cagd940

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 [DTD5] P.1(1-30) by:violeta p. 1



Available online at www.sciencedirect.com



Computer Aided Geometric Design ••• (••••) •••-•••



www.elsevier.com/locate/cagd

Discrete one-forms on meshes and applications to 3D mesh parameterization

Steven J. Gortler^a, Craig Gotsman^{a,*}, Dylan Thurston^b

^a Computer Science, Harvard University, 33 Oxford St., Cambridge, MA, USA ^b Mathematics Department, Harvard University, 1 Oxford St., Cambridge, MA, USA

Received 4 June 2004; received in revised form 9 May 2005; accepted 9 May 2005

Abstract

We describe how some simple properties of discrete one-forms directly relate to some old and new results concerning the parameterization of 3D mesh data. Our first result is an easy proof of Tutte's celebrated "spring-embedding" theorem for planar graphs, which is widely used for parameterizing meshes with the topology of a disk as a planar embedding with a convex boundary. Our second result generalizes the first, dealing with the case where the mesh contains multiple boundaries, which are free to be non-convex in the embedding. We characterize when it is still possible to achieve an embedding, despite these boundaries being non-convex. The third result is an analogous embedding theorem for meshes with genus 1 (topologically equivalent to the torus). Applications of these results to the parameterization of meshes with disk and toroidal topologies are demonstrated. Extensions to higher genus meshes are discussed.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Computer graphics; Parameterization; Embedding; One-form; Manifold mesh

1. Introduction

In 1963, Tutte (1963) proved his celebrated "spring embedding" theorem for planar graphs. This theorem maintains that a 3-connected planar graph may be easily drawn in the plane by embedding the

* Corresponding author.

E-mail address: gotsman@eecs.harvard.edu (C. Gotsman).

^{0167-8396/\$ –} see front matter @ 2005 Elsevier B.V. All rights reserved. doi:10.1016/j.cagd.2005.05.002

S0167-8396(05)00053-1/FLA AID:940 Vol. •••(••• COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

2

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••

graph boundary as a strictly convex polygon and solving a linear system for each of the two coordinates of the interior vertices. The linear system forces each interior vertex to lie at the centroid of its neighbors. Tutte proved that the result is indeed a straight line plane drawing, and, furthermore, the faces are non-degenerate, bounding convex regions in the plane.

Tutte's simple procedure remains a popular planar graph drawing method to date. It was generalized by Floater (1997) who showed that the theorem still holds when the boundary is not strictly convex (i.e., adjacent boundary vertices may be collinear), and when each interior vertex is positioned at a general convex combination of its neighbors coordinates. The method has established itself as the method of choice for parameterizing a three-dimensional mesh with the topology of a disk to the plane in geometric modeling and computer graphics, along with a multitude of variations on this theme (e.g. (Colin de Verdière et al., 2003; Desbrun et al., 2002; Kanai et al., 2002)). The main reason for the method's popularity is that it is computationally simple, and also guarantees an injective parameterization homeomorphic to a disk, meaning that the individual planar polygons are convex and do not intersect. The latter is crucial for the correctness of many algorithms relying on an underlying parameterization. As such, Tutte's theorem is the basis for solutions to other computer graphics problems, such as morphing (e.g. (Floater and Gotsman, 1999; Gotsman and Surazhsky, 2001; Kanai et al., 2002)). Many recipes exist for the convex combination weights in order to achieve various effects in the parameterization. Typically, it is desirable to reflect the geometry of the original 3D mesh in the 2D parameterization, so the 2D version should be a minimally distorted 2D version of the 3D original. Depending on how distortion is measured, different weights are used. For more details, see the recent survey by Floater and Hormann (2004).

Inspired by recent work on the theory of discrete *one-forms* (Benjamini and Lovasz, 2003; Gu and Yau, 2003; Mercat, 2001) and their use in mesh parameterization (Gu and Yau, 2003) as well as related results in vector field visualization (Polthier and Preuss, 2003; Tong et al., 2003), we show how some properties of these one-forms on meshes can be used to prove the injectivity of a number of mesh parameterization algorithms.

In a nutshell, a one-form is essentially an assignment of a value to each edge of the mesh. A simple counting argument then produces an Index theorem for these discrete one-forms; this is a discrete analog of the Poincaré–Hopf index theorem for smooth vector fields on surfaces. Imposing additional balancing conditions on the one-forms results in a linear subspace which is shown to be related to spring embeddings of the type described by Tutte.

One-forms on meshes turn out to be a very useful tool for mesh processing. In particular, the central result of Tutte's planar embedding theorem, when formulated in terms of the vector valued differences along *edges* of the graph, follows as a special case of the Index theorem with no more than simple counting arguments and elementary geometry. The techniques used in our proof are considerably simpler than those used in proofs of different versions of Tutte's theorem which evolved over the years (Becker and Hotz, 1987; Colin de Verdière et al., 2003; Floater, 2003a; Floater and Pham-Trong, 2005; Richter-Gebert, 1996; Thomassen, 2004). Moreover, the same arguments allow us to relax the conditions on the embedding of the mesh boundary, and even allow multiple boundaries. We show that it is sufficient that the embedding is well behaved (in a manner to be made precise later) at the vertices along the boundaries, even if they are non-convex, in order that the entire embedding be well-behaved. Since the requirement of a (pre-determined) convex boundary is the (only) major drawback of Tutte's method, this result could make Tutte's method even more popular than it already is. It introduces extra degrees of freedom into the solution, which may be used to produce less distorted parameterizations.

S0167-8396(05)00053-1/FLA AID:940 Vol. •••(••• COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

3

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••

While variants of Tutte's theorem for meshes with the topology of a disk (namely genus 0 with one boundary) are easy consequences of our Index theorem, novel and more interesting results may be obtained for meshes with higher genus. Particularly important results may be obtained for the toroidal (genus g = 1) case. Since, due to the different topologies, it is impossible to map the torus homeomorphically to the plane without cutting it, the most we can hope for is a parameterization method which has this behavior locally. Gu and Yau (2003) showed how to generate local parameterizations with socalled "conformal" structure. In the torus case, their parameterizations have the following properties: (1) Any connected submesh having the topology of the disk is mapped to a disk in the plane. (2) Any two connected submeshes with non-empty intersection, all (the two submeshes and their intersection) having the topology of the disk, are mapped to disks in the plane, such that the two parameterizations coincide, up to a translation of the plane, on their intersection. Gu and Yau do not prove that the resulting mappings of the disks to the plane are actually embeddings. Additionally Gu and Yau restricted their attention to a specific subset of the possible parameterizations, to triangular meshes. We apply our oneform theory to close this gap, providing a generalization of their basic algorithm, and prove that all these parameterizations are locally injective. This may be considered a "Tutte-like" embedding theorem for the torus.

For higher genus (g > 1) meshes, the situation is more complex. Any 2D parameterization must contain 2g - 2 "double wheels", which are neighborhoods of vertices whose faces wind twice around the vertex, so the parameterization cannot be locally injective at these vertices. While we do not yet fully understand this case, we are able to provide some mathematical and algorithmic insight into how to control and generate such parameterizations.

Beyond the theoretical interest, seamless local parameterization of higher genus meshes is useful for applications such as "cut and paste" operations (Biermann et al., 2002), texture mapping (Kraevoy et al., 2003) and meshing (Tewari et al., 2005).

2. Related work

The concept of a discrete one-form is identical to a one-cochain from simplicial cohomology (Hatcher, 2002). In fact, de Rham (1931) original work on the topology of differential forms effectively defined cochains as a way of discretizing differential forms. Whitney (1957) showed how co-chains could be interpreted as continuous forms. We will use the term one-form to emphasize this connection, as has been done before by others (Bossavit, 1999; Mercat, 2001; Hirani, 2003; Gu and Yau, 2003).

In the finite-element community, methods to discretize physical equations (Nedelec, 1980; Raviart and Thomas, 1977) led to discrete equations that manipulate discrete differential forms (Kotiuga, 1984; Bossavit, 1999; Hiptmair, 2002; Sen et al., 2000). This approach involves a simplicial complex, as well as its dual, which is used to define the various relevant operators, such as the *Hodge star*. The relationship between discrete forms and discrete vector fields is studied by Hirani (2003). Forman (1998) studies a different notion of discrete vector fields.

Harmonic discrete one-forms over a mesh are studied by Mercat (2001). His analysis is based on simultaneously looking at a mesh and a *perpendicular* (also called semi-critical) dualization. *Edge weights* are defined using the ratio of the length of a primal edge and its dual. As described by Pinkall and Polthier (1993), for the special case of *mid-perpendicular* (also called *circumcentric*) dualizations, these ratios will result exactly in the so-called *cotangent weights*. This relationship is made explicit by Hi-

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

4

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••

rani (2003). (These same cotangent weights can also be derived by computing the Dirichlet energy of a piecewise linear map (Pinkall and Polthier, 1993).) In this paper, we are unconcerned with the *conformal structure* of the input mesh, thus we will deal with arbitrary edge-weights. Edge weights can be used to define a *Hodge-star* operation (e.g. (Mercat, 2001)), which among other things, results in a definition of *co-closedness*, which we use in this paper. One-forms that are both closed and co-closed are called *harmonic*. Pairs of discrete harmonic one-forms that are related by the Hodge-star operator are used by Mercat to construct *holomorphic* one-forms. Mercat then shows how these discrete holomorphic one-forms converge in the limit to smooth holomorphic one-forms.

Benjamini and Lovasz (2003) prove a number of combinatorial properties of harmonic one-forms. Their work includes a definition equivalent to the *sign-changes* that we will use later in this paper. Independently, Lovasz (2004) describes an index theorem which we prove below in a similar way, and a Tutte-like embedding theorem for the torus, which we prove in quite a different way.

Gu and Yau (2003) use the same notion of harmonic, but define a different discrete Hodge-star operation over these harmonic one-forms, which does not involve a dual mesh. They then use this definition to form pairs of harmonic one-forms which are used to create planar parameterizations of meshes.

3. One-forms on meshes and the Index theorem

In this section we first review the concept of a discrete one-form over a mesh. We then prove a discrete analog of the Poincaré–Hopf index theorem that relates the number of singularities in the one-form to the Euler characteristic of the mesh.

3.1. Harmonic one-forms

Let $G = \langle \mathcal{V}, \mathcal{E}, \mathcal{F} \rangle$ be a mesh. \mathcal{V}, \mathcal{E} and \mathcal{F} are the sets of vertices, edges and faces of G, containing V vertices, E edges and F faces respectively. In this paper we will be concerned with *coherently oriented* meshes which are closed manifolds and have genus g. Since the mesh is oriented we can use the cyclic ordering of the vertices to define a cyclically ordered set of *half-edges*. Since the orientation is coherent, the two "twin" half-edges incident on two adjacent faces have opposite orientations. From now on, when we say "oriented", we will mean coherently oriented.

For face $f, \partial f$ is the *boundary* operator applied to f, yielding the set of half-edges bounding f. There is a well defined ordering of these half edges induced by f's orientation. For vertex $v, \delta v$ is the *coboundary* operator applied to the vertex, returning the set of half-edges emanating from v. There is a well defined ordering over these half edges induced by v's orientation.

Definition 3.1. A discrete one-form $[G, \Delta z]$ is an assignment of a real value Δz_{uv} to each half edge (u, v) of G such that $\Delta z_{uv} = -\Delta z_{vu}$. A half-edge (or edge) will be called *degenerate* if the one-form vanishes at that half-edge. A face or vertex is called *degenerate* if *all* of its half-edges are degenerate, otherwise it will be called *non-degenerate*. A one-form is called *degenerate* if *all* of its faces and vertices are degenerate, otherwise it is called non-degenerate. A one-form is called *vanishing* if *at least one* of its half-edges is degenerate, otherwise it is called *non-vanishing* (i.e., all its edges have non-zero values in the one-form).

cagd940

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●● COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 [DTD5] P.5(1-30) by:violeta p. 5

5

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-•

Since this paper will deal only with discrete one-forms, we will omit the word "discrete" from now on. As we will see, we will be able to characterize the scenario of Tutte's theorem in terms of non-vanishing one-forms on a genus-0 mesh. These one-forms will have certain balances, which will constrain the behaviors of the one-form on faces and vertices. We now restrict our attention to this subset of oneforms.

Definition 3.2. Given a set of (not necessarily symmetric) positive weights w_h associated with each half edge *h* in *G* and a one-form [*G*, Δz], a vertex *v* is called *co-closed with respect to (wrt) w* if

$$\sum_{h\in\delta v} w_h \Delta z_h = 0. \tag{1}$$

A face f is called *closed* if

$$\sum_{h\in\partial f}\Delta z_h=0.$$
(2)

A one-form whose faces are all closed and all vertices co-closed wrt some set of weights is called *har-monic*.

We finish by stating a well known general theorem concerning harmonic one-forms on meshes. This theorem is analogous to a classical theorem on continuous one-forms. The related fact that any vector field on a surface has a unique Hodge decomposition is the basis of the vector visualization techniques of Polthier and Preuss (2003) and Tong et al. (2003).

Theorem 3.3 (Benjamini and Lovasz, 2003; Gu and Yau, 2003; Mercat, 2001). If G is a closed oriented manifold mesh of genus g, then the linear space of harmonic one-forms wrt some set of positive weights has dimension 2g.

The reader may be convinced of the correctness of Theorem 3.3 by observing the following: The rank of the set of Eq. (1) is F - 1. This is despite the fact that there are F equations, since if all but one of the faces are closed, then this implies that the last face must be closed too. The same is true for the set of Eq. (2) – their rank is V - 1. The number of unknowns is E, hence the dimension of the solution space is E - (V + F - 2), which, by Euler's formula, is 2g.

Note that Theorem 3.3 implies that the only harmonic one-form on a closed spherical mesh is the degenerate (all zeros) one-form. Hence, to analyze Tutte parameterizations of disk-like meshes, we will be using one-forms which are harmonic except at boundary vertices, which may not all be co-closed.

3.2. Sign changes and indices

Of particular interest are the sign patterns of a non-vanishing one-form at the half-edges associated with a vertex or face of the mesh. We use these to classify the vertices and faces, as illustrated in Fig. 1.

Definition 3.4. Let $[G, \Delta z]$ be a non-vanishing one-form. The *index* of a vertex v in $[G, \Delta z]$ is ind(v) = (2 - sc(v))/2, where sc(v) is the number of sign changes in the values of Δz as one traverses the halfedges of δv in order. The *index* of a face f in $[G, \Delta z]$ is ind(f) = (2 - sc(f))/2, where sc(f) is the number of sign changes in the values of Δz as one traverses the half edges of ∂f in order. v is called

cagd940

[DTD5] P.6 (1-30)

by:violeta p. 6

S0167-8396(05)00053-1/FLA AID:940 Vol.●●(●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••



Fig. 1. Illustration of Definition 3.4. Classification of vertices and faces by number of sign changes. Arrows denote the orientation of the half-edge with a positive value of the one-form.

a non-singular vertex if ind(v) = 0 and a saddle vertex if ind(v) < 0. If ind(v) = 1, and all values of Δz at v are positive, v is called a source, otherwise v is called a sink. f is called a non-singular face if ind(f) = 0 and a saddle face if ind(f) < 0. If ind(f) = 1, f is called a vortex.

Note that since (by definition) the number of sign changes at a vertex or face must be even, the index is always an integer. Moreover, the index can never exceed +1. When the index is negative, its magnitude indicates the frequency of the one-form sign changes on its half-edges.

This notion of a non-zero index is a discretized notion of the index of a singularity of a smooth oneform. In the smooth case, the index of a smooth *vector field* at an isolated singularity is defined as the (signed) rotation number of the vectors as one traverses a small loop around a point where the field vanishes. The index of a singularity of a smooth *one-form* can be defined by a using some chosen metric to dualize the one-form to a vector field. It is easy to show that the index computed in this manner is invariant to the choice of metric. As is apparent in Fig. 1, sources, sinks, vortices and saddles *look* quite like their continuous counterparts. The index 0 cases correspond, in the smooth case, to regions with no singularities.

The following theorem now characterizes the global distribution of indices of vertices and faces of a mesh in terms of its genus:

Theorem 3.5 (Index theorem). If G is a closed oriented manifold mesh of genus g, then any nonvanishing one-form $[G, \Delta z]$ satisfies

$$\sum_{v \in V} ind(v) + \sum_{f \in F} ind(f) = 2 - 2g.$$

Proof.

$$\sum_{v \in V} ind(v) + \sum_{f \in F} ind(f) = \frac{1}{2} \sum_{v \in V} (2 - sc(v)) + \frac{1}{2} \sum_{f \in F} (2 - sc(f))$$

6

cagd940

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●● COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 [DTD5] P.7(1-30) by:violeta p. 7

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-•

$$= V + F - \frac{1}{2} \left(\sum_{v \in V} sc(v) + \sum_{f \in F} sc(f) \right)$$
$$= V + F - \frac{1}{2} (2E)$$
$$= V + F - E$$
$$= 2 - 2g.$$

The third equality is due to the total number of sign changes in the mesh (over vertices and faces) being equal to the number of half-edges in the mesh. To see this, consider a half edge h bordering some face f and co-bordering some vertex v. Consider h_f , the clockwise successor to h in f, and h_v , the counterclockwise successor to h in v. Clearly, h_f and h_v are half-edge mates, hence must have opposite signs in the one-form. Thus exactly one of these must account for a sign change with h in v and the other in f. The fifth equality is Euler's formula, which can be proven independently in numerous ways, including calculations of the homology groups (Giblin, 1981). \Box

This is simply a discretized version of the Poincaré–Hopf index theorem (Guillemin and Pollack, 1974). An equivalent statement of our Index theorem appeared independently in (Lovasz, 2004). A special case was obtained by Banchoff (1970) and used by Lazarus and Verroust (1999). However, their theorem applies only to so-called *null-cohomologous* one-forms (arising from the differences of a scalar potential field defined on the mesh vertices) over a triangle mesh, and thus does not consider faces of non-zero index. Consequently, their summation of indices is over vertices exclusively, whereas we sum over faces as well. Indeed, in Section 4 we deal with non-triangular interior faces (which we need to prove are not saddles), and more importantly, with a non-convex exterior face (that in fact *is* a saddle). Furthermore, in Section 5 we deal with closed meshes of genus one and higher, where we will deal exclusively with one-forms which are not null-cohomologous.

Closedness of faces and co-closedness of vertices are directly related to their indices:

Corollary 3.6. If a face f is closed in a non-vanishing one-form then $ind(f) \leq 0$. If a vertex v is coclosed wrt to some set of positive weights, then $ind(v) \leq 0$.

Proof. If f were a vortex, then all of the terms of the sum in (1) would be positive and thus could not sum to zero. The same holds for v a source or sink. \Box

4. Parameterizing a disk

Tutte's theorem may be stated as follows:

Theorem 4.1 (Tutte (1963)). Let $G = \langle \mathcal{V}, \mathcal{E}, \mathcal{F} \rangle$ be a 3-connected planar graph with boundary vertices $\mathcal{B} \subset \mathcal{V}$ defining a unique unbounded exterior face f_e . Suppose ∂f_e is embedded in the plane as a (not necessarily strictly) convex planar polygon, and each interior vertex is positioned in the plane as a strictly convex combination of its neighbors, then the straight-line drawing of G with these vertex positions is an embedding. In addition, this embedding has strictly convex interior faces.

7

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●● COMAID:m2 v 1.39 Prn:29/06/2005; 12:09

cagd940

8

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••

A graph is called 3-*connected*, if it remains connected after the removal of any two vertices and their incident edges. A manifold mesh $\{\mathcal{V}, \mathcal{E}, \mathcal{F}\}$ is called 3-*connected* if its one-skeleton, the graph $\{\mathcal{V}, \mathcal{E}\}$, is 3-connected. As we will see, this property is necessary to preclude various types of degeneracies in the drawing.

A drawing is an embedding if no two edges intersect, except at vertices.

Tutte constructs such an embedding by solving the following linear system for the x and y coordinate values of the vertices:

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

$$\sum_{v_j \in N(v_i)} w_{ij} y_j = y_i, \quad i = 1, ..., V - B,$$

$$x_i = b_i^x, \quad i = V - B + 1, ..., V,$$

$$y_i = b_i^y, \quad i = V - B + 1, ..., V,$$
(3)

where the interior vertices are labeled as $\{1, \ldots, V - B\}$, and the remaining boundary vertices as $\{V - B + 1, \ldots, V\}$. The b_i are the coordinates of the vertices of a convex polygon. The w_{ij} , associated with each half edge e_{ij} , are any set of positive numbers with unit row sums (hence the term *convex combinations*). We do not assume that w_{ij} are symmetric. $N(v_i)$ is the set of vertices neighboring v_i .

We denote by [G, x, y] the straight line plane drawing using this solution as coordinates for the vertices, and call it a *Tutte drawing*.

4.1. Single convex boundary

In Tutte's scenario, the boundary polygon, whose coordinates are $[b^x, b^y]$ is assumed to be (not necessarily strictly) convex. It is well known that (3) has a unique solution [x, y]. This follows directly from the fact that the linear system is irreducible, and is weakly diagonally dominant with at least one strongly diagonally dominant row (Varga, 2000).

Proving Tutte's theorem amounts to showing that a Tutte drawing is an embedding. We proceed in this direction by examining the properties of span(x, y)—all the different projections of the drawing, and constructing a one-form on *G*. For any choice of reals α and β , define for each vertex v_i at coordinates (x_i, y_i) , the quantity $z_i \equiv \alpha x_i + \beta y_i$.

At every interior vertex v_i , we have:

$$z_{i} \equiv \alpha x_{i} + \beta y_{i}$$

$$= \alpha \sum_{v_{j} \in N(v_{i})} w_{ij} x_{j} + \beta \sum_{v_{j} \in N(v_{i})} w_{ij} y_{j}$$

$$= \sum_{v_{j} \in N(v_{i})} w_{ij} (\alpha x_{j} + \beta y_{j})$$

$$= \sum_{v_{j} \in N(v_{i})} w_{ij} z_{j}.$$
(4)

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09

cagd940

[DTD5] P.9(1-30) by:violeta p. 9

9

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••



Fig. 2. The generic structure of a Tutte drawing, looking at the one-form which is the projection of the drawing along the vertical z. All vertices and faces are non-singular, except for one source and one sink on the boundary. Arrows mark the orientation of the half-edges possessing positive values of the one-form.

Now define the one-form $\Delta z_{ij} = -\Delta z_{ji} \equiv z_j - z_i$. Since the rows of w_{ij} sum to unity, (4) implies that every interior vertex v of $[G, \Delta z]$ is co-closed wrt w:

$$\sum_{h\in\delta v} w_h \Delta z_h = 0.$$

Furthermore, because the Δz are differences of the z values at vertices, they must sum to zero along any directed closed loop in G. In particular, each face f satisfies

$$\sum_{h\in\partial f}\Delta z_h=0$$

meaning $[G, \Delta z]$ is closed.

Fig. 2 illustrates the generic structure of the one-form Δz . Due to the convexity of the boundary \mathcal{B} , a property which is preserved under linear transformations, \mathcal{B} will generically contain one vertex with a maximum z and one with a minimum z. Every other vertex of \mathcal{B} has one neighbor in \mathcal{B} with a strictly greater z value, and its other neighbor in \mathcal{B} with a strictly smaller z value. As a result, none of these other vertices can be sources or sinks in $[G, \Delta z]$, i.e., they must have non-positive index.

The following lemma concerning one-forms characterizes to the one forms $[G, \Delta z]$ obtained from a Tutte drawing.

Lemma 4.2. If G has genus 0 and $[G, \Delta z]$ is a non-vanishing one-form such that all F faces are closed, V – B vertices are co-closed wrt some set of positive weights and of the remaining B vertices, B - 2 have index ≤ 0 , then $[G, \Delta z]$ has no saddle vertices and no saddle faces.

Proof. By the Index theorem (Theorem 3.5), the sum of the indices over all of the faces and vertices on a spherical mesh must total +2. By Corollary 3.6, all the faces and V - 2 vertices can contribute only non-positive values to this sum. The only way to achieve the sum of +2 is for all *F* faces and V - 2 of the vertices to have vanishing index, and for the remaining two vertices to have the maximal index value of +1. In other words, $[G, \Delta z]$ has no saddles (but it does have a source vertex and a sink vertex). \Box



S0167-8396(05)00053-1/FLA AID:940 Vol. •••(••• COMAID:m2 v 1.39 Prn:29/06/2005; 12:09

cagd940

10

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••

It is possible for a one-form, Δz as constructed above from a Tutte drawing [G, x, y], to be vanishing. For example, if an interior edge has zero length, then it will produce a degenerate one-form in all projections. Such a one-form will be vanishing and Lemma 4.2 will not directly apply. Fortunately it is possible to slightly perturb any vanishing one-form into a non-vanishing one-form *without* altering the signs of the one-form values on the non-degenerate edges. Appendix A shows how this is achieved. In particular it defines a *consistent perturbation* with the following properties: the perturbed one-form $\Delta z'$ is non-vanishing, no saddles are removed by the perturbation, and no new positive-indexed vertices or faces are created. The following three lemmas are proven in Appendix A.

Lemma A.5. Let Δz be a one-form derived from a Tutte drawing as in Section 4.1 using any α and β . Then there exists some consistent perturbation of Δz .

Lemma A.6. Let Δz be a one-form derived from a Tutte drawing as in Section 4.1 using any α and β . Then for any consistent perturbation of Δz , the perturbed one-form is non-vanishing and has at most two vertices with positive index.

Lemma A.7. For any consistent perturbation $\Delta z'$ of Δz , if there were sc sign changes around ∂f (δv respectively) in Δz , ignoring zeros, then $ind(f) \leq 1 - sc/2$ ($ind(v) \leq 1 - sc/2$ respectively) in $\Delta z'$.

From Lemma A.6, using the same reasoning as in the proof of Lemma 4.2, we obtain:

Corollary 4.3. If [G, x, y] is a Tutte drawing, then for any α , β and (any consistent perturbations of) $[G, \Delta z]$ constructed as in (4), no vertex or interior face is a saddle in $[G, \Delta z]$.

To proceed, we will need to rely on the fact that there are no degeneracies in a Tutte drawing. The following lemma is proved in Appendix B, using the 3-connectedness of G, and we will proceed under this assumption.

Lemma B.5. In a Tutte drawing there can be no face with zero area, no edge of zero length and no angle of 0 or π within any interior face.

(For arbitrary *n*-gons, these notions are all distinct.) Essentially this is proven by showing that if there were any such degeneracy, then there must exist a consistent perturbation that *does* result in a saddle, contradicting Corollary 4.3.

Next, we show that the faces and vertex one-rings in the drawing [G, x, y] must be *properly behaved*, as defined below. In particular, we show that if they are not properly behaved, then there must be a $z \in span(x, y)$ such that $[G, \Delta z]$ has either a saddle vertex or a saddle face (which would be in contradiction to Corollary 4.3).

Definition 4.4. A face f of G is called a *convex* face of [G, x, y] if its boundary is a simple, strictly convex polygon in the plane with non-zero area.

Definition 4.5. Let v be a vertex of G. Let α_i be the signed angles between adjacent half-edges in δv (angles are measured by going the "short" way between half edges, so $0 < |\alpha_i| < \pi$). v is called a *wheel* vertex of [G, x, y] if α_i all have the same sign and $|\sum \alpha_i| = 2\pi$.

S0167-8396(05)00053-1/FLA AID:940 Vol.●●(●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

[DTD5] P.11(1-30) by:violeta p. 11

11

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••



Fig. 3. Illustration of Definitions 4.4 and 4.5 and proof of Theorems 4.6 and 4.7. The horizontal dashed lines demonstrate that there always exists a line that intersects a non-wheel in more than two wedges and a non-convex face at more than two edges.

Fig. 3 shows some examples of non-convex faces and vertices which are not wheels.

Theorem 4.6. If [G, x, y] is a Tutte drawing, then all interior faces of G are convex.

Proof. Suppose the face f of G is non-convex in [G, x, y]. Then there exists a line l in the plane that intersects four or more edges of f. Rotate the drawing in the plane using matrix

$$\begin{pmatrix} w_i \\ z_i \end{pmatrix} = \begin{pmatrix} \beta & -\alpha \\ \alpha & \beta \end{pmatrix} \begin{pmatrix} x_i \\ y_i \end{pmatrix}$$

such that *l* is horizontal (see the dashed lines in Fig. 3). The z_i component now represents the vertical component of the rotated drawing. This means that the half-edges in ∂f exhibit at least four sign changes (ignoring zero values). If Δz is non-vanishing, then *f* is a saddle face in $[G, \Delta z]$. (By Lemma A.5, there exist consistent perturbations, and by Lemma A.7, for any consistent perturbation, *f* will remain a saddle face. If Δz is vanishing, then by Lemma A.7, for any consistent perturbation, *f* will be a saddle face.) The existence of this saddle contradicts Corollary 4.3.

Theorem 4.7. If [G, x, y] is a Tutte drawing, then all interior vertices of G are wheels.

Essentially we want to show that if a particular vertex v is not a wheel, then we can always find a line l through v that intersects four or more wedges (an angular extent between two topologically neighboring edges at v) of the drawing. Then, as argued above, the appropriate choice of α and β will make l horizontal and v a saddle in $[G, \Delta z]$, again contradicting Corollary 4.3.

To be more precise, for a particular vertex v, define a map $m: S^1 \to S^1$ as follows. Imagine placing the *n* edges that are incident to *v* evenly spaced around the unit circle (the domain of *m*), according to the orientation of *v*. In the Tutte embedding, these edges each point in some direction, with adjacent edges separated by angles α_i . We use this to define the map, *m*, for each of the evenly spaced edges in the domain. Between each pair of adjacent edges along a wedge, we can complete the map *m* to be

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●● COMAID:m2 v 1.39 Prn:29/06/2005; 12:09

cagd940

12

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••

angularly linear. The map *m* will have some integer degree: *d*. Since the wedges form a cycle, we must have $|\sum \alpha_i| = 2\pi d$,

Any non-wheel falls into one of the three distinct cases:

 $d \ge 2$ (i.e., the edges cycle around v more than once): Then any direction +q in the drawing must have at least two orientation preserving (or two orientation reversing) preimages under m. So must -q. This yields a line l intersecting at least four wedges.

d = 0 (i.e., the edges must wind and unwind the same amount): By Lemma B.5, v's neighbors cannot lie on a single line. Since v is in the strict convex hull of its neighbors, there must exist a vector such that both +q and -q pass through at least one wedge of v. But since d = 0, +q must have an equal number of orientation preserving and reversing preimages in m. This means that +q must intersect at least two wedges. The same holds for -q. This yields a line l intersecting at least four wedges.

d = 1 (i.e., the edges wind around once before meeting up) and there are two adjacent wedges with opposite signed angles: We can find a direction +q that intersects both adjacent wedges. This direction must have at least one orientation preserving and one orientation reversing preimages in m. In order to sum to d = 1 there also must be at least one more preimage. This implies that +q intersects at least three wedges. The opposite direction -q must also have at least one preimage in m, and thus must intersect at least one wedge. Thus the line l spanned by +q intersects at least four wedges. \Box

(Similar reasoning shows that the "half-ring" of interior faces around each boundary vertex maps homeomorphically to a "half-disk".)

Corollary 4.8. Two faces that share an edge are disjoint in a Tutte drawing.

Proof. *G* is a manifold, and so each interior edge *e* is on the boundary of two faces. By Theorem 4.6 these faces are convex. And so both faces must lie either completely in the same or opposite half-plane defined by *e*. By Theorem 4.7, these two faces must also be part of a wheel in a neighborhood of the any of the two vertices of *e*. Thus these two convex faces must lie on opposite sides of *e* and are therefore disjoint. \Box

As a result of Corollary 4.8 we say that the drawing is *locally* an embedding, or, as Floater (2003a) calls it—*locally injective*. The following theorem establishes that a Tutte drawing is also a global embedding (or, what Floater calls *globally* injective), namely that *any* two faces in the drawing are disjoint.

Theorem 4.9. All of the faces in a Tutte drawing are disjoint.

Proof. Each point in the plane is contained in a finite number of bounded convex faces—its *face count*. We will show that this number must be *one* at every point inside the convex hull of the boundary vertices \mathcal{B} .

Each interior vertex is in the strict convex hull of its neighbors, and so cannot be an extremal vertex of $CH(\mathcal{V})$ —the convex hull of all of the vertices of *G*. As a result, $CH(\mathcal{B}) = CH(\mathcal{V})$. As a result, no interior vertex can be outside of $CH(\mathcal{B})$. Thus the face count must be zero for any point outside $CH(\mathcal{B})$.

We can incrementally compute the face count for any point p inside the convex hull, by starting at any point outside CH(\mathcal{B}) and walking along a straight path to p. This path can cross the convex exterior face f_e only once at one boundary edge. Crossing this boundary edge must increment the face count by

cagd940

S0167-8396(05)00053-1/FLA AID:940 Vol. •••(••• COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 [DTD5] P.13(1-30) by:violeta p. 13

13

one. Since G is manifold, whenever the path crosses an interior edge, it must be incident to exactly two faces. By Theorem 4.8 these faces are disjoint, and therefore, the edge count must remain at one. Hence the face count at p must be one. \Box

Corollary 4.10. *No two edges of a Tutte drawing can intersect, except at a vertex.*

Corollary 4.10 and Theorem 4.6 together are equivalent to Tutte's Theorem 4.1, thus our proof of Tutte's theorem is complete.

To summarize the basic steps of the proof: Given a set of edge weights, and convex boundary mapping, there is a unique Tutte drawing [G, x, y]. For any projection (and any perturbation of it) we construct a non-vanishing one-form $[G, \Delta z]$. By Corollary 4.3, (derived from the Index theorem) due to the lack of sources, sinks, and vortices, there can be no saddles in any such $[G, \Delta z]$. By Theorems 4.6 and 4.7, if there were any non-convex faces or non-wheel vertices in [G, x, y], then we *could* construct such a $[G, \Delta z]$ with saddles, which would be a contradiction. Since our drawing has all convex faces and wheel vertices, Corollary 4.8 states that this implies local injectivity. This local property coupled with the convexity of the boundary mapping implies global injectivity in Theorem 4.9.

4.2. Multiple non-convex boundaries

While the scenario of Tutte's theorem requires the single boundary to be convex in order that the resulting drawing be a planar embedding (i.e., contain only wheel vertices and convex faces), our theory permits the presence of non-convex vertices (so-called *reflex* vertices) in the boundary. In fact, we permit the presence of *multiple* non-convex boundaries.

We can thus think of our mesh G as a topological sphere with some faces labeled as *exterior*. One of these exterior faces will be *unbounded* in the planar drawing, while the rest will be *bounded*. The boundaries of these exterior faces will be mapped to the plane, and impose boundary conditions in the system of equations (3). In order to produce a correctly oriented drawing, we will require that the turning number of the boundary of the unbounded exterior face be $+2\pi$ and the turning number of the boundaries of -2π . We will also require that every reflex vertex be in the strict convex hull of its neighbors. As we will see, these restrictions force the drawing to be an embedding. First a few definitions.

Note: our discussion here assumes that the unbounded exterior face of P is oriented clockwise, while the bounded exterior faces are oriented counter-clockwise. The entire discussion is, of course, also true if we consistently reverse these conventions.

Definition 4.11. Let P be a straight line mapping of an oriented polygon to the plane, with all edges having positive length (but allowed to cross). The *turning angle* of P at vertex v is the external angle at v as one traverses P consistent with its orientation. This angle is positive if the turn at v is a right turn in the plane and negative if the turn is a left turn. The *turning number* of P is the sum of the turning angles at the vertices of P.

Definition 4.12. Let P be as above. A vertex v is called *convex* in P if the turning angle at v is non-negative. Otherwise v is called *reflex* in P.

S0167-8396(05)00053-1/FLA AID:940 Vol. •••(••• COMAID:m2 v 1.39 Prn:29/06/2005; 12:09

cagd940

14

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-•••

Note that a bounded exterior face that is drawn as a convex polygon with turning number -2π , will in fact have *all* reflex vertices (since these boundaries have been drawn counterclockwise).

Definition 4.13. Let P be as above. A vertex v of P is called *extremum* relative to a direction d in the plane, if the edges emanating from v all project positively or all negatively onto d.

Lemma 4.14. Let P be as above with turning number $+2\pi$ (-2π respectively). Denote by C its set of convex extrema, and by R its set of reflex extrema with respect to any direction d. Then |C| - |R| = 2 (|R| - |C| = 2 respectively).

Proof. Imagine a continuous "flattening" operation applied to the polygon which simply scales *P* by a factor of *s* along the direction orthogonal to *d* with decreasing *s*. Since the scaling is orthogonal to *d*, it cannot change any dot products with *d*, and thus cannot create or remove any extrema. At the limit $s \rightarrow 0$, the turning angle will be $+\pi$ at each $v \in C$, $-\pi$ at every $v \in R$, and 0 at all other vertices. Since the total turning number is $+2\pi$ (-2π respectively), this means that $\pi(|C| - |R|) = 2\pi$ ($\pi(|C| - |R|) = -2\pi$ respectively), namely |C| - |R| = 2 (|R| - |C| = 2 respectively).

In our parameterization setting, the boundary of each exterior face is mapped to the plane. This mapping imposes boundary conditions on the linear system of equations (3). Under this mapping (and chosen orientation for G), the turning number of each external face's boundary is well-defined, and its vertices may be classified as convex or reflex.

Lemma 4.15. Suppose that: (1) *G* is an oriented 3-connected mesh of genus 0 having multiple exterior faces. (2) The boundary of the unbounded exterior face is mapped to the plane with positive edge lengths and turning number 2π . (3) The boundaries of the bounded exterior faces are mapped to the plane with positive edge lengths and turning number -2π . (4) [*G*, *x*, *y*] is the straight line drawing of *G* where each internal vertex is positioned as a convex combination of its neighbors. (5) In [*G*, *x*, *y*] the reflex vertices of all of the exterior face boundaries are also in the convex hull of their neighbors. Then for any α , β and [*G*, Δz] constructed as in (4), no vertex or interior face is a saddle in [*G*, Δz].

Proof. As illustrated in Fig. 4, assume the *N* boundaries of [G, x, y] form one unbounded and N - 1 bounded polygonal exterior faces in the plane with B_i vertices each, of which C_i are convex extremum vertices, and R_i are reflex extremum vertices. Consider the one-form $[G, \Delta z]$. Since all the extremal vertices of the exterior faces produce sign changes of $[G, \Delta z]$ in those faces, and only those do, the indices of the exterior faces are $1 - (C_i + R_i)/2$. The interior faces are closed and the interior vertices are co-closed, hence their indices are ≤ 0 . Having both positive and negative values of Δz on their co-boundaries, the indices of the non-extremal boundary vertices are ≤ 0 . Being inside the convex hull of their neighbors, the indices of the reflex extrema are also ≤ 0 . Trivially, the convex extrema have indices ≤ 1 . Denote by *s* the sum of the (negative) indices of the saddle interior faces and vertices. So the sum of the indices over the entire mesh is $\leq \sum(1 - (C_i + R_i)/2) + \sum C_i + s$. But the Index theorem maintains that this sum is +2, implying $2 \leq N + \sum(C_i - R_i)/2 + s$. Now Lemma 4.14 applied to the exterior faces implies that $\sum(C_i - R_i) = 2 - 2(N - 1) = 4 - 2N$. This means that $s \geq 0$, but since, by definition $s \leq 0$, we conclude that s = 0, namely, saddles do not exist in the interior faces or vertices.

Using arguments identical to those of the previous section, we conclude:

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

[DTD5] P.15(1-30) by:violeta p. 15

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-•



Fig. 4. The multiple non-convex boundary scenario. The one-form considered is the projection on the vertical (namely $\Delta z = \Delta y$). Arrows mark the half-edges having positive values of the one-form. Vertices are labeled with their indices.

Theorem 4.16. Under the conditions of Lemma 4.15, all interior faces of [G, x, y] are convex and all interior vertices are wheels. Moreover, if all the exterior faces are embedded as disjoint simple polygons (edge crossings are not allowed), then all faces of [G, x, y] are disjoint.

A natural question is whether Theorem 4.16 is of any use in practical parameterization scenarios, as Tutte's theorem is. Forcing the boundary vertices to form a convex shape, as in the Tutte scenario, is easy, but is it possible to relax that requirement, yet force these boundary reflex vertices to be in the convex hulls of their neighbors without compromising the same property of neighboring vertices? This seems to be quite difficult, since it is impossible to determine apriori which boundary vertices will be convex, and which reflex. Fortunately, Theorem 4.16 implies that the injectivity of the drawing is determined entirely by the behavior of the boundary vertices, hence we have to worry only about these. One way to do this is to replace the 2*B* linear boundaries, expressing the fact that all the wedges formed by these faces have the correct orientation. This involves considerably less inequalities than what would have been required had Theorem 4.16 not been true, as in that case, at least *F* inequalities would have been required, one for each face. However, being inequalities (as opposed to equalities), this will introduce many degrees of freedom into the solution.

A more practical approach is to require more of the boundary. In many applications, a *conformal* parameterization is sought, meaning one that preserves angles as much as possible in the transition from 3D to 2D. Traditional linear conformal parameterizations use the Tutte method with *cotangent* weights (Pinkall and Polthier, 1993), however, these weights can sometimes be negative, hence are actually in-

15

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●● COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

[DTD5] P.16(1-30) by:violeta p. 16

16

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••

appropriate for the Tutte scenario. Floater (2003b) recently proposed to replace these cotangent weights with so-called *mean-value* weights, which are always positive and seem to result in parameterizations which are close to conformal. The Angle-Based Flattening (ABF) method of Sheffer and de Sturler (2001) addresses the problem in the triangular case by solving directly for the angles of the triangles, and reconstructing the embedding from that. Forcing all the angles to assume values in the interval $(0, \pi)$ (along with other constraints) guarantees that the resulting 2D triangulation is injective. This approach is non-linear, but has the advantage of a free boundary.

Here we describe one alternative we have explored. Given a 3D triangle mesh M as input, with angles $0 \ll \alpha_i \ll \pi$ incident on the boundaries, and β_i elsewhere, we construct an injective 2D parameterization of M, such that the 2D angles are as close as possible to α_i and β_i . Theorem 4.16 implies that in order to obtain an embedding, it suffices to force the boundary angles to be close to α_i , and compute the interior vertices' positions by solving a linear system using mean-value weights derived from the β_i . If the resulting boundary angles are between 0 and π , then reflex vertices will naturally be in the convex hull of their neighbors, hence the result an embedding. Since we do not know in advance which vertices will be reflex, we apply our constraint on the angles around *all* boundary vertices.

Fig. 5 shows some results of this parameterization algorithm on two 3D input meshes. The first input is the "ear" mesh, which is embedded as a triangulation with a non-convex boundary. Note how the 2D boundary has a shape very close to that of the 3D boundary and how the angles are very similar. The second input is the "face" mesh, containing multiple boundaries ("holes"). Note that the hole corresponding to the "mouth" has a non-convex character in the 3D input, which is preserved in the resulting 2D embedding. The third input is a hemisphere, in which three slits have been cut. The resulting 2D embedding takes advantage of these slits when forming the boundary for the resulting conformal parameterization.

A linear system for free boundary parameterization is described by Desbrun et al. (2002) for a triangulated mesh with single boundary. However, that method may result in drawings which are not embeddings, for some inputs, even when all of the interior edge weights are positive.

Recent work of Karni et al. (2005) shows how to possibly improve on the results of Desbrun et al. (2002) by *iterating* a linear system of equations. Karni et al. show how Theorem 4.16 implies that if this iteration converges, then the limit drawing is guaranteed to be an embedding.

The following remains an important open problem: Is it possible to generate a straight-line embedding of a 3-connected triangulated 3D mesh with a single exterior boundary face, by posing a single set of *linear* equations on all the 2D mesh vertex coordinates, including the boundary vertices, possibly coupling the *x* and *y* coordinates. The coefficients in the equations should be derived from the 3D geometry, such that if the input mesh is already a planar embedding, the output should be identical to the input (this is called a 2D-*reproducing* scheme).

Colin de Verdiere (personal communication) has recently shown us how to prove Theorem 4.16 directly from Tutte's theorem. However, his proof does not distinguish between local and global injectivity of the embedding. Our approach has the advantage of cleanly separating the conditions required for local injectivity (Theorem 4.15) from the extra ones needed for global injectivity (simplicity of the boundaries, as in Theorem 4.16).

The proof of Colin de Verdiere proceeds as follows: Consider the convex hull of the boundary \mathcal{B} of the unbounded exterior face. Since \mathcal{B} is simple, the difference between the two is the union of simple polygons. Triangulate each of these polygons, as well as the bounded exterior faces. The result is a straight-line plane drawing of a new 3-connected graph with no unbounded exterior faces, whose boundary is embedded to a convex shape. Each interior vertex in the new drawing is connected to at least all the

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

[DTD5] P.17 (1-30) by:violeta p. 17

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••

17



Fig. 5. Parameterizing a mesh with free boundaries: (a), (c), (e) Original 3D meshes. (b), (d), (f) 2D parameterizations of (a), (c) and (e) when boundaries (bounded and unbounded) are free but forced to have boundary angles as close as possible to 3D originals. Mean-value weights were used in the harmonic equations for the interior vertices.

vertices it was connected to in the old drawing. Each interior vertex of the old drawing is still an interior vertex in the new drawing. Each reflex boundary vertex of the old drawing is now an interior vertex in the new drawing, positioned inside the convex hull of its neighbors. Each convex boundary vertex of the

S0167-8396(05)00053-1/FLA AID:940 Vol. •••(••• COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

18

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-•••

old drawing is now either a convex boundary vertex of the new drawing, or an interior vertex of the new drawing positioned inside the convex hull of its neighbors. Thus the new drawing satisfies the conditions of Tutte's theorem, hence is an embedding.

5. Parameterizing a torus

While our use of one-forms on meshes allowed us to obtain Tutte's theorem for disks and extensions, it has a more natural application in the case of a toroidal mesh. This case is actually easier to analyze because, in its simplest form, there is no boundary to complicate matters. On the other hand, it is more difficult to envision a parameterization in this case, since the torus obviously cannot be mapped in an injective manner to the plane. The traditional solution to this is to cut the mesh along an artificial boundary to form a disk, and then parameterize as any other disk-like mesh. While this is certainly possible, cutting the mesh introduces new problems such as optimization of this boundary, and obvious discontinuities in the parameterization along the boundary.

Another way to parameterize a toroidal mesh without encountering the cutting problem, is to parameterize it *locally*, meaning injectively embed any submesh with disk topology, yet in a way such that all local parameterizations fit together in a seamless manner. So, while one never attempts to parameterize the entire mesh, if two intersecting disk-like regions (whose intersection is also disk-like) are parameterized to the plane, the parameterization coincides on the intersection, possibly after an appropriate translation. See Fig. 6.

Seamless local parameterization is important for a variety of mesh processing applications, in particular "cutting and pasting" between meshes (Biermann et al., 2002), texture mapping (Kraevoy et al., 2003), meshing (Tewari et al., 2005) and remeshing.

Gu and Yau (2003) showed how seamless local parameterizations for the torus can be achieved using one-forms. In particular, the parameterization is driven by a pair of harmonic one-forms. Each of the two one-forms provides the information needed to synthesize each of the two coordinate values for the vertices in the plane. Their paper did not address whether the algorithm would produce a locally-planar embedding.

Here, we prove, based on the Index theorem (Theorem 3.5) that this algorithm will in-fact succeed in producing an injective parameterization. In this sense, this is a Tutte-like embedding theorem for the torus.

(Our theorem applies to a slightly more general case than that originally explored in (Gu and Yau, 2003). Gu and Yau dealt specifically with special *cotangent* weights in their Eq. (1). They also only dealt



Fig. 6. Seamless local parameterization of disk-like submeshes of the torus.

cagd940

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●● COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 [DTD5] P.19(1-30) by:violeta p. 19

with pairs of one-forms related by a special Hodge-star operator which they define, and only considered meshes with triangular faces. Our theorem applies to *arbitrary* positive weights, *any* pair of linearly independent harmonic one-forms, and applies to meshes with *arbitrary* sized faces.)

A recent paper of Steiner and Fischer (2004) makes observations similar to ours, in particular that linearly independent harmonic one-forms generate locally-injective parameterizations of the torus. The proofs they give, however, are rather complicated. Another quite different (and also more complicated) proof for this appeared independently in (Lovasz, 2004).

As opposed to the disk case, where we considered the planar coordinates of the embedding, and then converted them to a one-form to prove Tutte's theorem, in the toroidal case the algorithm starts off with one-forms, and then synthesizes the local embeddings from that.

The algorithm begins by picking two linearly independent harmonic one-forms. Theorem 3.3 implies that the space of harmonic one-forms on the torus is two-dimensional. A basis for this space can be found by simply solving for the null-space of the matrix representing the closedness and co-closedness constraints. This matrix is typically quite large, but also very sparse (six entries per row on the average), hence its nullspace may be computed using efficient numerical methods for sparse matrices. Two linearly independent one-forms may be then sampled from this space in a variety of ways. (Note that unlike the original description of Gu and Yau (2003), this process does not require computing any *basis loops*.)

To construct a local parameterization for the torus, the algorithm chooses any two linearly independent harmonic one-forms, Δx and Δy , which are linearly independent solutions to (1) and (2). Now, given a submesh with the topology of a disk, we assign the coordinates (0, 0) in the planar parametric domain to an arbitrary vertex v_0 of that submesh. Any other vertex is assigned coordinates by *integrating* (summing) the one-form along a directed path from v_0 to v. Since the one-form is closed, it does not matter which path is used:

$$(x_0, y_0) = (0, 0), \tag{5}$$

$$(x_i, y_i) = \left(\sum_{h \in P(v_0, v_i)} \Delta x_h, \sum_{h \in P(v_0, v_i)} \Delta y_h\right).$$
(6)

Because the space is two-dimensional, it does not really matter which two harmonic one-forms are used, as long as they are independent. Any other pair of one-forms will be a linear combination of these, meaning the resulting parameterizations will be related to each other by an affine transformation.

The proof that this algorithm will produce an injective mapping proceeds as follows. We begin with an analog of Lemma 4.2 for a torus:

Lemma 5.1. If G is a closed oriented manifold mesh with genus 1 and Δz a non-vanishing harmonic one-form on G, then $[G, \Delta z]$ has only non-singular vertices and faces.

Proof. By Theorem 3.5, the sum of the indices of the vertices and faces of $[G, \Delta z]$ is 0. Since all vertices are co-closed and all faces are closed, their indices are all non-positive. Thus the only way these indices can sum to zero is if they are all zero. \Box

We now use Lemma 5.1 to prove the analog of Theorems 4.6 and 4.7 for the torus:

19

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09

cagd940

[DTD5] P.20(1-30) by:violeta p. 20

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-•••



Fig. 7. Parameterization of a torus containing 32 vertices and 64 faces. (a) 3D torus. (b) Parameterization of the torus to the plane using two harmonic one-forms generated with uniform weights. Vertices are numbered. The color coded edges along the boundary correspond. (c) Double periodic tiling of the plane using the drawing in (b). (For interpretation of the references in color in this figure legend, the reader in referred to the web version of this article.)

Theorem 5.2. If G is a 3-connected oriented manifold mesh with genus 1, Δx and Δy two nondegenerate and linearly independent harmonic one-forms on G, and G' a submesh of G with the topology of a disk, then all faces of [G', x, y] are convex and all vertices of [G', x, y] are wheels, where x and y are constructed as in (5).

Proof. Identical to the proofs of Theorems 4.6 and 4.7, using Lemma 5.1 instead of Lemma 4.2. Note that if the two harmonic one-forms are linearly dependent, then the whole mapping will collapse to a line. In this case, there exist projections for which the one-form is *completely* degenerate, and the arguments in Appendix B cannot be applied. \Box

Note that if the method of (5) is run twice, each time with a different vertex as v_0 , the two resulting parameterizations will, by definition, be related to each other by a simple translation in the plane. The translation vector will be the coordinates of the second origin in the first parameterization (or vice versa).

Theorem 5.2 is a statement of local injectivity. It can also be shown that a pair of harmonic one-forms, in fact creates a globally injective mapping from the universal cover of the torus to the entire plane. As a result, there can be no edge crossings in these parameterizations. The proof relies on Theorem 5.2, but also on some notions from algebraic topology, and is thus omitted.

Fig. 7 shows such an parameterization of a toroidal mesh, where the disk-like submesh is actually the entire mesh, after it was cut twice along two *basis* loops of the handle. Because of the periodicity of the torus, the resulting embedding can be used to tile the plane in a doubly-periodic seamless manner.

Although we will not prove this, we believe that a toroidal mesh with boundaries ("holes") may be locally parameterized in an injective embedding in a manner similar to that of Section 4.2, namely, by allowing some vertices along the boundaries to be non-harmonic. This will hold if these boundaries have turning number -2π , and the reflex vertices on the boundaries are contained in the convex hulls of their neighbors.

20

S0167-8396(05)00053-1/FLA AID:940 Vol. •••(••• COMAID:m2 v 1.39 Prn:29/06/2005; 12:09

cagd940

21

6. Higher genus

While we have applied our Index theorem (Theorem 3.5) only to the disk and to the genus 1 case to prove Tutte-like embedding theorems, the theory is applicable also to higher genus meshes, except there matters are more complicated. By Theorem 3.3, the dimension of the space of harmonic one-forms for g > 1 is at least 4 and so there are many fundamentally different "pairs of harmonic one forms" that can be chosen from the space in order to create a parameterization.

In addition, for g > 1, the Index theorem implies the existence of at least two saddle vertices and/or faces in every such harmonic one-form. So *something* must go awry when we apply the parameterization approach of Section 5.

As observed by Gu and Yau (2003), the "best" one can hope for is to have all of the "badness" in the parameterization isolated at 2g - 2 vertices or faces. In this case we may have a vertex that is *doubly wheel*, where the co-boundary edges will cycle in the drawing around the vertex twice. Or we may have a face that is *doubly convex*, where the boundary edges cycle around the face twice (see the third column of Fig. 3). At these bad spots, the embedding cannot be even locally injective. Note that if the mesh has only triangular faces, doubly convex faces cannot occur.

Here we prove that if a pair of harmonic one-forms is chosen such that it has 2g - 2 such *doubles*, then the rest of the resulting parameterization is locally injective.

Theorem 6.1. If G is a closed oriented 3-connected manifold mesh of genus g, and Δx and Δy are harmonic one-forms on G and [G, x, y]—the corresponding drawing in the plane—contains 2g - 2 vertices or faces that are doubly wheel/convex, then all other vertices of [G, x, y] are wheels and all faces are convex.

Proof. In *any* projection of [G, x, y], a vertex that is doubly wheel, or face that is doubly convex, will be a saddle. If there are 2g - 2 doubles, then the Index theorem (Theorem 3.5) implies that there can be no other saddles. So all other vertices and faces are always non-singular in every projection, hence all other vertices must be wheels and all faces convex. \Box

Of course, since the mapping is not everywhere locally injective, it is obviously not globally injective. But the drawing will at least contain faces which are all oriented consistently.

While we do not have a closed characterization of which two independent one-forms from the 2g-dimensional solution space will, when used as x and y coordinates in the plane, form such an embedding, we can generate them using the following randomized (Las-Vegas) algorithm:

- 1. Compute a basis of the 2g-dimensional space of harmonic one-forms on G.
- 2. Select a one-form Δx at random from the space whose basis was computed in (1), e.g., a random linear combination of the basis functions. When integrated, this one-form will generate the *x* coordinate of the embedding.
- 3. Solve for another one-form Δy such that $[G, \Delta x, \Delta y]$ has 2g 2 "doubles". Since Δx has been fixed in step 2, this can be written as a linear program. In this program, every pair of adjacent edges at the saddles of Δx is constrained to have a positive angle (i.e., positive cross product of the two one-forms of those edges). If it exists, this Δy is called a *mate* of Δx .
- 4. If step 3 failed, goto step 2.

cagd940

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09

[DTD5] P.22(1-30) by:violeta p. 22

22

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••



Fig. 8. Parameterization of the two-hole torus. Left: the 3D mesh, containing 495 triangular faces. Middle: The parameterization of the mesh to the plane using uniform weights. The two double-wheel vertices are marked in red and green. The green one appears twice along the boundary. Right: Zoom into the red double-wheel vertex. (For interpretation of the references in color in this figure legend, the reader in referred to the web version of this article.)

It is easy to see that if Δx and Δy are mates, then any two linear combinations of Δx and Δy are also mates. Hence, in practice, it is possible to impose orthogonality of Δx and Δy in the linear program solved in step 3.

We emphasize that step 3 does indeed fail for some random choices made in step 2, meaning there do exist one-forms Δx for which there is *no* mate Δy (not even the Δy generated by applying the discrete Hodge star operation of Gu and Yau (2003) to Δx). However, in practice, the linear program fails to find a mate in step 3 only rarely, so the algorithm usually terminates after a small number of steps. We have observed experimentally, though, that this failure rate increases with the genus.

The embeddings in Fig. 8 were generated using this algorithm.

We conclude this discussion by stating some natural questions which are left open:

- 1. Which vertices in *G* can be saddles in a harmonic one-form? (We know that vertices with valence 3 cannot be saddles, because three edges cannot generate more than 2 sign changes of the one-form.)
- 2. Which vertices in G can appear as double wheels? (We known vertices with valence ≤ 4 cannot be saddles since each of the four angles must be $< \pi$, yet their sum must be 4π).
- 3. If a vertex can be a saddle or double wheel, how can we generate a one-form or pair of one-forms having this property?
- 4. Is there any natural characterization of the one-forms that have mates?

7. Conclusion

The concept of one-forms on meshes used in this paper, although simple, seems to be quite powerful. It unleashes a wealth of classical mathematical theory for which discrete analogs seem to exist. This is also related to some recent developments in polytopal graph theory (Mihalisin and Klee, 2000) and planar tilings using harmonic functions on graphs (Kenyon, 1998).

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●● COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

23

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••

This paper deals with the general case of asymmetric weights $w_{ij} \neq w_{ji}$ in (3). Many other papers (including Tutte (1963)) deal only with the symmetric case. This is appealing because then the system has a physical interpretation of a spring system, and the Tutte drawing minimizes the sum of the squares of the weighted spring lengths, hence the system's energy. Some of the recipes for generating barycentric coordinates for given embeddings yield symmetric weights, including the so-called cotangent weights (Pinkall and Polthier, 1993). However, most do not (e.g., the mean-value weights (Floater, 2003b)), and it seems that the one-form theory presented here is powerful enough to deal with this.

Theorem 4.16 raises hope for other possible applications. One of these is constrained parameterization (see e.g. (Kraevoy et al., 2003)), which is of paramount importance for texture mapping, guaranteeing that key features of a texture are mapped to the corresponding features on a mesh. This problem calls for an embedding a disk-like mesh in the plane, such that the parameterization is injective, but also satisfies positional constraints at a (usually small) subset of the interior vertices. The results of Section 4 indicate that if all the "problematic" regions of the mesh are embedded properly, then harmonicity will take care of the rest. In the scenario of Section 4—the problematic regions are the boundary vertices, which are taken care of either by forcing convexity or explicit positive angles. This can be generalized to the case of constrained parametrization by considering the constrained vertices to also be "problematic" (or, in other words, part of a non-connected boundary), so it seems that forcing *both* the boundary vertices *and* the constrained vertices to be wheels should solve the problem. However, it remains to see how this can be done in a computationally efficient manner. First results in this direction have been obtained recently by Karni et al. (2005).

A possible application of local parameterization of the torus is for parameterizing a disk-like mesh with a free boundary. Since Tutte's theorem requires a convex boundary, one common way around this is to "pad" the disk with a number of layers of "virtual" faces, forming a new "virtual" boundary. This larger mesh is embedded in the plane using the convex boundary method, and the extra padding then discarded. This method, due to Lee et al. (2002), gives the true boundary more flexibility, and it will typically end up being non-convex. It is, however, still influenced by the virtual boundary and the connectivity of the virtual faces, hence is not artifact-free. Making use of our method for local parameterization of the torus, we believe it is more natural to embed the original disk within a *torus*, rather than a larger disk, as the torus seems to be the "cleanest" mesh. After solving for a harmonic one-form on the torus, this is transformed into an embedding of just the original disk-like submesh. This procedure eliminates boundary conditions entirely, hence should contain less artifacts.

Acknowledgements

We thank Eric Colin de Verdiere for helpful comments on an early version of this paper. The work of C. Gotsman was partially supported by European grant FP6 IST Network of Excellence #506766 (AIM@SHAPE).

Appendices

The appendices provide some technical theorems which simplify the main results of this paper. The theorems will be formulated for the case of a mesh with a disk-like topology and convex boundary. With

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

24

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-•••

the appropriate modifications, these theorems carry over to the other cases (non-convex boundary, higher genus) as described.

Appendix A. Perturbing a vanishing one-form into a non-vanishing one-form

We will show how any (possibly vanishing) one-form on a disk-like mesh obtained from a Tutte drawing that contains vanishing values may be *consistently perturbed* into a non-vanishing one-form. Since the perturbation will be *sufficiently small*, it will not change the signs of any of the non-zero values. As a result, any such perturbation will not be able to remove any saddles from the original one-form. Moreover, because the perturbation is *consistent*, the perturbed one-form will not have any new sources, sinks or vortices (index +1) that were not in the original one-form.

The perturbed one-from may have non co-closed vertices. But as mentioned in Section 4, the key to Lemma 4.2 is that all but two of the vertices have non-positive values of the one-form on their co-boundaries. This is weaker than the co-closedness property. More formally:

Definition A.1. A vertex (respectively face) is *mixed* in a one-form if it has at least one positive and at least one negative value of the one-form on its co-boundary δv (respectively boundary ∂f).

Definition A.2. A one-form is called *mixed* if all its vertices and faces are mixed. A one-form is called *almost mixed* if all its faces are mixed and all its vertices are mixed, with the exception of at most two vertices.

The following key lemma follows as a special case of Theorem 2.2 of Linial et al. (1988):

Lemma A.3. If G is a 2-connected oriented manifold mesh $\langle \mathcal{V}, \mathcal{E}, \mathcal{F} \rangle$ and s and t any two distinct vertices of G, then there exists a non-vanishing one-form $[G, \Delta f]$, whose faces are all closed and whose vertices, (except for s, which is a source, and t, which is a sink), are all co-closed with respect to some set of positive edge weights w_{ij} .

Clearly such a $[G, \Delta f]$ is almost mixed.

Denote by f_e the outer face of the mesh. Tutte's method dictates that ∂f_e is embedded as a nondegenerate convex polygon, with no two vertices coincident. For any specific choice of α and β , in the rotated Tutte drawing [G, w, z] (where the rotation is determined by α and β), the boundary loop ∂f_e has vertices on its left side, and vertices on its right side. The left side has a top (and bottom) vertex, as does the right side. With respect to a generic choice of α and β , both sides will share their top (and bottom) vertex with each other. With respect to some non-generic choices of α and β , the upper or lower edges may be perfectly horizontal and so the two sides may not share their top (or bottom) vertices. In this case there will be a distinct top-left and top-right vertex. In addition, if the boundary is only weakly convex, then for a specific choice of α and β , there can also be "strictly top" and "strictly bottom" vertices that are strictly in between the left and right sides. (See Fig. A.1.)

Definition A.4. Pick any two vertices s and t on ∂f_e that are not strictly top or strictly bottom vertices. Let Δf be any non-vanishing almost mixed one-form with source s and sink t. Let ε be a non-zero

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

[DTD5] P.25(1-30) by:violeta p. 25

25

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••



Fig. A.1. Scenario of Lemma A.6. Arrows mark the orientation of the half-edges possessing positive values of the one-form. Edges without arrows have vanishing (zero) values.

scalar that is small enough such that adding $\varepsilon \Delta f$ to Δz will not to change the signs of the previously non-vanishing values of Δz . Any such perturbation is called a *consistent perturbation*. See Fig. A.1 for an illustration of this.

Lemma A.5. Let Δz be a one-form derived from a Tutte drawing as in Section 4.1 using any α and β . Then, there exists a consistent perturbation.

Proof. By Lemma A.3, for any choice of *s* and *t*, such a Δf must exist. Additionally, since Δz is defined over a finite number of edges, such an ε of sufficiently small magnitude must exist. \Box

Lemma A.6. Let Δz be a one-form derived from a Tutte drawing as in Section 4.1 using any α and β . Then, for any consistent perturbation, the resulting $\Delta z'$ is non-vanishing and almost mixed.

Proof. Since Δf is non-vanishing, and ε sufficiently small, the resulting $\Delta z' = \varepsilon \Delta f + \Delta z$ is non-vanishing.

Next we prove $\Delta z'$ is almost mixed by analyzing a small set of cases.

(1) Any degenerate face in Δz is determined completely by Δf , hence is mixed in $\Delta z'$.

(2) Any non-degenerate face f is mixed in Δz , hence for sufficiently small ε remains so in $\Delta z'$. Hence all faces are mixed in $\Delta z'$.

(3) The sign pattern of any degenerate vertex v in Δz is completely determined in $\Delta z'$ by Δf . Since s and t were explicitly chosen to be non-degenerate vertices, a degenerate v cannot be one of the chosen s or t. Hence such v must be co-closed (with respect to some weights) in Δf and thus mixed in $\Delta z'$.

(4) Any mixed vertex in Δz , for sufficiently small ε , must remain so in $\Delta z'$.

(5) The only non-degenerate, not mixed vertices possible in Δz are the extreme ones: the top-left (TL), bottom-left (BL), top-right (TR), bottom-right (BR), strictly top (ST), and strictly bottom (SB) vertices. We now show that there cannot be more than two non-mixed vertices from this set in the resulting $\Delta z'$.

If TL and TR coincide, then it can account for at most one non-mixed vertex in $\Delta z'$.

If TL and TR are distinct, then we use the following argument. By Lemma 4.2, there are no saddle faces in Δf , and so there can be only 2 sign changes as one circles ∂f_e (only at *s* and *t*). Hence $\Delta z'$

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●● COMAID:m2 v 1.39 Prn:29/06/2005; 12:09

cagd940

[DTD5] P.26(1-30) by:violeta p. 26

26

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••

must go in one direction (wlog say from left to right) along the top (resp bottom) of ∂f_e (see Fig. A.1). Therefore, the ST vertices and TR will have both positive and negative values on their co-boundaries in $\Delta z'$, hence be mixed. Thus TL can account for at most one non-mixed vertex.

An identical argument shows that BL and BR can account for at most one non-mixed vertex.

Lemma A.7. For any consistent perturbation $\Delta z'$ of Δz , if there were sc sign changes around ∂f (δv respectively) in Δz , ignoring zeros, then $ind(f) \leq 1 - sc/2$ ($ind(v) \leq 1 - sc/2$ respectively) in $\Delta z'$.

Proof. The sign of a non-degenerate edge of Δz is preserved in $\Delta z'$. Hence the number of sign changes can only increase and the index only decrease. \Box

This means that this consistent perturbation cannot remove any saddles.

For the non-convex boundary, in Lemma A.6, we choose for *s* and *t* any two boundary vertices that have at least one boundary edge with non-zero value in Δz . The resulting perturbed one-form $\Delta z'$ will have the appropriate properties for Lemma 4.2 to apply.

For the torus, (and higher genus) case we must assume that the one-form is not degenerate. Then for a consistent perturbation, we can choose *any* two non-degenerate vertices as *s* and *t*. The resulting perturbed one-form $\Delta z'$ will have *all* mixed vertices and faces and thus appropriate for Lemma 5.1 (respectively Theorem 6.1) to apply.

Appendix B. No degenerate vertices or faces

In this appendix we prove that a Tutte embedding does not contain degenerate elements. First we define what these are:

Definition B.1. Let Δz be a one-form on a mesh. A *degenerate* corner is a (vertex, face) pair whose two associated edges are degenerate.

For this appendix we need a slightly stronger version of Lemma A.3, which also follows directly from a variant of the arguments in (Linial et al., 1988).

Lemma B.2. If G is a 2-connected oriented manifold mesh $\langle \mathcal{V}, \mathcal{E}, \mathcal{F} \rangle$, s and t any two distinct vertices of G, and p any directed simple path from s to t, then there exists a non-vanishing one-form $[G, \Delta f]$, with all positive values along the half-edges of p, whose faces are all closed and whose vertices, (except for s, which is a source, and t, which is a sink), are all co-closed with respect to some set of positive edge weights w_{ij} .

We now prove a series of lemmas leading to the desired result.

Lemma B.3. Let [G, x, y] be a Tutte drawing. Then in any projected one-form $[G, \Delta z]$ there can be no non-degenerate interior vertex participating in a degenerate corner.

cagd940

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 [DTD5] P.27 (1-30) by:violeta p. 27

27

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••



Fig. B.1. Scenario of Theorem B.3. Arrows mark the orientation of the half-edges possessing positive values of the one-form. Edges without arrows have vanishing (zero) values.

Proof. Let v_0 be an interior vertex. Since it is non-degenerate, it must be mixed. We wish to show that if it participates in a degenerate corner, then we can find $\Delta z'$ —an appropriate perturbation of Δz of the form described in Lemma A.6—which is a non-vanishing almost mixed one-form *with* a saddle at v_0 . This would contradict Lemma 4.2.

Call the two edges of the degenerate corner e_{01} , e_{02} . These edges connect v_0 to v_1 and v_2 . Since G is 3-connected, the graph $G - \{v_0\}$ is 2-connected. Therefore for any two vertices s and t we can find two vertex-disjoint paths connecting v_1 and v_2 to s and t (we cannot say in advance which will be connected to which) such that v_0 is not in either path (Linial et al., 1988). By including all of the edges in these two paths in addition to the edges e_{01} and e_{02} , we obtain a simple directed path from s to t that proceeds in the order $[v_1, v_0, v_2]$ or $[v_2, v_0, v_1]$. See Fig. B.1. By Lemma B.2, there exists Δf : a non-vanishing almost mixed one-form with source s and sink t that either passes in order $[v_1, v_0, v_2]$ or $[v_2, v_0, v_1]$. Consider $\Delta z' = \Delta z + \varepsilon \Delta f$. This is a non-vanishing almost-mixed one-form with non-vanishing values on the edges e_{01} and e_{02} . With the proper choice of sign for ε , we can increases the number of sign changes around v_0 , creating a saddle at v_0 in $\Delta z'$. This contradicts Lemma 4.2. \Box

Next we show that if there was any degenerate corner at an interior face, there would have to be some degenerate corner with a participating non-degenerate vertex.

Lemma B.4. Let [G, x, y] be a Tutte drawing. Then in any projected one-form $[G, \Delta z]$ there can be no degenerate corner at an interior face.

Proof. Assume (v_0, f) is a degenerate corner. If it is not part of a triangle (which would have to be degenerate by closedness), then introduce a new edge e_{12} between vertices v_1 and v_2 (the neighbors of v_0 at the degenerate corner); this splits f into a degenerate triangle and some remainder face. This *face-split* operation cannot change the 3-connectedness of G. The addition of this edge cannot change the co-closedness of v_1 and v_2 . Repeat this operation for any degenerate corner. In the final G, any interior face with a degenerate corner must be a degenerate triangle.

Since [G, x, y] is not degenerate, there must be some interior face f that is not degenerate sharing an edge e_{01} , between vertices v_0 and v_1 , with a degenerate face. Either v_0 or v_1 must be an interior vertex, otherwise e_{01} would be an interior edge connecting two boundary vertices, which is impossible in a 3-connected planar graph.

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09

cagd940

[DTD5] P.28(1-30) by:violeta p. 28

28

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••



Fig. B.2. The scenario of Lemma B.4.

Without loss of generality v_0 is an interior vertex. f has another vertex— v_3 —that shares an edge e_{03} with v_0 . The one-form must be non-vanishing on this edge, otherwise e_{01} and e_{03} would have been a degenerate corner and f would have been a degenerate triangle. See Fig. B.2. So v_0 is an interior non-degenerate vertex at a degenerate corner, in contradiction of Lemma B.3. \Box

Lemma B.5. In a Tutte drawing there can be no face with zero area and no edge of zero length and no angle of 0 or π within any interior face.

Proof. Suppose there was a face with zero area. It is then possible to pick a projection (α, β) such that the resulting one-form vanishes on all edges of this face. Similarly, if there is an edge *e* with zero length, then pick (α, β) such that the resulting one-form will vanish on one of the edges neighboring on *e*. In both cases, we will have a degenerate corner in the one-form, in contradiction of Lemma B.4. \Box

The proofs of Appendix B apply directly to the non-convex boundary case of Theorem 4.16. They also apply directly to one-forms on the torus. The equivalent to Lemma B.4 for the torus will state that in *any* (non-degenerate) harmonic one-form $[G, \Delta z]$, there can be no degenerate corner. This implies no geometric degeneracies in any drawing [G', x, y] integrated from a pair of non-degenerate harmonic one-forms.

Harmonic one-forms on higher genus meshes have saddle vertices or faces and are thus more complicated. In this case there will exist one-forms $[G, \Delta z]$ with degenerate corners. But, in the special cases treated by Theorem 6.1, the saddles are all "accounted for", so again, no degenerate corners can exist.

References

Banchoff, T.F., 1970. Critical points and curvature for embedded polyhedral surfaces. Amer. Math. Monthly 77, 475–485. Becker, B., Hotz, G., 1987. On the optimal layout of planar graphs with fixed boundary. SIAM J. Comput. 16 (5), 946–972. Benjamini, I., Lovasz, L., 2003. Harmonic and analytic functions on graphs. J. Geometry 76 (1–2), 3–15.

Biermann, H., Martin, I., Bernardini, F., Zorin, D., 2002. Cut-and-paste editing of multiresolution subdivision surfaces. ACM Trans. Graph. 21 (3), 312–321.

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

29

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-••

- Bossavit, A., 1999. Computational electromagnetism and geometry. J. Japan Soc. Appl. Electromagn. Mech. 7, 150–159; 7, 294–301; 7, 401–408; J. Japan Soc. Appl. Electromagn. Mech. 8 (2000) 102–109; 8, 203–209; 8, 372–377.
- Colin de Verdière, E., Pocchiola, M., Vegter, G., 2003. Tutte's barycenter method applied to isotopies. Computational Geometry 26 (1), 81–97.
- Desbrun, M., Meyer, M., Alliez, P., 2002. Intrinsic parameterizations of surface meshes. Computer Graphics Forum 21 (2). (Proceedings of Eurographics).
- Floater, M.S., 1997. Parameterization and smooth approximation of surface triangulations. Computer Aided Geometric Design 14, 231–250.
- Floater, M.S., 2003a. One-to-one piecewise linear mappings over triangulations. Math. Comput. 72, 685-696.
- Floater, M.S., 2003b. Mean value coordinates. Computer Aided Geometric Design 20, 19–27.
- Floater, M.S., Gotsman, C., 1999. How to morph tilings injectively. J. Comput. Appl. Math. 101, 117–129.
- Floater, M.S., Hormann, K., 2004. Surface parameterization: A tutorial and survey. In: Dodgson, N.A., Floater, M.S., Sabin, M.A. (Eds.), Advances in Multiresolution for Geometric Modelling. Springer, Berlin, pp. 157–186.
- Floater, M.S., Pham-Trong, V., 2005. Convex combination maps over triangulations, tilings, and tetrahedral meshes. Adv. Comput. Math., in press.
- Forman, R., 1998. Morse theory for cell complexes. Adv. Math. 134, 90-145.
- Giblin, P.J., 1981. Graphs, Surfaces and Homology. Chapman & Hall.
- Gotsman, C., Surazhsky, V., 2001. Guaranteed intersection-free polygon morphing. Computers and Graphics 25 (1), 67–75.
- Guillemin, V., Pollack, A., 1974. Differential Topology. Prentice-Hall, Englewood Cliffs, NJ.
- Gu X., Yau S.-T., 2003. Global conformal surface parameterization. In: Proceedings of the ACM/Eurographics Symposium on Geometry Processing, Aachen.
- Hatcher, A., 2002. Algebraic Topology. Cambridge Univ. Press, Cambridge.
- Hirani, A.N., 2003. Discrete exterior calculus. PhD Thesis, California Institute of Technology.
- Hiptmair, R., 2002. Finite elements in computational electromagnetism. In: Acta Numerica. Cambridge Univ. Press, Cambridge, pp. 237–339.
- Kanai, T., Suzuki, H., Kimura, F., 2002. Metamorphosis of arbitrary triangular meshes. IEEE Comput. Graph. Appl. 20, 62–75.
- Karni, Z., Gotsman, C., Gortler, S.J., 2005. Free-boundary linear parameterization of 3D meshes in the presence of constraints. In: Proceedings of Shape Modeling International.
- Kenyon, R., 1998. Tilings and discrete Dirichlet problems. Israel J. Math. 105, 61-84.
- Kotiuga, P.R., 1984. Hodge decompositions and computational electromagnetics. PhD Thesis, McGill University.
- Kraevoy, V., Sheffer, A., Gotsman, C., 2003. Matchmaker: Constructing constrained texture maps. ACM Trans. Graph. 22 (3), 326–333. (Proceedings of SIGGRAPH 2003).
- Lazarus, F., Verroust, A., 1999. Level set diagrams of polyhedral objects. In: Proceedings of ACM Solid Modeling '99.
- Lee, Y., Kim, H.S., Lee, S., 2002. Mesh parameterization with a virtual boundary. Computers and Graphics 26 (5), 677–686.
- Lovasz, L., 2004. Discrete analytic functions: An exposition. In: Grigor'yan, A., Yau, S.-T. (Eds.), Eigenvalues of Laplacians and Other Geometric Operators. In: Surveys in Differential Geometry, vol. IX. International Press.
- Linial, N., Lovasz, L., Wigderson, A., 1988. Rubber bands, convex embeddings and graph connectivity. Combinatorica 8 (1), 91–102.
- Mercat, C., 2001. Discrete Riemann surfaces. Comm. Math. Phys. 218 (1), 77-216.
- Mihalisin, J., Klee, V., 2000. Convex and linear orientations of polytopal graphs. Discrete Comput. Geom. 24, 421-435.
- Nedelec, J.C., 1980. Mixed finite elements in R^3 . Numer. Math. 35, 315–341.
- Pinkall, U., Polthier, K., 1993. Computing discrete minimal surfaces and their conjugates. Experimental Math. 2 (1), 15-36.
- Polthier, K., Preuss, E., 2003. Identifying vector fields singularities using a discrete Hodge decomposition. In: Hege, H.C., Polthier, K. (Eds.), Proceedings of Visualization and Mathematics III. Springer, Berlin, pp. 113–134.
- Raviart, P.A., Thomas, J.M., 1977. A Mixed Finite Element Method for Second Order Elliptical Problems. In: Lecture Notes in Math., vol. 606. Springer, Berlin.
- de Rham, G., 1931. Sur l'analysis situs des varietes a *n* dimensions. J. Math. Pures Appl. 10, 200. C. R. Acad. Sci. (Paris) 188 (1929) 1651–1652.
- Richter-Gebert, J., 1996. Realization Spaces of Polytopes. Lecture Notes in Math., vol. 1643. Springer, Berlin.
- Sen, S., Sen, S., Sexton, J.C., Adams, D.H., 2000. Geometric discretization scheme applied to the abelian Chern–Simons theory. Phys. Rev. E (3) 61 (3), 3174–3185.

S0167-8396(05)00053-1/FLA AID:940 Vol.●●●(●●●) COMAID:m2 v 1.39 Prn:29/06/2005; 12:09 cagd940

[DTD5] P.30(1-30) by:violeta p. 30

30

S.J. Gortler et al. / Computer Aided Geometric Design ••• (••••) •••-•••

- Sheffer, A., de Sturler, E., 2001. Parameterization of faceted surfaces for meshing using angle based flattening. Engineering with Computers 17 (3), 326–337.
- Steiner, D., Fischer, A., 2004. Planar parameterization for closed 2-manifold genus-1 meshes. In: Proceedings of ACM Solid Modeling, Genova.
- Tewari, G., Gotsman, C., Gortler, S.J., 2005. Meshing genus-1 point clouds. Harward University Computer Science Technical Report TR-14-05.
- Thomassen, C., 2004. Tutte's spring theorem. J. Graph Theory 45, 272-280.
- Tong, Y., Lombeyda, S., Hirani, A., Desbrun, M., 2003. Discrete multiscale vector field decomposition. In: Proceedings of SIGGRAPH.
- Tutte, W.T., 1963. How to draw a graph. Proc. London Math. Soc. 13 (3), 743-768.
- Varga, R.S., 2000. Matrix Iterative Analysis. Springer, Berlin.
- Whitney, H., 1957. Geometric Integration Theory. Princeton Univ. Press, Princeton, NJ.