

ONE-TO-ONE PIECEWISE LINEAR MAPPINGS OVER TRIANGULATIONS

MICHAEL S. FLOATER

ABSTRACT. We call a piecewise linear mapping from a planar triangulation to the plane a *convex combination mapping* if the image of every interior vertex is a convex combination of the images of its neighbouring vertices. Such mappings satisfy a discrete maximum principle and we show that they are one-to-one if they map the boundary of the triangulation homeomorphically to a convex polygon. This result can be viewed as a discrete version of the Radó-Kneser-Choquet theorem for harmonic mappings, but is also closely related to Tutte's theorem on barycentric mappings of planar graphs.

1. INTRODUCTION

This paper establishes a condition for the injectivity of piecewise linear mappings over triangulations which satisfy a discrete maximum principle, and has application in geometric modelling, computer graphics, and numerical grid generation.

The theorem can be viewed as a discrete analog of a property of harmonic mappings known as the Radó-Kneser-Choquet theorem; see Figure 1. Recall that a function u defined on some closed region $D \subset \mathbb{R}^2$ is said to be *harmonic* if it satisfies the Laplace equation

$$u_{xx} + u_{yy} = 0$$

in the interior of D , and we say that a mapping $\phi = (u, v) : D \rightarrow \mathbb{R}^2$ is harmonic if both its components $u(x, y)$ and $v(x, y)$ are harmonic. The following theorem was first conjectured by Radó [19] and later proved independently by Kneser [12] and Choquet [3].

Radó-Kneser-Choquet theorem. *Suppose $\phi : D \rightarrow \mathbb{R}^2$ is a harmonic mapping which maps the boundary ∂D homeomorphically into the boundary $\partial\Omega$ of some convex region $\Omega \subset \mathbb{R}^2$. Then ϕ is one-to-one.*

Various aspects of the theorem and some generalizations are discussed by Duren and Hengartner [5]. A key element of Kneser's proof is that harmonic functions satisfy a maximum principle [18].

In this paper, we establish a similar property (Theorem 4.1) for certain piecewise linear mappings; see Figure 2. Specifically, we establish the injectivity of mappings $\phi : D \rightarrow \mathbb{R}^2$ which are piecewise linear over a triangulation \mathcal{T} of D and have the

Received by the editor July 9, 2001.

2000 *Mathematics Subject Classification.* Primary 65N30, 58E20; Secondary 05C85, 65M50.

Key words and phrases. Triangulation, harmonic map, maximum principle, parameterization, numerical grid generation, planar embedding.

©2002 American Mathematical Society




FIGURE 1. Harmonic map




FIGURE 2. Convex combination map

property that the image under ϕ of every interior vertex of \mathcal{T} is a convex combination of the images of its neighbouring vertices. We will call such mappings *convex combination mappings*, and analogously to the continuous case, a key element of the proof is that they satisfy a *discrete* maximum principle.

Convex combination maps appear to be more abundant than harmonic maps because, as shown in Proposition 3.3, all one-to-one piecewise linear maps are convex combination maps. Nevertheless, as we will see, certain convex combination maps arise as discrete approximations to Laplace's equation and it is in this sense that we can regard Theorem 4.1 as a discrete version of the RKC theorem. Discretizations of harmonic maps have been used for several years in numerical grid generation [2], [8], [14]. More recently, discrete harmonic-like maps (whose domain is a surface) have turned out to be very important in geometric modelling and computer graphics, where they have been used to construct *parameterizations*, i.e., parametric representations of surfaces in [9], [6], [13], and [16]. Such maps have also been applied to image morphing [10]. In all these applications, it is crucial that the mapping is one-to-one and the need for a condition to ensure injectivity was emphasized in [4].

The question of whether piecewise linear maps are one-to-one is closely related to a classical problem of graph theory: Can a (simple) planar graph G be embedded in the plane with straight line edges such that the only intersections between edges are at common endpoints? This question was first answered affirmatively by Fáry [7]. Tutte [20] later gave a constructive proof for so-called 3-connected graphs (which includes triangulations), based on what he called a *barycentric mapping*. The boundary nodes of G are mapped to the boundary of a convex polygon in the plane and each interior node is mapped to the barycentre of the images of its neighbouring nodes. It was later argued in [9] that Tutte's result generalizes to allow arbitrary barycentric (convex) combinations. However, Tutte's proof is

long and relies on the nontrivial fact that planar graphs contain no Kuratowski subgraphs. We establish Theorem 4.1 without needing Kuratowski’s theory, using instead the discrete maximum principle.

Finally, we remark that Theorem 4.1 also holds when the domain triangulation is replaced by a surface triangulation and is thus directly applicable to parameterization methods, such as those in [9], [6], and [13].

2. TRIANGULATIONS AND CONNECTIVITY

By a *triangle* we mean the convex hull $T = [v_1, v_2, v_3]$ of three points $v_1, v_2, v_3 \in \mathbb{R}^2$. If v_1, v_2, v_3 are collinear, T reduces to a line segment or point, in which case we say that T is *degenerate*.

Definition 2.1. Let \mathcal{T} be a finite set of nondegenerate triangles and let $D_{\mathcal{T}} = \bigcup_{T \in \mathcal{T}} T$. We will call \mathcal{T} a triangulation if

- (i) the intersection of any pair of triangles is either empty, a common vertex, or a common edge, and
- (ii) the edges in \mathcal{T} belonging to only one triangle form a simple closed polygon $\partial D_{\mathcal{T}}$.

From condition (ii), $D_{\mathcal{T}}$ is simply connected with boundary $\partial D_{\mathcal{T}}$. We denote by V and E the sets of vertices and edges of \mathcal{T} , respectively. We call vertices and edges contained in $\partial D_{\mathcal{T}}$ *boundary* vertices and edges and otherwise *interior* vertices and edges.

By a *path* in \mathcal{T} we mean a sequence P of one or more distinct vertices

$$(2.1) \quad v_1, v_2, \dots, v_k,$$

in which v_i and v_{i+1} are neighbours for each i , i.e., $[v_i, v_{i+1}]$ is an edge in \mathcal{T} . If $k = 1$, we call P a *null* path, as it contains only one vertex. Otherwise, we say that P *connects* the two vertices v_1 and v_k . Essential to the question of connectivity of triangulations will be paths P in which all the intermediate vertices (if any) v_2, v_3, \dots, v_{k-1} are interior vertices of \mathcal{T} . We call such paths *interior paths*. In particular it will be important whether a given boundary vertex can be connected to a given interior vertex by an interior path. For example, in Figure 3, only the four boundary vertices v_1, v_2, v_5 , and v_6 can be connected by an interior path to the interior vertex v . We will call an edge $[v, w]$ of \mathcal{T} a *dividing edge* if it is an




FIGURE 3. Connectivity and dividing edges

interior edge and both v and w are boundary vertices. The triangulation in Figure 3 contains one dividing edge, $[v_2, v_5]$.

Lemma 2.2. *Every interior vertex v_1 of a triangulation \mathcal{T} can be connected to at least three boundary vertices by an interior path. If \mathcal{T} contains no dividing edges, then v_1 can be connected to every boundary vertex by an interior path.*

Proof. Due to the finiteness of \mathcal{T} , it is clear that there must exist at least one interior path P of the form (2.1) connecting a given interior vertex v_1 to a boundary vertex v_k . We will show that there are at least two boundary vertices in addition to v_k which can be connected to v_1 by an interior path. Let w_1 and w_2 be the two boundary vertices of \mathcal{T} adjacent to v_k and let Q be the path consisting of the neighbours of v_k which connects w_1 to w_2 . The interior vertex v_{k-1} splits Q into two subpaths, Q_1 connecting v_{k-1} to w_1 and Q_2 connecting v_{k-1} to w_2 . Let y_1 be the first boundary vertex of Q_1 and let y_2 be the first boundary vertex of Q_2 . By combining the path v_1, \dots, v_{k-1} with the part of Q_1 from v_{k-1} to y_1 , we obtain an interior path which connects v_1 to y_1 . Similarly, we obtain an interior path which connects v_1 to y_2 . This completes the first part of the lemma.

Next suppose that \mathcal{T} contains no dividing edges. Then all neighbours of v_k other than w_1 and w_2 are interior vertices. Therefore $y_1 = w_1$ and $y_2 = w_2$ and so both w_1 and w_2 can be connected to v_1 by an interior path. Continuing in this way around the boundary of \mathcal{T} , we see that all boundary vertices can be connected to v_1 by interior paths. \square

The second part of the lemma motivates the following concept of connectivity, which we will use later.

Definition 2.3. *We say that a triangulation \mathcal{T} is strongly connected if it contains no dividing edges.*

3. CONVEX COMBINATION FUNCTIONS AND MAPPINGS

Given a triangulation \mathcal{T} , we will call a function $f : D_{\mathcal{T}} \rightarrow \mathbb{R}$ a *piecewise linear function* if it is continuous over $D_{\mathcal{T}}$ and is linear over each triangle in \mathcal{T} . The set of all such functions forms a linear space, which in spline theory is often written as $S_0^1(\mathcal{T})$, where the 0 indicates C^0 continuity and the 1 indicates the polynomial degree of f in each triangle. As is well known, any function in $S_0^1(\mathcal{T})$ is uniquely determined by its values at the vertices of \mathcal{T} and the dimension of $S_0^1(\mathcal{T})$ is equal to the number of vertices in \mathcal{T} .

We will study functions f in $S_0^1(\mathcal{T})$ with the property that the value of f at every interior vertex v is a convex combination of its values at the neighbours of v . For any vertex v of \mathcal{T} , let N_v be the set of neighbours of v in \mathcal{T} .

Definition 3.1. *Let $f : D_{\mathcal{T}} \rightarrow \mathbb{R}$ be a piecewise linear function, and suppose, for every interior vertex v of \mathcal{T} , that there exist weights $\lambda_{vw} > 0$, for $w \in N_v$, such that*

$$(3.1) \quad \sum_{w \in N_v} \lambda_{vw} = 1,$$

and

$$(3.2) \quad f(v) = \sum_{w \in N_v} \lambda_{vw} f(w).$$

Then we will call f a convex combination function.

Similarly, we will understand a *piecewise linear mapping* $\phi : D_{\mathcal{T}} \rightarrow \mathbb{R}^2$ to be any mapping $\phi = (u, v)$ such that both components u and v belong to $S_0^1(\mathcal{T})$. We call ϕ a *convex combination mapping* if, for every interior vertex v of \mathcal{T} , there are positive weights λ_{vw} satisfying (3.1) such that

$$\phi(v) = \sum_{w \in N_v} \lambda_{vw} \phi(w).$$

Thus the point $\phi(v)$ lies in the convex hull of the neighbouring points $\phi(w)$, $w \in N_v$.

Functions (and mappings) of this kind arise in various ways. Typically the weights λ_{vw} are chosen according to the application in hand together with the values of f at the boundary vertices of \mathcal{T} . Then the values of f at the interior vertices are found by solving the linear system (3.2). It can easily be shown that the system has a unique solution. This was done in [9] using the discrete maximum principle [1], but it can also be deduced by appealing to the theory of M -matrices [21].

For example, the finite element approximation in $S_0^1(\mathcal{T})$ to the solution of Laplace's equation with Dirichlet boundary conditions is the function f in $S_0^1(\mathcal{T})$ which minimizes

$$(3.3) \quad \int_{D_{\mathcal{T}}} |\nabla f(x)|^2 dx$$

subject to the boundary conditions. Letting h_v in $S_0^1(\mathcal{T})$, for $v \in V$, be the nodal basis function defined by $h_v(w) = \delta_{vw}$, for $w \in V$, the minimization is equivalent to finding f in $S_0^1(\mathcal{T})$ such that

$$(3.4) \quad \int_{D_{\mathcal{T}}} \nabla f(x) \cdot \nabla h_v(x) dx = 0$$

for all interior vertices v of \mathcal{T} ; see [11]. Noting that

$$f(x) = \sum_{w \in V} f(w) h_w(x),$$

we find that equation (3.4) is precisely (3.2), where

$$\lambda_{vw} = - \frac{\int_{D_{\mathcal{T}}} \nabla h_v \cdot \nabla h_w}{\int_{D_{\mathcal{T}}} \nabla h_v \cdot \nabla h_v}.$$

Using the fact that the nodal functions h_v sum to one, we find that

$$\int_{D_{\mathcal{T}}} \nabla h_v \cdot \nabla h_v = - \sum_{w \in N_v} \int_{D_{\mathcal{T}}} \nabla h_v \cdot \nabla h_w,$$


and it follows that the weights λ_{vw} , $w \in N_v$, sum to one. The remaining question is whether they are positive.

Some elementary calculations yield the well-known formula

$$(3.5) \quad \int_{D_{\mathcal{T}}} \nabla h_v \cdot \nabla h_w = -\frac{1}{2}(\cot \theta_1 + \cot \theta_2),$$

where θ_1 and θ_2 are the angles at the corners v_1 and v_2 in the two triangles $[v_1, v, w]$ and $[v_2, v, w]$ containing the edge $[v, w]$ [11]; see Figure 4. Now note that

$$\cot \theta_1 + \cot \theta_2 = \frac{\sin(\theta_1 + \theta_2)}{\sin \theta_1 \sin \theta_2},$$

FIGURE 4. Formula for weights λ_{vw}

and so the weight λ_{vw} is nonnegative if and only if $\theta_1 + \theta_2 \leq \pi$. If the quadrilateral $Q = [v_1, v, w] \cup [v_2, v, w]$ is convex, it follows from the well-known empty circumcircle property of Delaunay triangulations [17], that λ_{vw} will be nonnegative if and only if Q satisfies the Delaunay criterion. Since $\theta_1 + \theta_2$ is clearly less than π when Q is not convex, we conclude as follows.

Proposition 3.2. *The piecewise linear finite element approximation to Laplace's equation is a convex combination function if and only if \mathcal{T} is a Delaunay triangulation and no quadrilateral of \mathcal{T} has cocircular vertices.*

The formula in (3.5) also gives rise to convex combination mappings under the same assumption as in Proposition 3.2. Indeed the same formula was derived in [15] and in a different form in [6] from minimizing (3.3) over piecewise linear mappings from \mathbb{R}^2 to \mathbb{R}^3 and from \mathbb{R}^3 to \mathbb{R}^2 , respectively.

An alternative approach to constructing convex combination maps for use in parameterizing surfaces was proposed in [9], where, instead of attempting to approximate harmonic maps, sets of weights λ_{vw} were explicitly constructed which are not only positive and sum to one but also have “linear precision”, i.e., satisfy

$$v = \sum_{w \in N_v} \lambda_{vw} w.$$

The next proposition was the idea behind the method for morphing triangulations in [10].

Proposition 3.3. *Every one-to-one piecewise linear mapping $\phi = (u, v) : D_{\mathcal{T}} \rightarrow \mathbb{R}^2$ is a convex combination mapping.*

Proof. If ϕ is one-to-one, then for any interior vertex v of \mathcal{T} , the restriction of ϕ to $D_v = \bigcup_{T \in \mathcal{T}} T$, of v is also one-to-one. Thus it not difficult to see that the edges containing v , ordered around v , are mapped by ϕ to an ordered sequence of edges containing $\phi(v)$ and the angles between them sum to 2π . So the neighbours w of v are mapped to the vertices $\phi(w)$ of a star-shaped polygon containing $\phi(v)$ in its kernel. Thus positive weights λ_{vw} which sum to one can be found such that $\phi(v)$ can be expressed as a linear combination (3.2). A construction of such weights was given in [9]. \square

An important tool in our analysis will be the following (strong) discrete maximum principle.

Discrete Maximum Principle. *Let f be a convex combination function over a triangulation \mathcal{T} . For any interior vertex v_0 of \mathcal{T} , let V_0 denote the set of all boundary vertices which can be connected to v_0 by an interior path. If $f(v_0) \geq f(v)$ for all $v \in V_0$, then $f(v) = f(v_0)$ for all $v \in V_0$.*

Proof. Let W be the set of vertices of \mathcal{T} , including v_0 itself, which can be connected to v_0 by an interior path. Let v_1 be an interior vertex of W at which f attains its maximum M over W . Such a vertex must exist since $V_0 \subset W$ and we have $M = f(v_1) \geq f(v_0) \geq f(v)$ for all $v \in V_0$. Since v_1 can be connected to v_0 by an interior path in W , it can also be connected to any vertex v in V_0 by an interior path

$$v_1, v_2, \dots, v_k = v,$$

in W . From (3.2) it follows that $f(y) = M$ for all neighbours y of v_1 , and in particular $f(v_2) = M$. Continuing in this fashion, we see that $f(v_3) = M$ and so on. Thus $f(v_k) = M$, which implies $f(v) = M = f(v_0)$. \square

Note that this implies that the maximum of a convex combination function is attained at a boundary vertex. This is the “weak” form of the maximum principle.

4. LOCAL INJECTIVITY

We will establish the following property of convex combination mappings.

Theorem 4.1. *Suppose \mathcal{T} is a strongly connected triangulation and that $\phi : D_{\mathcal{T}} \rightarrow \mathbb{R}^2$ is a convex combination mapping which maps $\partial D_{\mathcal{T}}$ homeomorphically into the boundary $\partial\Omega$ of some (closed) convex region $\Omega \subset \mathbb{R}^2$. Then ϕ is one-to-one.*

Our approach to proving the theorem is to follow the main idea of Kneser’s proof of the Radó-Kneser-Choquet theorem, which is to first show that the harmonic mapping $\phi = (u, v)$ is locally one-to-one, in the sense that the Jacobian

$$\begin{vmatrix} u_x & u_y \\ v_x & v_y \end{vmatrix}$$

never vanishes. In the discrete case, we show initially that ϕ is locally one-to-one in the sense that it is one-to-one on each quadrilateral in \mathcal{T} . By a *quadrilateral* we understand the union $Q = T_1 \cup T_2$ of any two triangles T_1 and T_2 sharing a common edge.

Theorem 4.2. *If \mathcal{T} and ϕ are as in Theorem 4.1, then ϕ is locally one-to-one in the sense that $\phi|_Q$ is one-to-one for every quadrilateral Q in \mathcal{T} .*

We start to prove Theorem 4.2 with the following lemma.

Lemma 4.3. *If \mathcal{T} and ϕ are as in Theorem 4.1, then for every interior vertex v of \mathcal{T} , the point $\phi(v)$ lies in the interior of Ω .*

Proof. It is sufficient to show that $\phi(v)$ lies to the left of each edge of $\partial\Omega$, w.r.t. the anticlockwise direction around $\partial\Omega$. Suppose then that $[v_1, v_2]$ is an edge of $D_{\mathcal{T}}$ and let $ax_1 + bx_2 + c = 0$ be the equation of the infinite line L passing through the mapped edge $[\phi(v_1), \phi(v_2)]$. Then the function $h : D_{\mathcal{T}} \rightarrow \mathbb{R}$ defined by

$$h(x) = a\phi_1(x) + b\phi_2(x) + c$$




FIGURE 5. Mapping a quadrilateral

is a convex combination function satisfying

$$h(v_1) = h(v_2) = 0,$$

and we can assume that $h(w) \geq 0$ for all boundary vertices of \mathcal{T} . Since $\partial\Omega$ is not contained in L , there must be at least one boundary vertex w of \mathcal{T} such that $h(w) > 0$. Since \mathcal{T} is strongly connected, w can be connected to v by an interior path. Therefore, we cannot have $h(v) \leq 0$, for this would contradict the Discrete Maximum Principle (or rather, the analogous Discrete Minimum Principle) applied to h . Thus $h(v) > 0$ as required. \square

From Lemma 4.3 it follows that every triangle T of \mathcal{T} containing a boundary edge of \mathcal{T} is mapped by ϕ to a nondegenerate triangle $\phi(T)$, i.e., $\phi|_T$ is one-to-one. Indeed, if all three vertices of T are boundary vertices, we must have $\mathcal{T} = \{T\}$, for otherwise one of the edges of T would be a dividing edge of \mathcal{T} , contrary to the assumption that \mathcal{T} is strongly connected. Thus in this case $\phi(T)$ is trivially nondegenerate. Otherwise, $T = [v_1, v_2, v_3]$, where $[v_1, v_2]$ is a boundary edge of \mathcal{T} , and v_3 is an interior vertex of \mathcal{T} . Then from Lemma 4.3, $\phi(v_3)$ cannot be collinear with $\phi(v_1)$ and $\phi(v_2)$.

Lemma 4.4. *Let \mathcal{T} and ϕ be as in Theorem 4.1 and suppose $T_1 \cup T_2$ is a quadrilateral in \mathcal{T} . If $\phi|_{T_1}$ is one-to-one, then $\phi|_{(T_1 \cup T_2)}$ is one-to-one.*

Proof. Suppose the triangles are $T_1 = [v_1, v_2, v_3]$ and $T_2 = [v_2, v_3, v_4]$ with common edge $[v_2, v_3]$, as in Figure 5. If $\phi(T_1)$ is nondegenerate, the point $\phi(v_1)$ lies on one side of the infinite straight line L passing through the points $\phi(v_2)$ and $\phi(v_3)$. Our task is to show that the point $\phi(v_4)$ lies on the side of L opposite to $\phi(v_1)$. We will suppose the contrary, namely that $\phi(v_4)$ either lies on the same side as $\phi(v_1)$, as in Figure 5, or lies on L itself. Let $ax_1 + bx_2 + c = 0$ be the equation of L and define the function $h : D_{\mathcal{T}} \rightarrow \mathbb{R}$ as

$$h(x) = a\phi_1(x) + b\phi_2(x) + c.$$

Then h is a convex combination function satisfying

$$h(v_2) = h(v_3) = 0,$$

and we may assume without loss of generality that

$$h(v_1) > 0 \quad \text{and} \quad h(v_4) \geq 0.$$

Next we will construct a path, called a *rising* path, from v_1 to a boundary vertex. First we notice that if v_1 is an interior vertex, it must have some neighbour, say w_1 , such that $h(w_1) > h(v_1)$. This follows from (3.2) and the fact that $h(v_1) > h(v_2)$.

By a similar argument, since $h(w_1) > h(v_1)$, if the vertex w_1 is not a boundary vertex, it must have a neighbour, say x_1 , such that $h(x_1) > h(w_1)$. Continuing in this fashion, we obtain a sequence of interior vertices along which h is strictly increasing. Since each new vertex in the sequence must be distinct from all the previous ones, the sequence must eventually terminate at a boundary vertex, say u_1 . We have therefore constructed a path P_1 which connects v_1 to u_1 . Since h is strictly increasing along P_1 , we have that $h(u_1) > 0$. Following the terminology of Tutte [20], we call P_1 a *rising path* with respect to h . Note that in the case that v_1 is itself a boundary vertex, we have $u_1 = v_1$ and P_1 is the null path consisting of just the one vertex v_1 .

In an analogous manner, we can also construct a falling path P_2 from v_2 to some boundary vertex u_2 , due to the fact that $h(v_2) < h(v_1)$. Along this falling path, h is strictly decreasing, and in particular $h(u_2) < 0$ if v_2 is an interior vertex and $u_2 = v_2$ and $h(u_2) = 0$ if v_2 is a boundary vertex. Analogously to the construction of P_2 , we can construct a falling path P_3 from v_3 to a boundary vertex u_3 with similar properties.

Next observe that the rising path P_1 is clearly distinct to both the falling paths P_2 and P_3 . However, it is possible that P_2 and P_3 meet at some common vertex. We will show that this cannot happen. For if w is the first vertex P_2 and P_3 have in common, starting from v_2 and v_3 , respectively, then the subpath Q_2 of P_2 from v_2 to w and the subpath Q_3 of P_3 from v_3 to w together with the edge $[v_2, v_3]$ form a closed path, say Q . The triangles inside Q form a subtriangulation \tilde{T} of \mathcal{T} . The closed path Q cannot enclose v_1 for then Q would have to cross P_1 . Thus Q must enclose v_4 . But since $h(v_4) \geq 0$, the Discrete Maximum Principle for h over \tilde{T} yields a contradiction. This is because by Lemma 2.2, v_4 must be connected to some vertex v in Q other than v_2 and v_3 and due to the construction of the falling paths, we must have $h(v) < 0$.

We have now established that the three paths P_1, P_2, P_3 are distinct, and therefore u_1, u_2, u_3 are distinct and we have $h(u_1) > 0$, $h(u_2) \leq 0$, and $h(u_3) \leq 0$. Note that due to the assumption that \mathcal{T} is strongly connected, the interior edge $[v_2, v_3]$ cannot be a dividing edge, which means that at least one of its endpoints, say v_2 , is an interior vertex. Thus we can assume that $u_2 \neq v_2$ and $h(u_2) < 0$.

Now let Q be the closed path consisting of the paths P_2 and P_3 , the edge $[v_2, v_3]$, and the part of the boundary of \mathcal{T} connecting u_2 and u_3 which does not contain u_1 . Due to the convexity of the image of $\partial D_{\mathcal{T}}$, and since $h(u_2) < 0$ and $h(u_3) \leq 0$, we have $h(v) < 0$ for all boundary vertices v of \mathcal{T} in Q other than u_2 and u_3 . Therefore, $h(v) < 0$ for every vertex v in Q other than v_2 and v_3 . Since the vertex v_4 is enclosed by Q and $h(v_4) \geq 0$, the Discrete Maximum Principle for h inside Q yields a contradiction. \square

Proof of Theorem 4.2. Due to the connectedness of $D_{\mathcal{T}}$ implied by Definition 2.1, any two triangles in \mathcal{T} can be connected by a path of triangles, i.e., a sequence of triangles

$$T_1, T_2, \dots, T_k,$$

in which T_i and T_{i+1} share a common edge for each i . Thus any given triangle T_1 of \mathcal{T} can be connected in this way to any triangle T_k which contains a boundary edge of \mathcal{T} . By Lemma 4.3, $\phi(T_k)$ is nondegenerate. Then, by Lemma 4.4, $\phi(T_{k-1})$ is also nondegenerate, and likewise $\phi(T_{k-2})$ and so on, and so $\phi(T_1)$ is also nondegenerate, which means that ϕ is one-to-one on every triangle in \mathcal{T} . It follows from Lemma 4.4 that ϕ is one-to-one on every quadrilateral in \mathcal{T} . \square

5. GLOBAL INJECTIVITY

In Theorem 4.2 we established that ϕ in Theorem 4.1 is locally one-to-one. We now complete Theorem 4.1 by showing that global injectivity follows from local injectivity. We assume in Lemmas 5.1 to 5.4 that \mathcal{T} and ϕ are as in Theorem 4.1.

Lemma 5.1. *For any two distinct triangles T and S in \mathcal{T} , the triangles $\phi(T)$ and $\phi(S)$ have disjoint interiors.*

Proof. If $x \in \text{Int}(\phi(T)) \cap \text{Int}(\phi(S))$, then $x \in \text{Int}(\Omega)$ by Lemma 4.1. Let L be any semi-infinite line starting at x which does not pass through any point $\phi(v)$, $v \in V$. The last mapped edge $\phi(e)$ crossed by L must be the image of some unique boundary edge e of \mathcal{T} . Set $T_1 = T$ and let e_1 be the unique edge of T_1 such that $\phi(e_1)$ and L intersect. If e_1 is not a boundary edge of \mathcal{T} , let T_2 be the triangle in \mathcal{T} sharing the edge e_1 and let e_2 be the other edge of T_2 for which $\phi(e_2)$ and L intersect. Since ϕ is one-to-one on $T_1 \cup T_2$, the intersection point $\phi(e_2) \cap L$ is further from x than the intersection point $\phi(e_1) \cap L$. We continue in this way until the next edge in the sequence is e . We thus obtain a unique sequence of distinct triangles and distinct edges

$$T = T_1, T_2, \dots, T_k, \quad \text{and} \quad e_1, e_2, \dots, e_k = e.$$

The edges intersect L in an ordered sequence of distinct points.

In an identical way, we also construct a unique sequence of distinct triangles and distinct edges

$$S = S_1, S_2, \dots, S_\ell, \quad \text{and} \quad f_1, f_2, \dots, f_\ell = e.$$

It now follows that $T_k = S_\ell$. We then deduce that $e_{k-1} = f_{\ell-1}$ and $T_{k-1} = S_{\ell-1}$ and so on. Suppose now that $k = \ell$. Then we conclude that $T = S$, which is a contradiction. Otherwise, we may assume that $k > \ell$, but then we have $T_{k-\ell+1} = S$, which is impossible since $T_{k-\ell+1}$ does not contain x . \square

Lemma 5.2. *If two triangles T_1 and T_2 in \mathcal{T} do not share an edge, then the intersection $\phi(T_1) \cap \phi(T_2)$ is either empty or a point $\phi(v_1) = \phi(v_2)$ for some vertices $v_1 \in T_1$ and $v_2 \in T_2$.*

Proof. Suppose not. By Lemma 5.1, $\phi(T_1) \cap \phi(T_2)$ contains a point x in the relative interior of $\phi(e)$ for some edge e in T_1 or T_2 . We may assume that $e \subset T_1$. Observe that e cannot be a boundary edge. Thus there is a second triangle T_3 sharing e with T_1 and $T_3 \neq T_2$. By Theorem 4.2 applied to T_1 and T_3 , we see that the interiors of the two triangles $\phi(T_3)$ and $\phi(T_2)$ must intersect, which contradicts Lemma 5.1. \square

Lemma 5.3. *If the intersection of two triangles T_1 and T_2 in \mathcal{T} is a common vertex v , then the intersection of the two triangles $\phi(T_1)$ and $\phi(T_2)$ is the common vertex $\phi(v)$.*

Proof. By Lemmas 5.1 and 5.2, the intersection between $\phi(T_1)$ and $\phi(T_2)$ is either empty or a common vertex. Since v belongs to both T_1 and T_2 , the intersection of $\phi(T_1)$ and $\phi(T_2)$ must therefore be $\phi(v)$. \square

Lemma 5.4. *If the intersection of two triangles T_1 and T_2 in \mathcal{T} is empty, then the intersection of the two triangles $\phi(T_1)$ and $\phi(T_2)$ is also empty.*

Proof. Suppose not. Then by Lemmas 5.1 and 5.2, the triangles $\phi(T_1)$ and $\phi(T_2)$ intersect in a common vertex $\phi(v_1) = \phi(v_2)$ for some vertices $v_1 \in T_1$ and $v_2 \in T_2$. Clearly, both v_1 and v_2 are interior vertices of \mathcal{T} . Let S_1, S_2, \dots, S_m be the triangles in \mathcal{T} containing v_1 in some cyclic order around v_1 starting with $S_1 = T_1$. All the triangles S_1, \dots, S_m are distinct from T_2 since T_1 and T_2 are disjoint. By Theorem 4.2 applied to each pair S_i and S_{i+1} , there is some $i \in \{2, \dots, m\}$ such that the interiors of $\phi(S_i)$ and $\phi(T_2)$ intersect, which, due to Lemma 5.1, is a contradiction. \square

Proof of Theorem 4.1. Let x_1 and x_2 be distinct points in $D_{\mathcal{T}}$. If they belong to a common triangle T then Theorem 4.2 shows that $\phi(x_1) \neq \phi(x_2)$. Otherwise $x_1 \in T_1$ and $x_2 \in T_2$ for distinct triangles T_1 and T_2 . But by Theorem 4.2, and Lemmas 5.3 and 5.4, ϕ is one-to-one on $T_1 \cup T_2$ and so $\phi(x_1) \neq \phi(x_2)$. \square

6. GENERAL TRIANGULATIONS

The condition of Theorem 4.1 demands that \mathcal{T} contains no dividing edges. Yet in applications such as geometric modelling, we often encounter triangulations containing such edges. In order to continue to ensure the injectivity of ϕ in Theorem 4.1 we must then place a mild restriction on how the boundary vertices of \mathcal{T} are mapped to $\partial\Omega$.

Theorem 6.1. *Suppose \mathcal{T} is any triangulation and that $\phi : D_{\mathcal{T}} \rightarrow \mathbb{R}^2$ is a convex combination mapping which maps ∂D homeomorphically into the boundary $\partial\Omega$ of some convex region $\Omega \subset \mathbb{R}^2$. Then ϕ is one-to-one if and only if no dividing edge $[v, w]$ of \mathcal{T} is mapped by ϕ into $\partial\Omega$.*

Proof. For the necessity part, observe that every dividing edge $[v, w]$ partitions \mathcal{T} into two disjoint subtriangulations \mathcal{T}_1 and \mathcal{T}_2 . Thus, due to the convexity of $\partial\Omega$, if the line segment $[\phi(v), \phi(w)]$ is contained in $\partial\Omega$, all the boundary vertices of one of the two subtriangulations, say \mathcal{T}_1 , are mapped by ϕ into the line segment $[\phi(v), \phi(w)]$, and so ϕ is clearly not one-to-one on $D_{\mathcal{T}_1}$.

Now we prove that the condition is sufficient. If \mathcal{T} contains no dividing edges, then Theorem 4.1 guarantees that ϕ is one-to-one. Otherwise, the dividing edges of \mathcal{T} partition \mathcal{T} into a set of disjoint subtriangulations $\mathcal{T}_1, \dots, \mathcal{T}_k$, $k \geq 2$, none of which has a dividing edge. The restriction ϕ_i of ϕ to $D_{\mathcal{T}_i}$ is a convex combination mapping. Thus, if we can show that the boundary of \mathcal{T}_i is mapped by ϕ_i to a convex closed curve, Theorem 4.1 applied to ϕ_i will ensure that ϕ_i is one-to-one. But since the cyclically ordered boundary vertices of \mathcal{T}_i are a subsequence of the cyclically ordered boundary vertices of \mathcal{T} , it is sufficient simply to show that the boundary of \mathcal{T}_i is not mapped by ϕ_i to a straight line. To this end, it is clear that at least one boundary edge $[v, w]$ of \mathcal{T}_i is a dividing edge of \mathcal{T} . Since, by hypothesis, $[\phi(v), \phi(w)]$ is not contained in $\partial\Omega$, no boundary vertex u of \mathcal{T}_i other than v and w is mapped by ϕ into the straight line passing through $\phi(v)$ and $\phi(w)$. But \mathcal{T}_i must contain at least one boundary vertex other than v and w . \square

Note that the assumption in Theorems 4.1 and 6.1 that $\partial D_{\mathcal{T}}$ is mapped homeomorphically to $\partial\Omega$ implies that a subsequence of m of the cyclically ordered boundary vertices v_1, \dots, v_n of \mathcal{T} are mapped to the cyclically ordered vertices of an m -sided convex polygon. The condition on the dividing edges in Theorem 6.1 simply says that no dividing edge must be mapped wholly into one side of the polygon. In the special case that $m = n$, this can never happen, for none of the mapped

boundary vertices are then collinear. Thus we deduce the following corollary of Theorem 6.1.

Corollary 6.2. *Suppose \mathcal{T} is any triangulation and that $\phi : D_{\mathcal{T}} \rightarrow \mathbb{R}^2$ is a convex combination mapping which maps the the cyclically ordered boundary vertices v_1, \dots, v_n of \mathcal{T} to the cyclically ordered vertices $\phi(v_1), \dots, \phi(v_n)$ of an n -sided convex polygon. Then ϕ is one-to-one.*

This stronger assumption on how the boundary is mapped was used by Tutte in [20]. In this sense, Theorems 4.1 and 6.1 are improvements over Tutte's results.

REFERENCES

- [1] D. Braess, *Finite elements*, Cambridge University Press, Cambridge, 1997. MR **98f**:65002
- [2] J. E. Castillo (ed.), *Mathematical Aspects of Numerical Grid Generation*, SIAM, 1991. MR **92h**:65148
- [3] G. Choquet, *Sur un type de transformation analytique généralisant la représentation conforme et définie au moyen de fonctions harmoniques*, Bull. Sci. Math. **69** (1945), 156–165. MR **8**:93a
- [4] T. Duchamp, A. Certain, A. DeRose and W. Stuetzle, *Hierarchical computation of PL harmonic embeddings*, vol. preprint, 1997.
- [5] P. Duren and W. Hengartner, *Harmonic mappings of multiply connected domains*, Pac. J. Math. **180** (1997), 201–220. MR **99f**:30051
- [6] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle, *Multiresolution analysis of arbitrary meshes*, Computer Graphics Proceedings, SIGGRAPH 95 (1995), 173–182.
- [7] I. Fáry, *On straight line representations of planar graphs*, Acta Sci. Math. Szeged **11** (1948), 229–233. MR **10**:136f
- [8] D. Field, *Laplacian smoothing and Delaunay triangulation*, Comm. Num. Meth. Eng. **4** (1988), 709–712.
- [9] M. S. Floater, *Parametrization and smooth approximation of surface triangulations*, Comp. Aided Geom. Design **14** (1997), 231–250. MR **98a**:65018
- [10] M. S. Floater and C. Gotsman, *How to morph tilings injectively*, J. Comp. Appl. Math. **101** (1999), 117–129. MR **99k**:52032
- [11] A. Iserles, *A first course in numerical analysis of differential equations*, Cambridge University Press, Cambridge, 1996. MR **97m**:65003
- [12] H. Kneser, *Lösung der Aufgabe 41*, Jahresber. Deutsch. Math.-Verien. **35** (1926), 123–124.
- [13] A. W. F. Lee, W. Sweldens, P. Schröder, L. Cowsar, and D. Dobkin, *MAPS: Multiresolution adaptive parameterization of surfaces*, Computer Graphics Proceedings, SIGGRAPH 98 (1998), 95–104.
- [14] C. W. Mastin and J. F. Thompson, *Elliptic systems and numerical transformations*, J. Math. Anal. Appl. **62** (1978), 52–62. MR **80i**:65134
- [15] U. Pinkall and K. Polthier, *Computing Discrete Minimal Surfaces and Their Conjugates*, Exp. Math. **2** (1993), 15–36. MR **94j**:53009
- [16] E. Praun, W. Sweldens, and P. Schröder, *Consistent mesh parameterizations*, Computer Graphics Proceedings, SIGGRAPH 01 (2001), 179–184.
- [17] F. P. Preparata and M. I. Shamos, *Computational geometry*, Springer, New York, 1985. MR **87d**:68102
- [18] M. H. Protter and H. F. Weinberger, *Maximum principles in differential equations*, Springer, New York, 1984. MR **86f**:35034
- [19] T. Radó, *Aufgabe 41*, Jahresber. Deutsch. Math.-Verien. **35** (1926), 49.
- [20] W. T. Tutte, *How to draw a graph*, Proc. London Math. Soc. **13** (1963), 743–768. MR **28**:1610
- [21] R. S. Varga, *Matrix iterative analysis*, Springer, New York, 2000. MR **2001g**:65002

SINTEF APPLIED MATHEMATICS, P.O. BOX 124 BLINDERN, 0314 OSLO, NORWAY
E-mail address: mif@math.sintef.no