Algorithmic Details

Ashish Myles^{1*}

Nico Pietroni^{2†}

Denis Zorin^{1‡}

¹New York University ²ISTI, CNR, Italy

1 Non-existence of quadrangulations

We show why the field examples of Section 3 are not topologically compatible with quadrangulations. The example with boundary was considered in [Myles and Zorin 2013]. We consider the two torus examples.

A quadrangulation Q of a surface naturally defines an associated cross-field $V(Q)$: assume that each quad is mapped to a unit square in (u, v) domain, with consistency on shared edges of quads up to a $k\pi/2$ rotation in the (u, v) domain. Then a cross-field is defined by the gradients of u and v .

A holonomy type of a vector field refers to the values of the turning number of the field on all homotopy classes of loops on the surface; that is, field singularity indices and turning numbers along a basis of noncontractible loops.

The immediate consequence of this definition is that the input cross-field and the one computed from the parametrization as above should have identical singularity indices as well as equal holonomies on topologically equivalent loops.

Torus with a single 3-5 pair. The nonexistence of a quadrangulation with a single 3-5 pair for a torus has been proved in [Barnette et al. 1971]. Remarkably, [Jucovič and Trenkler 1973] shows that for any choice of non-regular vertex valences, constrained by the formula for Euler’s characteristics, and any genus other than 1, a quadrangulation with vertices with these valences exists (if vertices of valence 4 can be added); and for tori, the 3-5 pair is the only exception. We note that this does *not* guarantee that for any vector field there is a quadrangulation consistent with it, as the topology of the field also includes holonomies along non-contractible loops.

A torus with field rotation $\pi/2$ along a loop. If the field on a torus has no singularities, but undergoes $\pi/2$ rotation around a non-contractible loop, there is no quadrangulation with the same field topology. [Kurth 1986] shows that any regular quadrangulation of a torus is obtained from a fundamental parallelogram with vertices at integer points in the plane, with opposite sides glued together. We observe that u and v gradient fields as a result are globally defined on the torus, and turning number of either field is zero along any loop. We conclude that regular quadrangulations are not compatible with the field. As the field has no singularities, no other quadrangulation can be compatible.

Finally, we note that although due to [Jucovič and Trenkler 1973], a choice of singularity indices is rarely an obstacle to quadrangulation existence, a combination of choices of turning numbers and indices is much more likely not to correspond to a quadrangulation.

2 Partition construction

Basic form of the motorcycle graph algorithm. We generate a motorcycle graph G consisting of a set $\{T_i\}$ of traced curves, where T_i is a sequence of points, with adjacent points lying on a common triangle. We represent each point \mathbf{p}_i on the surface in barycentric coordinates within a mesh triangle, and consider different representations of the same point (on edges shared by triangles) to be equivalent. At every step, we maintain a set of active (trace curve, direction) pairs (T_i, \mathbf{d}_i) where \mathbf{d}_i denotes one of the four

cross-field directions at the last point \mathbf{p}_i in T_i . We assume that the operation $(\tilde{\mathbf{p}}_i, \tilde{\mathbf{d}}_i, \text{stop}) \leftarrow \text{tracefield}(\mathbf{p}_i, \mathbf{d}_i, G)$, traces the field, advancing through the next triangle, or to the nearest intersection with already traced lines in G if the triangle contains them. The function returns a new (point, direction) pair, as well as an indicator if the resulting point is an intersection with another line. The basic form of the algorithm is given by the following pseudocode.

```
1: Motorcycle graph tracing
2: Initialize the active set  $A$  with active traces  $T_i$  emanating from all singular points and feature line endpoints paired with all non-feature field directions at these points.
3: Initialize the motorcycle graph  $G$  with all traces in  $A$  and inactive traces corresponding to feature lines.
4: while  $A$  is not empty do
5:   for all traces and directions  $(T_i, \mathbf{d}_i)$  in  $A$  do
6:      $(\tilde{\mathbf{p}}_i, \tilde{\mathbf{d}}_i, \text{stop}) \leftarrow \text{tracefield}(\mathbf{p}_i, \mathbf{d}_i, G)$  where  $(\mathbf{p}_i, \mathbf{d}_i)$  is the last point in  $T_i$ .
7:     Append  $(\tilde{\mathbf{p}}_i, \tilde{\mathbf{d}}_i)$  to  $T_i$ .
8:     if stop then
9:       Remove  $T_i$  from  $A$ .
10:    end if
11:  end for
12: end while
```

Extension of the motorcycle graph algorithm with termination guarantees.

Recall that on a disk topology domain with no singularities, any integral line starts and ends on the boundary.

We construct a cut of the mesh passing through all field singularities and cutting the surface to a disk. On a cut mesh with disk topology, we can separate the cross-field globally into four fields, labeled \mathbf{u} , \mathbf{v} , $-\mathbf{u}$, and $-\mathbf{v}$.

The key observation is that we can trace parallel lines in u and v directions on the cut mesh, to form a sufficiently fine grid, and we can extend the lines of the grid across the cut until they reach other lines forming the grid; in this way, we can obtain a fine partition. A subset of edges of this partition can be added to the motorcycle graph constructed as described above to ensure that every line terminates.

We assign an orientation to the boundary of the cut mesh. The cut is partitioned into *monotone* segments; a segment is monotone if for all points of the segment, the projection of its direction on \mathbf{u} and \mathbf{v} has a constant sign.

Next, we determine all points on the cut for which at least the projection on \mathbf{u} is zero, and we trace integral lines through these points till they intersect with an already traced line or the cut. As a result, any segment of a cut which is a part of a face boundary is monotone in \mathbf{u} , and the cut mesh is split into *slabs* along the u direction. A slab has either two monotone boundary segments on the cut, or one boundary segment that can be split into two.

Similarly, we build slabs in v direction. The intersection of two slabs is either a quad bounded by integral lines, or has a single boundary segment monotonic both in \mathbf{v} and \mathbf{u} . We refer to these parts as *cells* forming a *cell partition* (Figure 1). It is easy to see that the cell partition can be extended to a proper partition: each integral line terminating at the cut is extended to the cell C on the other side; as the field is regular on C , the integral line has to reach the boundary again.

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

†e-mail: nico.pietroni@isti.cnr.it

‡e-mail: dzorin@cs.nyu.edu

Finally, for any traced integral line stopped at the cut, we continue it through the cut into the cell on the other side until it terminates on the cell boundary. We also add the cell boundary to the partition and extend it until it stops at a partition line or on the cut. In the latter case, we add it to the set of unterminated lines.




Figure 1: Cell partition

Proof of Proposition 1. First we prove a Lemma.

Lemma 1. *A connected component of a boundary of any face of a motorcycle graph is either a closed integral line, or is a sequence of integral line segments forming convex angles.*

Proof. If a boundary curve has no corners, it has to be on a single closed integral line. Suppose it has corners. Each corner can be of one of two types, either a starting point, in which case all integral lines from this point were traced, and cannot be in a face interior; or it is a T-joint, with one line stopped by the other; in this case, both corners formed by these lines are convex. \square

Now we proceed with the proof of Proposition 1.

Proof. By the version of the Poincaré-Hopf theorem for N -symmetry fields, the total turning number on all boundary segments should be equal to the Euler characteristic of the face. Because the boundary consists of integral lines, the turning numbers are equal to the sums of the angles of the boundary divided by 2π . As all angles are $\pi/2$, the total turning number of the boundary is equal to $n/4$ and is non-negative, and Euler characteristic is $2 - 2g - b$, where g is the genus and b is the number of boundary components. If the genus is greater than 1, then the Euler characteristic is negative, so this situation is not possible. So the genus is either 0 or 1. If the genus is 1, then the Euler characteristic is non-positive, and is only zero if $b = 0$, and the face has no boundary. This is the case of the torus, with no traced lines (not possible by assumption). Finally, if $g = 0$, then $b = 1$ or $b = 2$. If $b = 1$, then the turning number of the boundary component has to be one, i.e., it has four corners and the face is a quad. If $b = 2$, then the turning number of the boundary is zero, and we have two boundary curves with no corners, i.e. a cylinder.

We note that generically the latter case (cylinder) is only possible in the presence of closed feature curves, as all other curves in the motorcycle graph start at a singularity or an endpoint of a feature curve, and form convex corners with all other integral lines passing through these points. \square

3 Field tracing

Nonuniform field definition. While uniform-rotation fields defined in the paper are adequate for many cases, nonuniform rotation may be needed, e.g., with multiple feature lines at a singularity. Suppose in the chart domain at a vertex, we choose m directions β_i , with assigned field angles ψ_i , $\sum_i \beta_i = 2\pi$, and define $\Delta\psi_i = \psi_{i+1} - \psi_i$, $\Delta\beta_i = \beta_{i+1} - \beta_i$.

Then we define, for an angle ϕ in the chart, the

$$v(\phi) = \psi_j + \Delta\psi_j \frac{\phi - \beta_j}{\Delta\beta_j}$$

if ϕ is between β_{i+1} and β_i .

The map $z \rightarrow z^{\frac{\Theta}{2\pi}}$, maps this angle function to the function on the isometric unfolding of the 1-neighborhood of the vertex given by

$$v'(\phi) = v(\phi) + \phi \left(\frac{\Theta}{2\pi} - 1 \right).$$

The expression in terms of the angle ϕ' on the surface is obtained by substitution $\phi = 2\pi\phi'/\Theta$.

(Almost) uniqueness analysis. Consider a partition of the triangle into curved strips, obtained by tracing the field from each tangency point to intersection with the boundary both ways (we assume that the field is generic and no field line is tangent to the border twice). This partition has the same properties as the linear partition we have constructed. Note that strips are attached to each other against an aligned division, with each strip attached to either 1 or two other strips. This means they form a sequence, and the first and second strip in a sequence form a pair, corresponding to the $T - eA - T - eA - T - oA$ sequence we have described. After the removal of this pair, the same observation applies iteratively, which means that any field on a triangle allows for a decomposition of partition generated by the algorithm. If at each step, the choice of subsequence $T - eA - T - eA - T - oA$ is unique, then the algorithm will necessarily generate the same subsequence as the curved partition of the field.

It remains to show that the choice of the split sequence is unique. First, we observe that all aligned divisions on a triangle edge are either interior or exterior at the same time, because the field rotates in the same direction along the edge. If only one edge has interior divisions, the initial sequence of divisions has only one continuous subsequence of interior divisions, and one of exterior divisions, because each triangle vertex can be only an exterior division. The uniqueness of the sequence of strips generated immediately follows: there is only one way to define a strip near the points where two subsequences meet, and splitting off this strip results in a sequence of divisions with the same property.

Suppose now there are two interior-division edges, and one exterior; if the vertex between two interior-division edges is not a division, then the situation is the same in the case 1.

If this vertex is a division, then there are two separate sequences of exterior aligned divisions: one of length 1, eA_v , consisting of the vertex, and one of arbitrary length on the opposite edge, (eA_1, \dots, eA_k) . Suppose the edge exterior division sequence is of length k , two interior sequences on edges are of length n_1^i and n_2^i , such that $k - 1 = n_1^i + n_2^i$. We observe that chopping off strips at each end of the edge sequence, by eliminating an exterior and interior strip at a time, generates two uniquely defined collections of strips on two sides, until the remaining active face has two sequences of interior divisions of length 1. In addition, there is a single-vertex exterior sequence, and a second exterior sequence consisting of three exterior divisions. At this point, there are two possible ways to partition the remaining domain into strips, and the choice of partition cannot be determined by the boundary information only.

Proof of Proposition 2.

Proof. Consider the sequence of changes from exterior to interior aligned divisions. One can observe the following about the field rotation on the transversal division between two aligned divisions: if there is no change of type (both adjacent aligned divisions are exterior or interior) the rotation on the transversal division is π , or respectively $-\pi$; if there is a change, it is zero. Suppose we have a sequence of n_1^e exterior divisions followed by n_1^i interior, followed by n_2^e exterior and so on, up to n_k^i interior aligned divisions. Suppose the total numbers of interior and exterior divisions are n^e and n^i . Then we have a total turning number of the field along the border of the triangle equal to $(\pi \sum_{j=1}^k (n_j^e - 1) - \pi \sum_{j=1}^k (n_j^i - 1))$, which is equal to $\pi(n^e - n^i)$; as the field has no singularities in the

triangle, we conclude that it should be 2π , i.e. $n^e - n^i = 2$, which proves (a).

It immediately follows that at least two exterior aligned divisions are sequential, and we can pick them next to an interior division, if there is one. If there are no interior divisions, there is a total of two aligned divisions, both exterior, and respectively two transversal divisions, i.e. they form a strip. \square

Conversion of field representations.




Figure 2: Cross-field representations and conversions. (a) The cross-field representation of [Ray et al. 2008; Bommes et al. 2009]. The primary direction of each cross (red) is represented by its angle θ_i from the reference vector \mathbf{r}_i . (b) The traceable representation using two angles per vertex within a facet, with angles across edges related by (1). (c) Labeling of angles for conversion.

Figure 2 compares facet-based cross-field representation of [Ray et al. 2008; Bommes et al. 2009] with the traceable representation described in Section 5. In both cases, cross-field values are represented in each triangle i as angles with respect to a reference vector \mathbf{r}_i , with $\kappa_{ij} = -\kappa_{ji} \in (-\pi, \pi]$ denoting the angle of rotation between adjacent reference vectors \mathbf{r}_i and \mathbf{r}_j . An integer matching variable $m_{ij} = -m_{ji}$ accounts for the $\pi/2$ ambiguity between adjacent triangles so that $\theta_i + \kappa_{ij} + \frac{\pi}{2}m_{ij}$ is the representation of the cross-field from facet i parallel-transported to the representation in triangle j . Thus, in the traceable representation,

$$\alpha_j = \beta_i + \kappa_{ij} + m_{ij} \frac{\pi}{2} \quad (1)$$

since the representation is continuous along edges. Recall that our traceable field representation rotates at a constant rate around the vertex, so that with angles around an interior vertex labeled as in Figure 2c,

$$\beta_i - \alpha_i = \frac{2\pi I - \Theta}{2\pi - \Theta} \Theta_i, \quad (2)$$

where I is the index of the field at the vertex, Θ_i is the vertex angle in triangle i , and Θ is the total angle at the vertex. These constraints leave only one degree of freedom around the vertex, which can be computed from the angles θ_i of an existing facet-based field by locally minimizing the quadratic energy $\sum_{i=1}^n (\frac{1}{2}(\alpha_i + \beta_i) - \theta_i)^2$.

For vertices along feature curves (including border vertices), the field is computed per sector delineated by adjacent feature curves around the vertex. Here, Θ in (2) denotes the total angle of the sector and ΘI is the total angle of rotation of the field around the sector. No minimization is required since one of the angles is fixed to be aligned to the feature curve.

For the converted field to meet the requirements for motorcycle graph tracing, the facet-based field must meet the following conditions.

1. All boundary curves must be feature curves.
2. The cross within a facet must be aligned to all adjacent feature edges adjacent to the facet.
3. All singularities must have field index $I < 1$.

4. The field's angle of rotation within a sector must be a multiple of $\pi/2$ with $I < 1$.

Conditions (3) and (4) ensure that no elliptic or parabolic neighborhoods are forced around a singularity or within a sector.

Techniques such as [Bommes et al. 2009] fix the field adjacent to feature curves and minimize a quadratic energy with integer variables to generate a smooth field. Such algorithms can be adapted to satisfy condition (2) by first splitting triangles adjacent to two non-orthogonal feature curves, and to satisfy condition (4) by constraining the field and matchings around sector corners formed by close-to-parallel feature curves to have a turning angle of $\pi/2$.

4 Partition simplification

4.1 Parametric edge lengths for T-mesh quads

Section 7 requires the computation of parametric edge lengths that satisfy consistency conditions of the form (3) from the paper for each pair of opposite edges of a quad. Denoting parametric length of all mesh edges b_k , $k = 1 \dots n$, we write (3) from the paper in the form $S\mathbf{b} = 0$, where matrix S has entries ± 1 .

We target parametric lengths proportional to surface edge lengths ℓ_k by minimizing the quadratic energy $E_{\text{len}} = \sum_{k=1}^n (b_k - s\ell_k)^2$, where $s > 0$ allows for uniform scaling, and $b_k \geq 0$. As $b_k = 0$ satisfies the constraints, this quadratic program always has a solution, but to generate a bijective parametrization, we require $b_k \geq 1$ except where the constraints force zero length.

Given a vector \mathbf{b}^{\min} of minimum lengths and $\mathbf{w} \in \{0, 1\}^n$, the procedure $\mathbf{b} \leftarrow \text{qpSolve}(\mathbf{w}, \ell, \mathbf{b}^{\min}, S)$ computes a minimizer $\arg\min_{\mathbf{b}} \sum_{k=1}^n w_k (b_k - s\ell_k)^2$ under the constraints $s > 0$, $\mathbf{b} \geq \mathbf{b}^{\min}$ (element-wise), and $S\mathbf{b} = 0$, or fails if there is no solution.

Parametric lengths \mathbf{b} are computed from given ℓ and an $m \times n$ constraint matrix S using the following function.

- 1: **function** $\mathbf{b} \leftarrow \text{computeParamLengthCompatible}(\ell, S)$
- 2: $\mathbf{b} \leftarrow \text{qpSolve}(\mathbf{1}, \ell, \mathbf{1}, S)$
- 3: **if** qpSolve failed **then**
- 4: $\mathbf{z} \leftarrow \text{whichForcedToZero}(S)$
- 5: $\mathbf{b}^{\min} \leftarrow n$ -vector with $b_{i \in \mathbf{z}}^{\min} = 0$ and $b_{i \notin \mathbf{z}}^{\min} = 1$
- 6: $\mathbf{b} \leftarrow \text{qpSolve}(\mathbf{1}, \ell, \mathbf{b}^{\min}, S)$
- 7: **end if**

{Non-zero lengths are iteratively removed from the energy to force all lengths not strictly forced to zero to become positive.}

- 8: **function** $\mathbf{z} \leftarrow \text{whichForcedToZero}(S)$
- 9: $\ell \leftarrow \mathbf{1}$, $\mathbf{b}^{\min} \leftarrow \mathbf{0}$, $\mathbf{b}^{\text{last}} \leftarrow \mathbf{0}$ (all n -vectors)
- 10: $\mathbf{b} \leftarrow \text{qpSolve}(\ell, \ell, \mathbf{b}^{\min}, S)$
- 11: **while** \mathbf{b} has more non-zero entries than \mathbf{b}^{last} **do**
- 12: **for all** non-zero entries i of \mathbf{b} **do**
- 13: $\ell_i \leftarrow 0$ and $b_i^{\min} \leftarrow 1$
- 14: **end for**
- 15: $\mathbf{b}^{\text{last}} \leftarrow \mathbf{b}$
- 16: $\mathbf{b} \leftarrow \text{qpSolve}(\ell, \ell, \mathbf{b}^{\min}, S)$
- 17: **end while**
- 18: $\mathbf{z} \leftarrow$ indices of zero entries in \mathbf{b}

Proofs of Proposition 3 and 4. We start with several basic properties of zero chains.

Property 1. Side quads S_j , $i = 0, 1$ of a maximal zero chain have a T-joint at one of the endpoints of e_0 or e_n . If this were not true, the chain could be extended across the side quad, which should have a zero edge opposite e_0 or e_n and would not be maximal.

Property 2. e_0 , e_n of a zero chain cannot have common edges. Suppose these edges share a common edge of length x and the sum of other edge lengths in e_0 and e_n is x_0 and x_n respectively. Then the only constraint on x is $x + 0 + x = x_n + x$, which allows x

to have any value, i.e. the chain is not forced by constraints to be a zero chain.

Property 3. A chain is admissible if its set of quads does not contain a zero loop and there are no valence 1 vertices. Indeed, if Q_i is the same as Q_{i+1} , unless it is a circular loop (which cannot be a zero chain), two adjacent edges of Q_i has to be identified, i.e. there should be a valence 1 vertex.

Proof of Proposition 3.

Proof. Suppose there is a zero chain; then it contains a side edge with T-joints; w.l.o.g. denote it by e_n . If there is a single T-joint, one of the subedges is an isolated zero edge. If there is more than one T-joint, there is a zero chain with e_0 with no T-joints, that starts on e_n : just pick the edge between two adjacent T-joints. Similarly, at the end of that chain, if there are T-joints, we can find another zero chain with e_0 with no T-joints. Continuing this process, by the finiteness of the number of edges, we either find a simple zero chain, or reach e_0 . Let x_i denote the parametric lengths of the shared zero edges between chains. Then the constraint equations for each chain have the form $x_i + y_i = x_{i+1} + z_i$ where $i + 1$ is modulo the number of chains. Then clearly if we add the same positive constant to all x_i , the equations are still satisfied, so the chain is not forced to be a zero chain. \square

Conditions required for an edge collapse to be possible. The edge collapse involving a triangle as one of the faces can be represented as a composition of a face(facet) join of Q_i and Q'_i across g_i with a vertex-join, using CGAL terminology. For the face join to be possible, Q'_i and Q_i should be distinct (true by assumption on the chain), and the valence of the endpoints of g_i (v_{i-1} and w_i) should be no less than 3. For the vertex-join to be possible, its endpoints v_i, w_i should be distinct, and faces on two sides (join of Q'_i and Q_i , and Q_{i+1}) should have at least 3 vertices.

Proof of Proposition 4.

Proof. Suppose (inductive assumption) before the face-join/vertex-join pair is applied, all T-joints have extended valence 3 or higher, and any vertex on the extent side of the chain has actual valence 3 or higher (because the actual valence of a vertex on the side had to be at least two, and one extension was added). Any cone on the quad chain has valence 3 or higher by assumption; finally, we assume that v_{i-1} has extended valence 4 or higher. At the beginning of a collapse step, Q_i is a triangle. Q'_i and Q_i have to be distinct, otherwise there is a zero loop contained in the zero chain. By assumption on valences of vertices v_{i-1} and w_i , the face join is possible. Because Q_i is a triangle and Q'_i is a quad, the merged face Q'_i has at least five edges and Q_{i+1} has four, so the edge between them can be collapsed, unless v_i, w_i coincide. If these vertices coincide, then there was a (possibly zero) sequence of quad collapses that brought them together; then the edge e_i and a part of this sequence form a zero loop in the zero chain. The complete operation affects valences of two vertices, v_{i-1} and v_i . The valence of v_{i-1} is decreased by 1, so as we assumed that its extended valence is 4, after the operation it is at least 3. Suppose w_i and v_i both have valence 3, before the complete operation. Both cannot be cones, as otherwise there is a zero path connecting them. If both of these are regular T-joints, then their extended valence is 4, and the edge collapse results in valence 3 (extended valence 4) vertex. If one of them is a valence 3 cone, and the other is a regular T-joint, its extended valence is 4; and after collapse, the resulting valence is still 3 (extended valence 4). If at least one vertex has valence greater than 3, the merged vertex is guaranteed to have valence 3 or higher. We conclude that the invariant of the inductive assumption is preserved.

It remains to handle the start and the end of the iteration. In the beginning, a single side edge collapse is used, collapsing v_0 and w_0 and converting Q_0 to a triangle. As one of these vertices is a T-joint, with stem along the extent direction, the resulting vertex has

valence at least 4. At termination, we reach a state with no Q_{i+1} , and similarly the collapse of e_n yields a vertex of valence at least 4.

As the face-join and vertex-join preserve mesh validity, it remains to verify that each face has four corners. This follows from the fact that the number of corners of quads outside the chain never changes by construction, and quads in the chain are completely eliminated at the end. \square

Cone insertion algorithm

The presence of zero parametric lengths in \mathbf{b} triggers the T-mesh simplification procedure described in Section 7. If zero lengths remain at the end of the procedure, the following function is used to compute strictly-positive parametric lengths, allowing for incompatible length assignments only where necessary. Section 7 then proceeds by inserting 2-6 cone pairs to resolve the incompatible quads. We used MATLAB notation $S(\mathbf{r}, :)$ for a submatrix consisting of a subset of rows of a matrix S .

```

1: function  $\mathbf{b} \leftarrow \text{computeParamLengthIncompatible}(\ell, S)$ 
2:  $\mathbf{b} \leftarrow \text{qpSolve}(\mathbf{1}, \ell, \mathbf{1}, S)$ 
3: if  $\text{qpSolve}$  failed then
4:    $\mathbf{z} \leftarrow \text{whichForcedToZero}(S)$ 
5:    $\mathbf{r} \leftarrow \text{maximalAllowableRows}(S, \mathbf{z})$ 
6:    $\tilde{S} \leftarrow S(\mathbf{r}, :)$  (sub-rows corresponding to indices  $\mathbf{r}$ )
7:    $\mathbf{b} \leftarrow \text{qpSolve}(\mathbf{1}, \ell, \mathbf{1}, \tilde{S})$ 
8: end if

{For each length forced to zero, one or both of the constraints
involving it are removed to allow a positive length.}
9: function  $\mathbf{r} \leftarrow \text{maximalAllowableRows}(S, \mathbf{z})$ 
10:  $\mathbf{r} \leftarrow \{1, \dots, m\}$ 
11:  $\mathbf{z}_{\text{left}} \leftarrow \mathbf{z}$ 
12: while  $\mathbf{z}_{\text{left}}$  is not empty do
13:    $k \leftarrow$  an index (of a zero-valued length) from  $\mathbf{z}_{\text{left}}$ 
14:    $\ell \leftarrow n$ -vector with  $\ell_k = 1$  and  $\ell_{i \neq k} = 0$ 
15:    $\mathbf{r}_k \leftarrow$  {rows of  $S$  with non-zero entry in  $k$ }
16:   for all  $\mathbf{r}_{\text{remove}} \subset \mathbf{r}_k$ , in increasing order of cardinality do
17:      $\tilde{S} \leftarrow S(\mathbf{r} \setminus \mathbf{r}_{\text{remove}}, :)$  (sub-rows  $\mathbf{r} \setminus \mathbf{r}_{\text{remove}}$ )
18:      $\mathbf{b} \leftarrow \text{qpSolve}(\mathbf{1}, \ell, \mathbf{0}, \tilde{S})$ 
19:     if  $\mathbf{b}_k > 0$  then
20:        $\mathbf{r} \leftarrow \mathbf{r} \setminus \mathbf{r}_{\text{remove}}$ 
21:        $\mathbf{z}_{\text{left}} \leftarrow \mathbf{z}_{\text{left}} \setminus$  {all non-zero indices in  $\mathbf{b}$ }
22:       Break out of for loop
23:     end if
24:   end for
25: end while

```

References

- BARNETTE, D., JUCOVIC, E., AND TRENKLER, M. 1971. Toroidal maps with prescribed types of vertices and faces. *Mathematika* 18, 01, 82–90.
- BOMMES, D., ZIMMER, H., AND KOBELT, L. 2009. Mixed-integer quadrangulation. *ACM Trans. Graph.* 28, 3, 77.
- JUCOVIĆ, E., AND TRENKLER, M. 1973. A theorem on the structure of cell-decompositions of orientable 2-manifolds. *Mathematika* 20, 01, 63–82.
- KURTH, W. 1986. Enumeration of platonic maps on the torus. *Discrete mathematics* 61, 1, 71–83.
- MYLES, A., AND ZORIN, D. 2013. Controlled-distortion constrained global parametrization. *ACM Transactions on Graphics (TOG)* 32, 4, 105.
- RAY, N., VALLET, B., LI, W., AND LÉVY, B. 2008. N-Symmetry direction field design. *ACM Trans. Graph.* 27, 2.