Numerical Algorithms 13 (1996) 201–223

Does contraction preserve triangular meshes?

P. Ciarlet, Jr

CEA/LV, F-94195 Villeneuve St Georges Cedex, France

F. Lamour*

CISI, Branche CEA-Défense, BP 28, F-91192 Gif sur Yvette Cedex, France

Received 7 February 1996
Communicated by C. Brezinski

A triangular graph is a planar graph in which each face is a 3-cycle, except possibly for the exterior face, and without articulation nodes. The embedding of a triangular graph in the plane is called a triangular mesh. More generally, a triangular graph with multiple contours is a planar graph without articulation nodes in which each face is a 3-cycle, except possibly for a fixed number of them. A contraction along an edge in a graph is the result of identifying the two endpoints of the edge. In this paper a necessary and sufficient condition is shown for which triangularity (possibly with multiple contours) is preserved after contraction. Moreover, when a licit contraction is performed, the question to answer is whether or not it is possible to derive the embedding of the contracted triangular graph from the original triangular mesh by redrawing only around the contraction zone.

Introduction

The contraction of an edge in a graph is the process which consists in merging its two endpoints in a single node while preserving the neighboring relations attached to them. The contraction of edges in a graph is a problem that one is faced with whenever one needs to derive smaller graphs from the original one. When a graph has some special properties such as planarity, triangularity, ..., it can be interesting to look for a contracted graph with the same characteristics. In this paper, our interest is focused on triangular meshes, that is, the result of meshing a plane region with triangles. In particular our goal is to coarsen meshes for the purpose of multigrid methods or domain decomposition methods in numerical analysis. In that scope, contracted meshes are required to be also triangular meshes and to be as close as possible to the original mesh.

* Part of this work was done when the author was in his sabbatical year at the University of California at Los Angeles.

In the following, we give first a necessary and sufficient condition that defines what we call a licit contraction, that is, the condition under which an edge of a triangular graph can be contracted such that the resulting contracted graph remains a triangular graph. In the second place, we address the problem of drawing a triangular graph in the plane with straight non-intersecting edges. Such a representation, called a triangular mesh, always exists, but an interesting question is to know if, being given a triangular mesh, the contracted triangular mesh can be easily deduced from it.

In the last part, complexity aspects of the edge contraction process together with the embedding of the contracted triangular graph are discussed. Finally, a generalization of this plane embedding problem to other surfaces is presented.

1. Preliminaries

This section summarizes some definitions and results on planar graphs. For more details on graph theory and planarity we refer the reader to [5] and [8].

Let G be a graph with node set V and edge set E, without loops or multiple edges. The *order* of G, denoted n, is the number of nodes in V. The number of edges is denoted m. The set of neighbors of a node v is denoted $\Gamma(v)$ and the degree of v, denoted $d(v)$, is the number of its neighbors.

A graph G is *planar* if it has an embedding in the plane, i.e., G can be drawn in such a way that no two edges intersect except at a common endpoint. Moreover, for every planar graph there exists an embedding such that every edge is a straight line segment (see, for example, [2, theorem 5.1]).

A *face* of an embedded planar graph is a portion of the plane bounded by edges and such that any two points in this region can always be linked by a continuous line that never encounters either a node or an edge. Note that one face is unbounded: it is called the *exterior* face. All other faces are called *interior* faces.

We now call *contour* of the graph the set of edges that delimits the exterior face. Its edges are then called *contour edges* while others are called *interior edges*. In the same way, nodes that define the contour of the graph are *contour nodes* while others are *interior nodes*.

Euler (1736) first established that the numbers of nodes (n), edges (m) and interior faces (f) of a planar graph are related by

$$n - m + f = 1. \qquad \text{(Euler)}$$

On the other hand, if we count the edges around each interior face we obtain another relation which concerns m and the number of contour edges m_c. The relationship is that if $m_{\mathcal{F}}$ is exactly the number of edges of a given interior face $\mathcal{F}$ then

$$\sum_{\mathcal{F}} m_{\mathcal{F}} = 2m - m_c. \qquad \text{(Faces)}$$

A consequence of the two relations is that an upper bound on the number of edges, for any planar graph of order greater than or equal to 3, is $3n-6$. Moreover, if

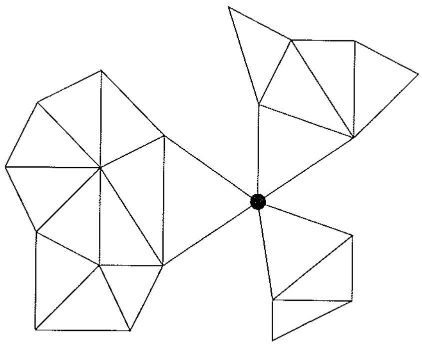

Figure 1. Articulation node with multiplicity 2.

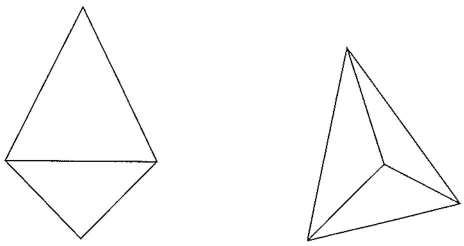

(a) (b)

Figure 2. (a) A triangular graph with 4 nodes ($n_c = 4$). (b) A triangulation with 4 nodes.

G is a *maximal planar* graph, that is, a planar graph for which it is impossible to add more edges without losing the planarity, then the number of edges is exactly $3n - 6$. Indeed, a maximal planar graph has all faces (including the exterior one) bounded by 3-cycles. Because 3-cycles are triangles in a straight line segments representation, the term triangle usually replaces the term 3-cycle. In the same way maximal planar graphs are also called *triangulations*.

Now it is possible to define another useful class of planar graphs. Let n_c be the number of contour nodes. Among them, we have ordinary contour nodes (endpoints of exactly two contour edges) and articulation nodes (endpoints of $2k$ contour edges, with $k \geqslant 2$). Let n_a be the number of articulation nodes, including their multiplicity ($= (k-1)$). For instance, the articulation node in figure 1 has multiplicity 2 ($n_c = 19$ and $n_a = 2$). Then it is straightforward to check that $m_c = n_c + n_a$.

A *triangular graph* is a planar graph in which each face is a triangle, except the exterior face and without any articulation node ($n_a = 0$). According to this definition a triangulation is a particular triangular graph (figure 2). We have the following result:

Lemma 1. A planar graph G with n_c contour nodes is a triangular graph if and only if

$$m = 3n - n_c - 3. \qquad \text{(Triangular)}$$

Proof. If G is a triangular graph, then $m_{\mathcal{F}} = 3$ for all interior faces. Thus (Faces) yields $3f = 2m - m_c$. As $n_a = 0$ by assumption, we deduce that $3f = 2m - n_c$. Now, using (Euler) leads to the result.

Conversely, suppose G is a planar graph such that (Triangular) holds. From (Faces), we infer that $3f \leqslant 2m - m_c$. If the inequality is sharp then, using the value of f given by the Euler formula and the one of m given by (Triangular), we get that $m_c < n_c$, which is impossible. So $3f = 2m - m_c$, which means that all interior faces in G are triangles. Moreover in that case $m_c = n_c$, that is, $n_a = 0$. Thus G is a triangular graph. □

2. Edge contraction

The *contraction* of a graph G along an edge (u, v) is the graph G' resulting from

(a) the deletion of one endpoint, for instance v, and the edges incident to v,

(b) the insertion of an edge (u, w) for each node $w \in \Gamma(v) - \Gamma(u)$.

This contraction is often called a *simplicial contraction* [4], as no multiple edge is created. An example is given in figure 3 with a contraction along the dark edge.

Before considering the contraction of edges in planar graphs, let us recall that K_5 is the complete graph with 5 nodes and that $K_{3,3}$ is the complete bipartite graph on two sets of three nodes each (see figure 4): both graphs are non-planar. In addition, we have to define *graph homeomorphism*, which stems from the usual notion of homeomorphism in topology: let E and F be two topological spaces, then $f: E \to F$ is an homeomorphism if and only if f and f^{-1} are continuous bijections. Following [4], we note that a graph G can be represented by a topological space in which each node is represented by a given point, and each edge by a distinct arc homeomorphic to the interval $[0, 1]$. Then two graphs are said to be *homeomorphic* if and only if their topological representations are homeomorphic as topological spaces.

An interesting result is that the contraction of a planar graph is also planar. To be convinced one can consider the two following results from [4]:

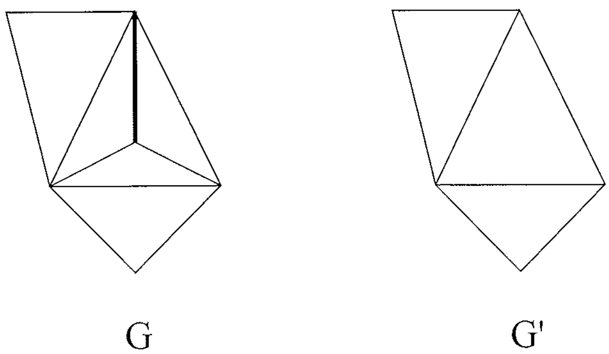

Figure 3. Contraction of one edge.

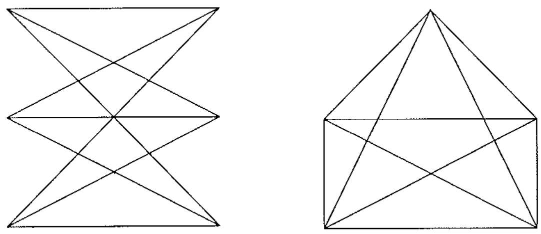

$K_{3,3}$ K_5

Figure 4. Basic non-planar graphs.

Kuratowski's theorem. A graph G is planar if and only if it does not contain any subgraph homeomorphic to K_5 or $K_{3,3}$.

Lemma. Let G be a graph that contains no homeomorph of K_5 or $K_{3,3}$, and let e be any edge in G. Then the result of contracting the graph G on edge e is a graph G' that contains no homeomorph of K_5 or $K_{3,3}$.

Knowing that, one can ask oneself whether contraction preserves other characteristics such as maximality or triangularity.

It has been shown that if the contraction is carried out along some special edges, the maximality of a planar graph is preserved [6, 10]. Indeed, Kampen in [6] defines the *T-contraction* of a triangulation as the contraction of an edge which belongs to two distinct faces (including the exterior one) that leads to a resulting triangulation. The edge that allows a T-contraction is said to be *T-contractable*. He then states the following result: *from each node of a triangulation of order greater than 3, there exists at least two T-contractable edges originating from it.* Let us briefly outline the proof of this result, which is obtained by induction on n, the order of the triangulation. The result is clearly verified for a triangulation of order 4. For a triangulation of order $n > 4$, Kampen first notices that if an edge (p, q) is not T-contractable then p and q have at least 3 common neighbors (see also remark 1), u, v and w, pqu and pqv being two faces. Then, he assumes that u lies in the interior of pqw (so v lies in the exterior). By deleting edges and nodes exterior to pqw, the result is a triangulation (pqw is the exterior face) of order lower than n. The result follows by induction. To get the second T-contractable edge, it suffices to consider the symmetric situation where the triangulation of order lower than n is obtained by deleting edges and nodes interior to pqw.

A precise way of proving that maximality is preserved during contraction is by using the relationship between the numbers of nodes and edges of a graph and its resulting contracted graph together with the relationship between the numbers of nodes and edges of a triangulation.

Clearly, the order of G' in case of a graph G of order n is $n - 1$. As for the number of edges, it can be determined by:

Lemma 2. Let G be a graph with m edges. G' has $m - (1 + d_{\mathrm{cn}})$ edges where d_{cn} is the number of common neighbors of the endpoints of the contracted edge.

Proof. Let us call u and v the endpoints of the contracted edge. The number of edges in G having u or v as an endpoint is equal to $d(u) + d(v) - 1$. Let d_{cn} be the number of common neighbors of u and v, and let $d_{\mathrm{ncn}}(u)$ (resp. $d_{\mathrm{ncn}}(v)$) be the number of neighbors of u (resp. v) that are not neighbors of v (resp. u). Now $d(u)$ (resp. $d(v)$) can be written as $1 + d_{\mathrm{cn}} + d_{\mathrm{ncn}}(u)$ (resp. $1 + d_{\mathrm{cn}} + d_{\mathrm{ncn}}(v)$). As we contract the edge (u, v), the number of edges in G' having u as an endpoint is equal to $d_{\mathrm{ncn}}(u) + d_{\mathrm{ncn}}(v) + d_{\mathrm{cn}} = d(u) - 1 + d(v) - 1 - d_{\mathrm{cn}}$. Then, the total number of edges in G' is equal to $m - 1 - d_{\mathrm{cn}}$. □

Consequently:

Lemma 3. Let G be a triangulation. G' is a triangulation if and only if the endpoints of the contracted edge have exactly 2 common neighbors.

Proof. G' is a triangulation leads to $m' = 3n' - 6 = 3(n - 1) - 6$. At the same time, as G is a triangulation then $m' = m - 3$ because $3n = m + 6$. Therefore as $m' = m - (1 + d_{\mathrm{cn}})$, by identification we deduce that $d_{\mathrm{cn}} = 2$.

Conversely G' being the result of the contraction of an edge of a planar graph is also planar. On the other hand, if the endpoints of the contracted edge have exactly 2 common neighbors then we deduce $m' = m - 3$ from lemma 2. As G is a triangulation $m = 3n - 6$ and then $m' = 3n - 6 - 3 = 3(n - 1) - 6 = 3n' - 6$ which proves that G' is a maximal planar graph, i.e., a triangulation. □

Remark 1. In particular, this proves Kampen's assumption that if an edge of a triangulation is not T-contractable then $d_{\mathrm{cn}} \geqslant 3$. Indeed, the previous lemma asserts that an edge of a triangulation is not T-contractable if and only if $d_{\mathrm{cn}} \neq 2$, but it is a simple matter to check that $d_{\mathrm{cn}} \geqslant 2$ for all the edges of a triangulation of order greater than 3.

According to Kampen's definition of a T-contraction, there is no distinction between edges bordering the exterior face (i.e., on the contour) and interior edges. This is simply explained by the fact that all faces, including the exterior one, are triangles. In the case of triangular graphs, this is no longer true: therefore, contour nodes play a particular role. More precisely, we have to introduce a rule of *contour preservation*.

Let (u, v) be the edge of G along which the contraction is performed and w a node which does not belong to $\{u, v\}$.

- If either u or v is a contour node of G, then u is on the contour of G'.
- If w is a contour (resp. an interior) node of G, then w is a contour (resp. an interior) node of G'.

In particular, a contour node of G other than u and v cannot become an interior node of G' and vice versa: $n'_{\mathrm{c}} = n_{\mathrm{c}} - 1$ if and only if both u and v are contour nodes (and

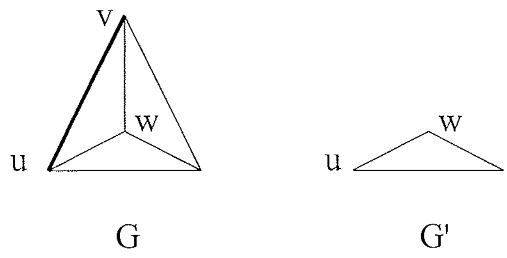

Figure 5. Differing contractibility notions.

$n'_c = n_c$ otherwise). Then, if the contour is reduced to three nodes, the contraction of a contour edge cannot lead to a triangular graph. For instance the contraction of (u, v) in figure 5 leads to a triangulation, but not to a triangular graph, as one would get $n'_c = 2$ or the node w would become a contour node in G'.

Now, the necessary and sufficient condition for preserving the triangularity is given by the following result:

Lemma 4. Let G be a triangular graph of order greater than 3. G' is a triangular graph if and only if one of the two following conditions holds:

- $n'_c = n_c$ and $d_{cn} = 2$,
- $n'_c = n_c - 1$ and $d_{cn} = 1$.

Proof. G and G' being triangular graphs, we have respectively $m = 3n - (n_c + 3)$ and $m' = 3n' - (n'_c + 3)$. From lemma 2 we have that $m - 1 - d_{cn} = 3n' - (n'_c + 3)$ or equivalently $3n - (n_c + 3) - 1 - d_{cn} = 3(n - 1) - (n'_c + 3)$ as $n' = n - 1$. Thus $d_{cn} = (n'_c - n_c) + 2$ and as a result either $d_{cn} = 2$ and $n'_c = n_c$ or $d_{cn} = 1$ and $n'_c = n_c - 1$.

Conversely, from lemma 2 and the fact that G is a triangular graph, we have $m' = m - (1 + d_{cn}) = 3n - n_c - 3 - (1 + d_{cn})$. As $n' = n - 1$, this yields $m' = 3n' - n_c - (1 + d_{cn})$.

G' being a planar graph, formula (Euler) holds together with (Faces) and thus we have $m' \leqslant 3n' - m'_c - 3$. Replacing m'_c by $n'_c + n'_a$ and m' by $3n' - n_c - (1 + d_{cn})$ in the previous inequality we get $n'_c - n_c + n'_a \leqslant d_{cn} - 2$.

Then, case 1 ($n'_c = n_c$ and $d_{cn} = 2$) gives $n'_a \leqslant 0$, from which $n'_a = 0$. As formula (Triangular) holds for G' and it has no articulation node then it is a triangular graph. Case 2 ($n'_c = n_c - 1$ and $d_{cn} = 1$) leads to the same conclusion, which proves the lemma. □

Generalizing Kampen's T-contractability, we say that an edge of a triangular graph G is *t-contractable* if the graph resulting from the contraction of G is also a triangular graph. Then, from the previous lemma one can derive the following corollary.

Corollary. An edge of a triangular graph is not t-contractable if and only if one of the following conditions holds:

- $n'_c = n_c$ and $d_{cn} \geqslant 3$,
- $n'_c = n_c - 1$ and $d_{cn} \geqslant 2$.

From this, we can prove the following lemma which concerns the existence of t-contractable edges in a triangular graph.

Lemma 5. Let G be a triangular graph of order greater than 3 and u one of its nodes.
If u is an interior node, there exist at least two t-contractable edges originating from it.
If u is a contour node, let v and w be its contour neighbors.

- If $(v, w) \in E$, there exists at least one t-contractable edge originating from it.
- If $(v, w) \notin E$, there exist at least two t-contractable edges originating from it.

Proof. Let u be a given node of G. We shall consider two cases, whether u is an interior node or a contour node.

Interior case. Let u be an interior node. We shall prove that when there exists one edge originating from u which is not t-contractable, two regions around u can be defined such that each one of them includes a t-contractable edge having u as one of its endpoints.

u being an interior node, all its incident edges are interior edges, and they belong to exactly two interior faces: triangles by definition. Also u being interior, we have $n'_c = n_c$ and both u and the endpoint of a t-contractable edge originating at u have necessarily (lemma 4) two common neighbors. Let d be the degree of u.

Note that if $d = 3$, then clearly the three edges having u as an endpoint are t-contractable ($n'_c = n_c$ and $d_{cn} = 2$).

If $d \geqslant 4$, let us give an ordering of the neighbors of u and of the faces it belongs to. So let us denote by v_0 one of the neighbors of u and among the two interior faces bordering (u, v_0), let us choose one which is denoted by T_1. Now, for $k = 1$

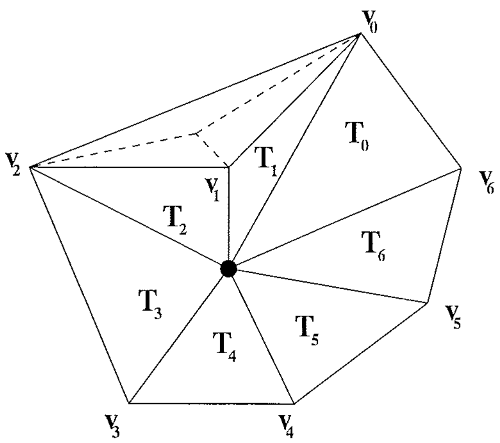

Figure 6. Numbering of neighbors and triangles.

to $k = d-2$, let v_k be the neighbor of u such that $T_k = \{u, v_{k-1}, v_k\}$ and let T_{k+1} be the other triangle bordering (u, v_k). Last, denote v_{d-1} the remaining neighbor (such that $T_{d-1} = \{u, v_{d-2}, v_{d-1}\}$) and T_0 the triangle $\{u, v_{d-1}, v_0\}$. Thus one gets a numbering of neighbors and triangles surrounding u, such that, modulo d, we have $T_k = \{u, v_{k-1}, v_k\}$ for k in $\{0, \dots, d-1\}$ as illustrated by figure 6 with $d = 7$.

Let us note that, by construction, both v_{k-1} and v_{k+1} are common neighbors of v_k (modulo d) for k in $\{0, \dots, d-1\}$.

Now suppose that (u, v_{i_0}) is not t-contractable (we can suppose that such an edge exists otherwise the lemma is proved). If we let $d_{\text{cn}}(u, v_{i_0})$ be the number of common neighbors of u and v_{i_0}, then, according to the corollary, $d_{\text{cn}}(u, v_{i_0}) \geqslant 3$. Thus there exists $j_0 \notin \{i_0 - 1, i_0 + 1\}$ such that v_{j_0} is a neighbor of v_{i_0} ((u, v_{j_0}) is also not t-contractable).

As the graph is planar, among the two interior faces leaning on (u, v_{i_0}) (the triangles T_{i_0} and T_{i_0+1}), one of the two is inside the cycle delimited by (u, v_{i_0}, v_{j_0}) and the other is outside this cycle: we can assume that T_{i_0+1} lies in the interior of the cycle (u, v_{i_0}, v_{j_0}). Thus v_{i_0+1} is also inside the cycle. From there, using again the planarity of the graph, one easily gets step by step that the nodes $\mathcal{C}_0^{\text{int}} = (v_k)_{k=i_0+1}^{k=j_0-1}$ are inside the cycle. By the same argument, it is clear that the nodes $\mathcal{C}_0^{\text{out}} = (v_k)_{k=j_0+1}^{k=i_0-1}$ are outside the cycle. Therefore, there is no edge between these two collections of nodes: $(\mathcal{C}_0^{\text{int}} \times \mathcal{C}_0^{\text{out}}) \cap E = \varnothing$.

From this point, we can prove that there exists at least one node v_k in each of these collections such that (u, v_k) is t-contractable.

Let us consider first the case of the interior nodes $\mathcal{C}_0^{\text{int}}$. Note that we have $|j_0 - i_0| \geqslant 2$, as $j_0 \notin \{i_0 - 1, i_0 + 1\}$. Let $i_1 = i_0 + 1$. If $|j_0 - i_0| = 2$ then $\mathcal{C}_0^{\text{int}}$ reduces to v_{i_1}. As there is no edge between $\mathcal{C}_0^{\text{int}}$ and $\mathcal{C}_0^{\text{out}}$, we get $d_{\text{cn}}(u, v_{i_1}) = 2$ and therefore (u, v_{i_1}) is t-contractable.

Suppose now that $|j_0 - i_0| > 2$. If (u, v_{i_1}) is t-contractable, the assertion is proved. Otherwise, there exists v_{j_1} such that v_{j_1} is a neighbor of v_{i_1}. From $(\mathcal{C}_0^{\text{int}} \times \mathcal{C}_0^{\text{out}}) \cap E = \varnothing$, we deduce that necessarily v_{j_1} is in $\mathcal{C}_0^{\text{int}} \cup \{v_{j_0}\}$. So $|j_0 - i_0| - 1 \geqslant |j_1 - i_1| \geqslant 2$ by construction. Again, we define two collections of nodes $\mathcal{C}_1^{\text{int}} = (v_k)_{k=i_1+1}^{k=j_1-1}$ and $\mathcal{C}_1^{\text{out}} = (v_k)_{k=j_1+1}^{k=j_0-1}$ and we get $(\mathcal{C}_1^{\text{int}} \times \mathcal{C}_1^{\text{out}}) \cap E = \varnothing$ (note that $\mathcal{C}_1^{\text{out}}$ can be empty).

Let $v_{i_2} = v_{i_1+1}$. If $|j_1 - i_1| = 2$, the assertion is proved for (u, v_{i_2}). Otherwise, the process goes on...

At step s, we define $v_s = v_{i_{s-1}+1}$. If $|j_{s-1} - i_{s-1}| = 2$ the assertion is proved for (u, v_s). Otherwise, if (u, v_s) is not t-contractable, we get again $|j_{s-1} - i_{s-1}| - 1 \geqslant |j_s - i_s| \geqslant 2$. By induction, we have $|j_0 - i_0| - s \geqslant |j_s - i_s| \geqslant 2$.

Therefore, as $|j_s - i_s|$ decreases to 2 when s increases, at some step S we obtain that (u, v_{S+1}) is t-contractable, either by direct checking or because $|j_S - i_S| = 2$. This proves the existence of a t-contractable edge inside the cycle (u, v_{i_0}, v_{j_0}).

The proof can be extended straightforwardly to the case of the collection $\mathcal{C}_0^{\text{out}}$, and this shows that at least two t-contractable edges can be exhibited from any interior node.

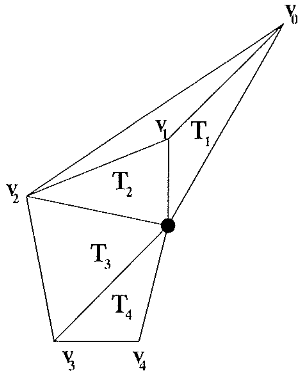

Figure 7. Numbering of neighbors and triangles.

Contour case. Let u be a contour node. We shall prove that when there exists one contour edge originating from u which is not t-contractable, at least one region around u can be defined such that it includes a t-contractable edge having u as one of its endpoints.

u being a contour node, two of its incident edges are contour edges while the other incident edges are interior edges. The former belong to one interior face, and the latter to two interior faces. Let us notice that in this case $n'_c = n_c$ or $n'_c = n_c - 1$. Let d be the degree of u.

Let us note that if $d = 2$, then clearly both edges having u as one of their endpoints are t-contractable.

If $d \geqslant 3$, denote by v_0 one of the contour neighbors of u and denote by T_1 the interior face bordering (u, v_0). Now, for $k = 1$ to $k = d - 2$, let v_k be the neighbor of u such that $T_k = \{u, v_{k-1}, v_k\}$ and T_{k+1} the other triangle bordering (u, v_k). Last, denote v_{d-1} the remaining (contour) neighbor. Thus one gets a numbering of neighbors and triangles surrounding u, such that, modulo d, we have $T_k = \{u, v_{k-1}, v_k\}$ for k in $\{1, \ldots, d-1\}$. See figure 7 with $d = 5$.

If both (u, v_0) and (u, v_{d-1}) are t-contractable, the lemma is proved in this case. Otherwise, let $i_0 = 0$ and suppose that (u, v_{i_0}) is not t-contractable. Then $d_{\mathrm{cn}}(u, v_{i_0}) \geqslant 2$, and there exists j_0 such that (v_{i_0}, v_{j_0}) is an edge. Again, let $\mathcal{C}_0^{\mathrm{int}} = (v_k)_{k=i_0+1}^{k=j_0-1}$ and $\mathcal{C}_0^{\mathrm{out}} = (v_k)_{k=j_0+1}^{k=d-1}$. As before, there is no edge between these two collections of nodes: $(\mathcal{C}_0^{\mathrm{int}} \times \mathcal{C}_0^{\mathrm{out}}) \cap E = \varnothing$. Following the proof given in the first part for $\mathcal{C}_0^{\mathrm{int}}$, one gets by induction that there is at least one t-contractable edge of the form (u, v_k), with $v_k \in \mathcal{C}_0^{\mathrm{int}}$, or equivalently $0 < k < j_0$. When (v_0, v_{d-1}) is an edge, no better result can be obtained because $\mathcal{C}_0^{\mathrm{out}}$ can be empty. Figure 8 shows an example with only one contractable edge.

If (v_0, v_{d-1}) is not an edge, then necessarily $j_0 < d - 1$. Thus $\mathcal{C}_0^{\mathrm{out}}$ cannot be empty in this case (see for instance figure 7). The same induction process can also be carried out on $\mathcal{C}_0^{\mathrm{out}}$, and as a result there exists at least one t-contractable edge of the form (u, v_l), with $v_l \in \mathcal{C}_0^{\mathrm{out}}$, or equivalently $j_0 + 1 < l \leqslant d - 1$. This ends the proof. □

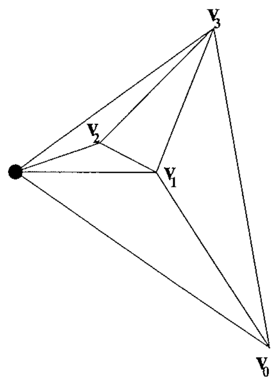

Figure 8. Case of a single t-contractable edge.

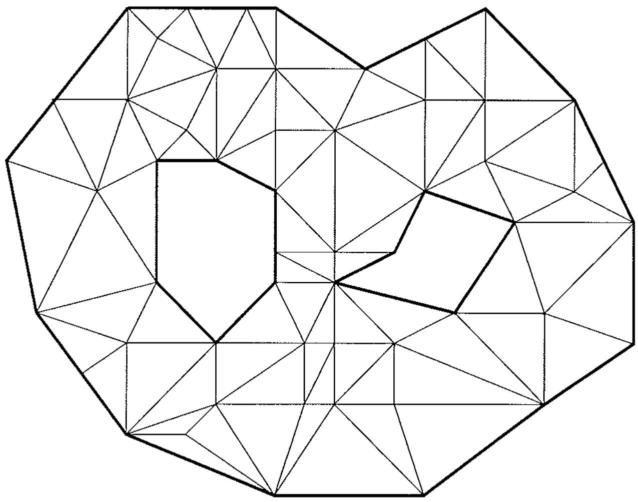

Figure 9. A triangular graph with two holes.

All preceding results can be generalized to the case of a triangular graph with *multiple contours*, i.e., a planar graph without articulation nodes in which each face is a triangle, except possibly for a fixed number of them. Indeed, if we call *holes* the regions of the triangular graph delimited by contours and surrounded by triangles then it is easy to show that formulas (Euler) and (Triangular) can be rewritten with respect to that feature as follows, while formula (Faces) remains unchanged. Note that in this more general case the notion of articulation nodes extends to the holes and $m_c = n_c + n_a$ still holds. Let h denote the number of holes, then

$$n - m + f = 1 - h. \qquad \text{(Euler bis)}$$

A planar graph with n_c contour nodes and h holes is a triangular graph if and only if

$$m = 3n - n_c - 3 + 3h. \qquad \text{(Triangular bis)}$$

As noted earlier, if a contour around a hole reduces to 3 edges (or nodes), then none of its edges is t-contractable. Thus, as the contours are fixed when they are reduced to 3 edges, the number of holes cannot be modified when a t-contraction is performed. From there, it is a simple matter to check that lemmas 2, 4 and 5 remain true.

3. Embedding the contracted graph

In the first place let us recall our goal. From a given triangular mesh, that is, an embedding of a triangular graph, we want, with minimum cost, to derive an embedding of the resulting triangular graph obtained by contracting a t-contractable edge. As we know from the previous section that, from any node in a triangular graph, there exists at least one t-contractable edge, we can always obtain a contracted triangular graph. To construct rapidly an embedding of the contracted triangular graph, it would be pleasant to need to do only local modifications on the original triangular mesh. The following propositions and lemmas lead to characterize a region of the plane in which simple moves of one of the endpoints of the contracted edge produce the triangular contracted mesh.

For the following let us note $\mathrm{Dr}(G)$ the drawing associated to the graph G. Each node u is associated with two coordinates (x_u, y_u) and the location in $\mathrm{Dr}(G)$ of u, called a *point*, is still denoted by u. In addition, two assumptions are made. First, that two distinct nodes cannot be located at the same point. Second, that to a contour (resp. interior) node corresponds a contour (resp. interior) point. An edge (u, v) is represented by a straight line segment $[u, v]$. An *embedding* of a triangular graph G is a plane drawing $\mathrm{Dr}(G)$ where there are no edge crossings nor triangles with a flat angle. As we said earlier, we call such a drawing of a triangular graph a *triangular mesh*.

To begin with, we need a well known result concerning the sign of an angle.

Proposition 1. Let (u, v, w) be three points in the plane such that $u \neq v$, and define θ_w as the angle $(\overrightarrow{uv}, \overrightarrow{uw})$, ($\theta_w \in]-\pi, \pi]$). Let $\mathcal{D}$ be the straight line containing u and v. Let w vary. Then $\sin\theta_w = 0$ iff $w \in \mathcal{D}$. In addition, the sign of θ_w (or of $\sin\theta_w$) is constant in each half-plane delimited by $\mathcal{D}$. In particular, if w spans any segment strictly embedded into one of the half planes, the sign of θ_w is constant.

Proof. See the appendix. □

Let us recall that for an interior point of degree d, the faces $(T_k)_{k=0}^{k=d-1}$ and neighbors $(v_k)_{k=0}^{k=d-1}$ can be labelled in such a way that, modulo d, we have $T_k = \{u, v_{k-1}, v_k\}$ for $k \in \{0, \ldots, d-1\}$. From there, it is possible to define θ_k, for $k = 0, \ldots, d-1$, as the angle $(\overrightarrow{uv_{k-1}}, \overrightarrow{uv_k})$ modulo d, with $\theta_k \in]-\pi, \pi]$ (see figure 10). Similarly, for a contour point of degree d, the faces $(T_k)_{k=1}^{k=d-1}$ and neighbors $(v_k)_{k=0}^{k=d-1}$ are numbered so that, modulo d, $T_k = \{u, v_{k-1}, v_k\}$. And, for $k \in \{1, \ldots, d-1\}$, θ_k can be defined as the angle $(\overrightarrow{uv_{k-1}}, \overrightarrow{uv_k})$ with $\theta_k \in]-\pi, \pi]$.

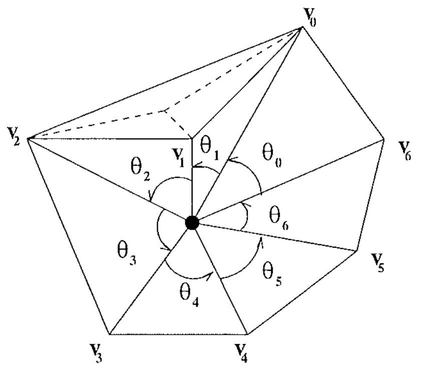

Figure 10. Angles.

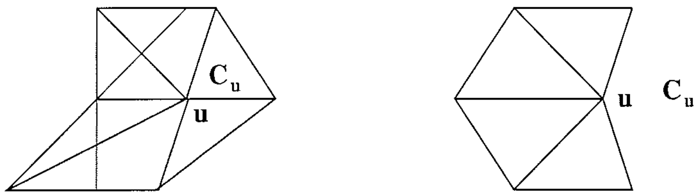

Figure 11. Examples of control regions C_u.

Here, the only difference between the interior and contour cases is that the faces and angles are labelled from 0 to $d-1$ in the former case and from 1 to $d-1$ in the latter. Whenever needed for enumerations, we use $\varepsilon = 0$ if the point is interior and $\varepsilon = 1$ if the point is on the contour.

Remark 2. Evidently, as there is no flat triangle in an embedding, there is no angle θ such that $\sin\theta = 0$.

Proposition 2. Let u be a point of an embedding of G. Then the sign of the angles $(\theta_k)_{k=\varepsilon}^{d-1}$ is constant. In addition, $|\sum_{k=\varepsilon}^{d-1}\theta_k| \leqslant 2\pi$ with equality if and only if u is an interior node.

Proof. See the appendix. □

Remark 3. In our case, an embedding of the initial triangular mesh is given. Therefore, by renumbering locally the neighbors of a point, it is possible to get only positive angles. Indeed, for a point u, if $\theta_\varepsilon > 0$, then all angles are nonnegative, and else it suffices to reorder the neighbors the other way round to obtain again nonnegative angles.

From this point on, it is assumed that, for all points u of a given embedding, $\theta_k > 0$, for all k.

We now call $(R_k)_{k=\varepsilon}^{d-1}$ the (open) half planes delimited by $\mathcal{D}_{v_{k-1},v_k}$ such that u belongs to R_k. Then we can define what we call the *control region of* u as the intersection of all (R_k), $C_u = \bigcap_{k=\varepsilon}^{d-1} R_k$ (see figure 11).

Proposition 3. C_u is a nonempty convex region of the plane defined by

$$u' \in C_u \quad \text{iff} \quad \theta'_k > 0,\ \forall k \in \{\varepsilon, d-1\}, \quad \text{where } \theta'_k = (\overrightarrow{u'v_{k-1}}, \overrightarrow{u'v_k}).$$

In addition, if $u' \in C_u$, then $\sum_{k=\varepsilon}^{d-1} \theta'_k \leqslant 2\pi$.

Proof. See the appendix. □

C_u might be unbounded in the case of a contour point. On the other hand, when u is an interior point, if the polygon delimited by $\{v_0, \dots, v_{d-1}\}$ is convex, then C_u corresponds exactly to its interior.

We want to characterize moves of points along edges of a triangular mesh that preserve embeddings. Assume that a point u is moved to a new location u', we say that the displacement of u from u to u' is *valid* if all drawings made with u spanning $[u, u']$ are embeddings. As a consequence of the previous propositions, we have the following lemma.

Lemma 6. Let G be a triangular mesh, and let u be a point of G. If u is moved to u', then the displacement is valid if and only if $u' \in C_u$.

Proof. If $\mathrm{Dr}'(G)$ is an embedding of G ($\mathrm{Dr}'(G)$ is a drawing of G when u moves to u'), then all angles at u' have the same sign. Moreover, if $\theta'_k < 0$ for all k, then as all θ_k are positive, proposition 1 yields the existence of $u'' \in [u, u']$ such that there exists k_0 such that $\sin\theta''_{k_0} = 0$. Thus $\mathrm{Dr}''(G)$ is not an embedding because flat angles are not allowed in embeddings, and therefore the displacement from u to u' is not valid. This contradicts the assumption. Thus $\theta'_k > 0$, $\forall k$, or, equivalently, u' belongs to C_u.

Conversely, as C_u is convex, it is clear that any u'' which belongs to $[u, u']$ also belongs to C_u. Thus, from proposition 3, we get $\theta''_k > 0$, $\forall k$. Now, as only the location of u is changed, the modifications in the drawing are restricted to u and the edges $(u, v_k)_{k=0}^{d-1}$. Thus it is enough to consider drawings of the graph G_u with nodes $V_u = \{u, v_0, \dots, v_{d-1}\}$ and those edges of G whose endpoints both belong to V_u.

From this point, there remains only to prove that if u is moved to u'' such that $\theta''_k > 0$, $\forall k$, then the drawing $\mathrm{Dr}''(G_u)$ is an embedding. If it is not, then there are at least two segments that intersect. There are two families of edges in G_u, namely $(v_{k-1}, v_k)_k$ and $(u, v_k)_k$. The location of $(v_k)_k$ is unmodified, thus there is no segment crossing between two segments of the first family. If the segments $[u'', v_{k_0}]$ and $[u'', v_{k_1}]$ for $k_0 \neq k_1$ intersect somewhere else than in u'', then $\sum_{k\leqslant k_0} \theta_k \equiv \sum_{k\leqslant k_1} \theta_k$, modulo 2π. This is not possible as $0 < \sum_k \theta_k \leqslant 2\pi$. Finally, there is no segment crossing such that $[v_{k_0-1}, v_{k_0}] \cap [u'', v_{k_1}] \neq \varnothing$ for $k_1 \notin \{k_0 - 1, k_0\}$, as again this would imply $\sum_k \theta_k > 2\pi$ (see figure 12). Therefore the drawing $\mathrm{Dr}''(G_u)$ is an embedding, which proves the lemma. □

In order to characterize the property $u' \in C_u$, a number of general purpose algorithms have been designed: see for instance [11, 12, 7] and references therein.

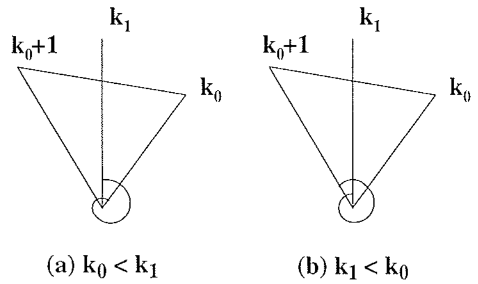

Figure 12. (a) $k_0 < k_1$: $\sum_{k=k_0+1}^{k=k_1} \theta_k > 2\pi$, (b) $k_1 < k_0$: $\sum_{k=k_1+1}^{k=k_0+1} \theta_k > 2\pi$.

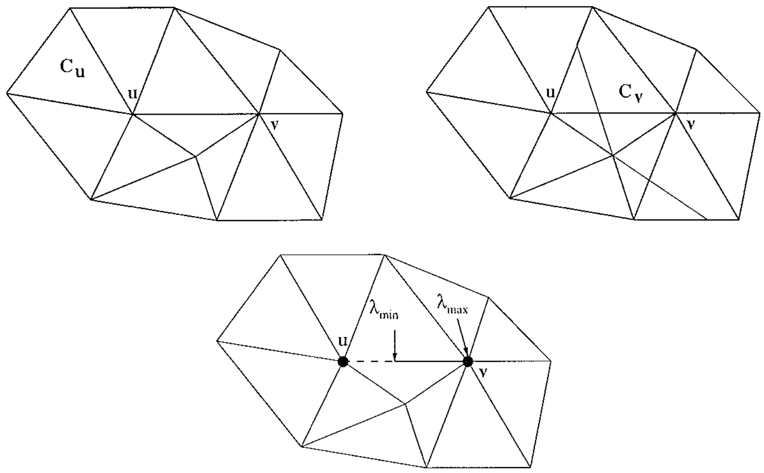

Figure 13. Obtaining $\lambda_{\min}$ and $\lambda_{\max}$.

All these algorithms require the vertices of C_u to be computed explicitly. In our case, as we consider only displacements along segments (segment representing the edge to be contracted) of the drawing, we only need to determine the maximal valid displacement along a segment.

Lemma 7. Let u be a node and let us choose one of its neighbors, for instance v_{k_0}. Then there exists a value $\lambda_{\max} > 0$ such that for $x \in [u, v_{k_0}]$

- if $||\overrightarrow{ux}|| < \lambda_{\max}$, then the displacement from u to $u' = x$ is valid,
- if $||\overrightarrow{ux}|| \geqslant \lambda_{\max}$, then the displacement from u to $u' = x$ is not valid.

Moreover, if we let

$$\lambda_k = ||\overrightarrow{uv_{k_0}}|| \frac{\det(\overrightarrow{uv_{k-1}}, \overrightarrow{uv_k})}{\det(\overrightarrow{uv_{k_0}}, \overrightarrow{v_{k-1}v_k})}, \quad \text{for } \varepsilon \leqslant k \leqslant d-1, \text{ modulo } d,$$

we have

$$\lambda_{\max} = \min_{\varepsilon \leqslant k \leqslant d-1} \lambda_k^+, \quad \text{where } \lambda_k^+ = \begin{cases} \lambda_k, & \text{if } 0 \leqslant \lambda_k \leqslant ||\overrightarrow{uv_{k_0}}||, \\ ||\overrightarrow{uv_{k_0}}||, & \text{else.} \end{cases}$$

Proof. Basically, using the previous lemma, this amounts to computing the intersection between $[u, v_{k_0}]$ and the boundary of C_u: for $k \in \{\varepsilon, \dots, d-1\}$, we compute $[u, v_{k_0}] \cap \mathcal{D}_{v_{k-1}, v_k}$.

For instance, suppose that we define an orthonormal basis $(\vec{e}_1, \vec{e}_2)$ such that $\overrightarrow{uv_{k_0}} = ||\overrightarrow{uv_{k_0}}||\vec{e}_1$. Let (α, β) be the coordinates of a point in $(u, \vec{e}_1, \vec{e}_2)$.

If $\beta_{k-1} = \beta_k$, then $\mathcal{D}_{u,v_{k_0}}$ and $\mathcal{D}_{v_{k-1},v_k}$ are strictly parallel ($\mathcal{D}_{u,v_{k_0}} = \mathcal{D}_{v_{k-1},v_k}$ is impossible as it would imply that the triangle (u, v_{k-1}, v_k) is flat).

If $\beta_{k-1} \neq \beta_k$, then an equation of $\mathcal{D}_{v_{k-1},v_k}$ is

$$\alpha = \mu_k \beta + \lambda_k, \quad \text{with } \mu_k = \frac{\alpha_k - \alpha_{k-1}}{\beta_k - \beta_{k-1}} \text{ and } \lambda_k = \frac{\beta_k \alpha_{k-1} - \beta_{k-1}\alpha_k}{\beta_k - \beta_{k-1}}.$$

Then $\mathcal{D}_{u,v_{k_0}} \cap \mathcal{D}_{v_{k-1},v_k}$ occurs at $\alpha = \lambda_k$. If $\lambda_k \notin [0, ||\overrightarrow{uv_{k_0}}||]$, $[u, v_{k_0}] \cap \mathcal{D}_{v_{k-1},v_k}$ is empty. Else, if we let $\{x_k\} = [u, v_{k_0}] \cap \mathcal{D}_{v_{k-1},v_k}$, $[u, x_k[$ is in C_u and $[x_k, v_{k_0}]$ is outside the region. Finally, defining λ_k^+ and $\lambda_{\max}$ as above leads to the conclusion.

It should be noted that $0 < \lambda_{\max}$ stems from two properties of C_u: $u \in C_u$ and C_u is topologically an open region. This proves the lemma. □

Suppose now that the t-contractable edge (u, v) is contracted, and that we want to derive an embedding of the contracted graph G' by simply modifying the location of the node u (of G') along the segment $[u, v]$: it is now possible to derive a sufficient and necessary condition to guarantee that the drawing of G' is indeed an embedding.

Theorem 1. In G, let d' be the degree of v and $(w_i)_{i=0}^{d'-1}$ be the neighbors of v and

$$\lambda_i = ||\overrightarrow{uv}|| \left(1 + \frac{\det(\overrightarrow{vw_{i-1}}, \overrightarrow{vw_i})}{\det(\overrightarrow{uv}, \overrightarrow{w_{i-1}w_i})}\right), \quad \text{for } \varepsilon' \leqslant i \leqslant d' - 1 \ (\varepsilon' = 0 \text{ or } 1),$$

and define

$$\lambda_{\min} = \max_{\varepsilon' \leqslant i \leqslant d'-1} \lambda_i^-, \quad \text{where } \lambda_i^- = \begin{cases} \lambda_i, & \text{if } 0 \leqslant \lambda_i \leqslant ||\overrightarrow{uv}||, \\ 0, & \text{else.} \end{cases}$$

Then it is possible to get an embedding of the contracted graph by modifying only the location of u along the segment $[u, v]$ iff $\lambda_{\min} < \lambda_{\max}$.

Proof. One gets the result by using the previous lemma from both sides of the contracted edge (see figure 13). □

Unfortunately, the sufficient and necessary condition is not always satisfied. Figure 14 illustrates such a case. In that situation, one can choose another point and restart the process. However, one can observe that, in some cases, the displacement of one or several well chosen points may invert the inequality connecting $\lambda_{\min}$ and $\lambda_{\max}$. In particular, if one can perform moves that lead to a new $\lambda_{\min}$ equal to zero (or a new $\lambda_{\max}$ equal to $||\overrightarrow{uv}||$), then the drawing of G' with u unchanged (or u moved to v) is necessarily an embedding. From these two considerations, we can make use of the following simple iterative heuristic:

When $\lambda_{\min} \geqslant \lambda_{\max}$

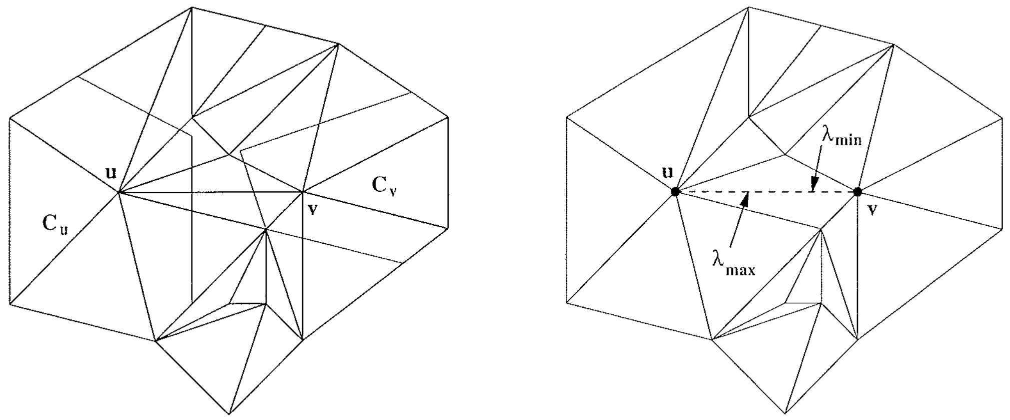

Figure 14. A case of $\lambda_{\min} > \lambda_{\max}$.

(I) $\lambda_{\max} = \min_k \lambda_k^+$.

(1) Let e_{k_0} be one edge such that $\lambda_{k_0}^+ = \lambda_{\max}$.

(2) Move the endpoint of e_{k_0} (the one closer to v) in such a way that $\lambda_{k_0}^+ = \|\overrightarrow{uv}\|$. If this is not possible, move it to improve $\lambda_{k_0}^+$.

(II) $\lambda_{\min} = \max_i \lambda_i^-$.

(1) Let e_{i_0} be one edge such that $\lambda_{i_0}^- = \lambda_{\min}$.

(2) Move the endpoint of e_{i_0} (the one closer to u) in such a way that $\lambda_{i_0}^- = 0$. If this is not possible, move it to improve $\lambda_{i_0}^-$.

For this heuristic process to be well defined, we have to explain the exact way of moving the endpoint of e_{k_0} (or e_{i_0}). Let the endpoints of e_{k_0} be x and y where y is the point to be moved, denote the angle $(\overrightarrow{xy}, \overrightarrow{xv})$ by α_0 and the half-line starting at x at angle α with $\overrightarrow{xy}$ by $\mathcal{D}_x^\alpha$.

In order to improve the value of the parameter $\lambda_{k_0}^+$, the following criteria should be satisfied, if possible: y' should belong to $\mathcal{D}_x^{\alpha_0}$; then $\lambda_{k_0}^+ = \|\overrightarrow{uv}\|$. As previously stated, the move is valid as long as y' is in C_y. If one calls $\alpha_{\min}$ (resp. $\alpha_{\max}$) the minimum (resp. maximum) value of α such that C_y and $\mathcal{D}_x^\alpha$ intersect, the condition to get an optimal move is simply $\alpha_{\min} < \alpha_0 < \alpha_{\max}$. In this case, y' is placed on $\mathcal{D}_x^{\alpha_p}$ with $\alpha_p = \alpha_0$. On the other hand, if $\alpha_0 \leqslant \alpha_{\min}$ or $\alpha_{\max} \leqslant \alpha_0$, then y can be moved nevertheless to improve $\lambda_{k_0}^+$, but then $\lambda_{k_0}^+$ remains lower than $\|\overrightarrow{uv}\|$. For instance, if $\alpha_0 < 0$ (resp. $\alpha_0 > 0$), y' can be placed onto $\mathcal{D}_x^{\alpha_p}$, with $\alpha_p = \alpha_{\min}/2$ (resp. $\alpha_{\max}/2$).

Finally as we know that $\mathcal{D}_x^{\alpha_p} \cap C_y$ is either a segment or a half-line (the latter can occur when C_y is unbounded), we choose to place y' either at the middle of the segment or at distance $\|\overrightarrow{xy}\|$ of the endpoint of the half-line.

If in addition y is a common neighbor of u and v, then the set of valid moves can be smaller than C_y. Indeed, if there exists a common neighbor w of v and y

(other than u), the value of λ_i^- corresponding to $\mathcal{D}_{y,w}$ can deteriorate. Then instead of considering C_y as the set of valid moves, the set of valid moves is restricted to $C_y \cap R_y$, where R_y is the half-plane delimited by $\mathcal{D}_{y,w}$ and not containing v ($C_y \cap R_y$

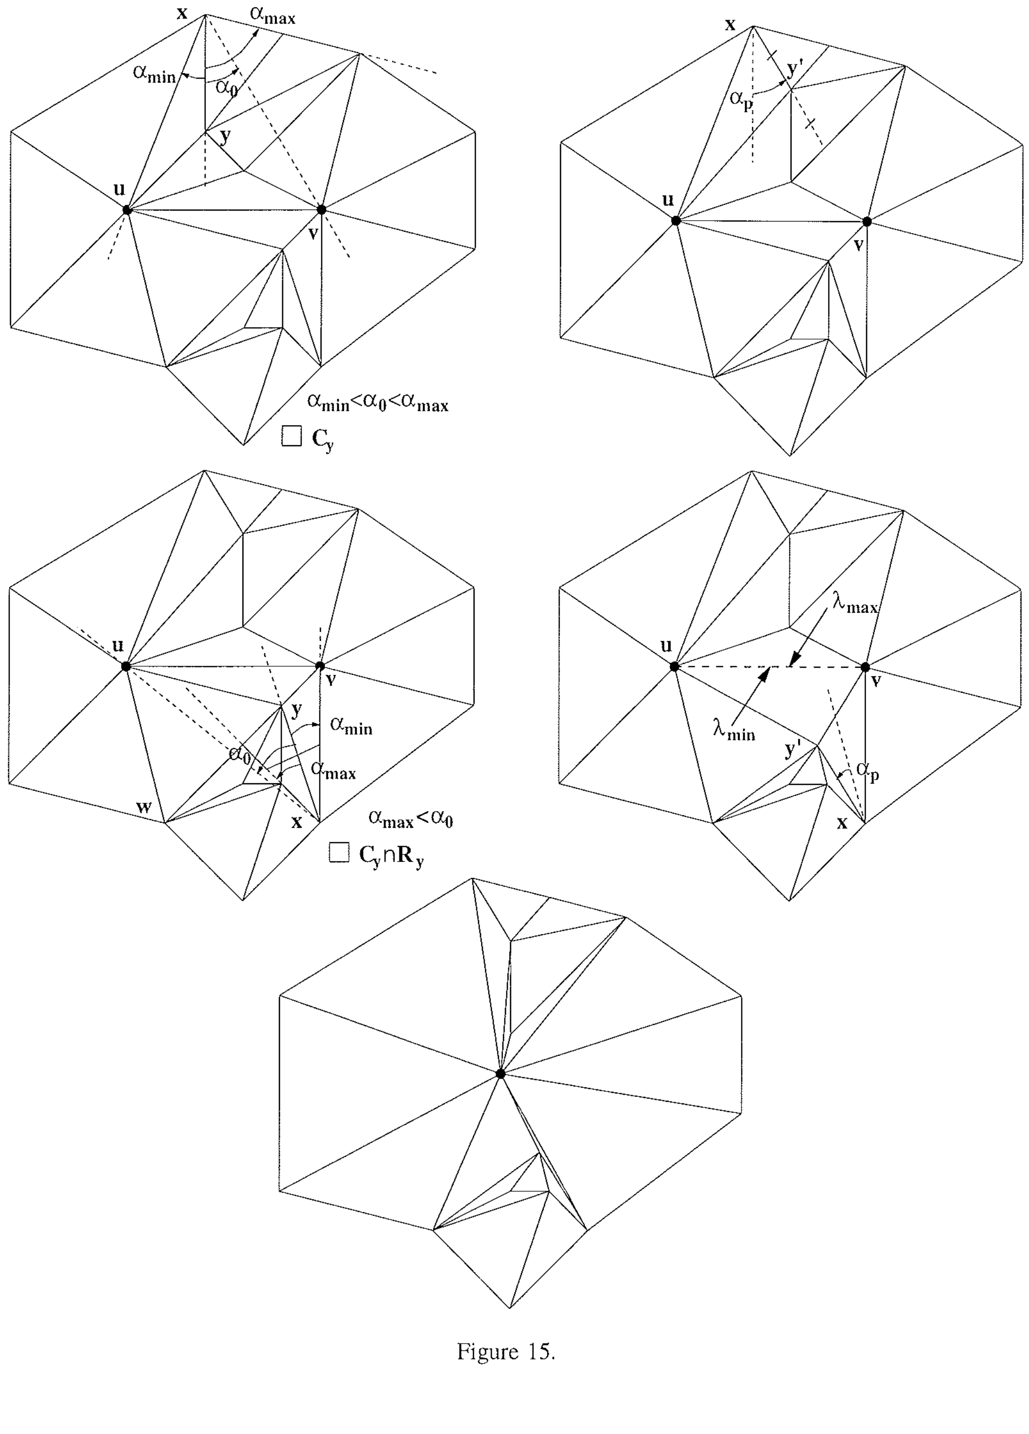

Figure 15.

is still a convex region of the plane). The new location of y is then determined as previously. In this way, the value of λ_i^- corresponding to $\mathcal{D}_{y,w}$ can not deteriorate any more, and the same holds for $\lambda_{\min}$. The different steps of the process are illustrated in figure 15.

4. Complexity and generalization

4.1. Complexity

In the previous section we have presented a method that draws an embedding of a contracted triangular mesh. This method leans on the drawing of the original triangular mesh and tries to deduce the embedding of the contracted mesh by moving only one of the endpoints of the chosen contracted edge along the segment of this edge. It has been shown that such moves are possible only when the control regions of u and v have a nonempty intersection on the segment $[u, v]$. Nevertheless, it is sometimes possible, if one accepts to move some neighboring points of u or v, to redraw the control region of u or v such that both regions intersect on $[u, v]$.

Let us now study the time complexity of the process of contraction and embedding of a given triangular mesh. We suppose that for each point u, the doubly-linked list of its neighbors, $\Gamma(u)$, is given counterclockwise (see [9]) together with its degree $d(u)$, its status (interior or contour) and its position (x_u, y_u).

The contraction and embedding process is made up of three stages. For a given node u the first stage consists in searching for a neighbor v that verifies the conditions stated in lemma 4. This can be done in $\mathrm{O}(d(u)\sum_{k=1}^{k=d(u)} d(v_k))$ by comparing $\Gamma(u)$ with $\Gamma(v_k)$ for each v_k. Following theorem 1, drawing an embedding of the contracted graph requires first an ordering of the neighbors of the two endpoints (ordering which is already at hand) and then the computation of the $d(u) + d(v)$ quantities (λ_k^+) and (λ_i^-). The complexity of this second stage is $\mathrm{O}(d(u) + d(v))$. If $\lambda_{\min} < \lambda_{\max}$ then the third step consists in updating the data corresponding to the contracted mesh. The update of the degrees (at most three), status (one) and position (one) is done in $\mathrm{O}(1)$ operations. Let v_1 (and possibly v_2) be the common neighbor(s) found during the first stage. If one assumes that the contracted node is called v in G', the neighbors list of a neighbor of v (other than v_1 and v_2) is not modified. In this way, updating all the neighbors lists requires first to eliminate the entry u in the neighbors list of the common neighbor(s) ($\mathrm{O}(d(v_1))$ or $\mathrm{O}(d(v_1)+d(v_2))$ operations), second to replace the entry u by v in the list of the other neighbors of u ($\mathrm{O}(\sum_{k\neq 1,2} d(v_k))$ operations), and third to build the neighbors list of the contracted node v ($\mathrm{O}(d(u) + d(v))$ operations). The complexity of stage three is then $\mathrm{O}(d(u) + \sum_{k=1}^{k=d(u)} d(v_k))$.

Therefore the general complexity is $\mathrm{O}(d(u)\sum_{k=1}^{k=d(u)} d(v_k))$. Let us notice that it is of interest to contract edges around nodes u such that $d(u)$ is bounded above by a constant. As in a planar graph there always exists a node with degree less than 6, this strategy makes sense.

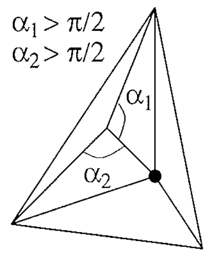

Figure 16.

On the other hand, if we need to use the heuristic process, we cannot say anything concerning the whole complexity because the number of inner iterations is not bounded a priori. Nonetheless, if we have recourse to the heuristic for attempting to draw the embedding of the contracted graph, it would be clever to couple it with heuristics such as those in [1, 3], which try to reduce the number of obtuse angles in a triangular mesh. Indeed, for an interior node u, when the polygon whose vertices are $\{v_0, \ldots, v_{d-1}\}$ is convex, it is in addition exactly the control region of u and consequently all displacements of u on the segment $[u, v]$ are valid. Moreover a sufficient condition for the associated polygon of u to be convex is that all the angles opposite to u of the triangles surrounding u are acute. Therefore any method, whose aim is to reduce the number of obtuse angles, contributes to the design of a good low-cost redrawing heuristic.

Evidently, as illustrated by figure 16, there exist cases where eliminating all (or even simply reducing the number of) obtuse angles is impossible.

4.2. Generalization

In this subsection, we derive some results for graph embeddings on the sphere. A graph G is said to be *spherical* if it has an embedding on the sphere. There is no exterior face for such graphs, but the concept of contours still applies. Defining holes as previously, the Euler formula for a spherical graph with h holes is

$$n - m + f = 2 - h. \qquad \text{(Euler ter)}$$

Let us notice that when drawing on a sphere, it can happen in particular that $h = 0$ or, equivalently, $n_c = m_c = 0$. A spherical graph without articulation nodes in which each face is a 3-cycle, except possibly for a fixed number of contours, is called a *spherical triangular* graph. Clearly, a spherical graph with n_c contour nodes and h holes is a triangular spherical graph if and only if:

$$m = 3n - n_c - 6 + 3h. \qquad \text{(Triangular ter)}$$

The notion of t-contractability is still valid and from the previous two formulas, it is a simple matter to check that the spherical version of lemmas 2, 4 and 5 holds.

Let us assume that the spherical graph is embedded on the unit sphere S centered at the origin o. In order to obtain an embedding of the spherical graph when contracting an edge (u_s, v_s), it is possible to use results of the previous section. Indeed, in the algorithms described for the plane embedding, when contracting the edge $(u_p,\ v_p)$, points are moved inside the convex envelope $\mathcal{E}_{u_p v_p}$ of the neighbors of u_p and v_p only. Thus, for $[u_s, v_s]$, the aim is to find a one-to-one mapping $\mathcal{M}$ from a part of the sphere, which includes at least $\mathcal{E}_{u_s v_s}$, to the plane. Then, in the plane, the segment $[\mathcal{M}(u_s), \mathcal{M}(v_s)]$ is contracted and the spherical embedding is simply deduced by using the inverse mapping (from the plane to the sphere).

For that, let z be a point of the space, z_0 a point of the sphere, $\mathcal{P}$ the plane tangent to $\mathcal{S}$ at z_0 (of equation $\overrightarrow{oz_0} \cdot \overrightarrow{oz} = 1$), and $\mathcal{S}^+$ the half-sphere defined by $\overrightarrow{oz_0} \cdot \overrightarrow{oz} > 0$ and $||\overrightarrow{oz}|| = 1$. Then there exists a one-to-one mapping $\mathcal{M}$ from $\mathcal{S}^+$ onto $\mathcal{P}$ defined by

$$\overrightarrow{o\mathcal{M}(z)} = \frac{1}{\overrightarrow{oz_0} \cdot \overrightarrow{oz}} \overrightarrow{oz}, \quad \text{and equivalently,} \quad \overrightarrow{o\mathcal{M}^{-1}(z)} = \frac{1}{||\overrightarrow{oz}||} \overrightarrow{oz}.$$

Thus, subject to the condition that there exists a half-sphere which includes all the neighbors of u_s and v_s, the method which associates a one-to-one mapping and the algorithms for plane embedding can be used to derive a spherical embedding of the contracted graph. In particular, the half-sphere condition is a very relaxed one as long as there are many vertices in the graph: taking $z_0 = u_s$ is a reasonable and straightforward choice.

This method could be generalized to other surfaces as soon as there exist one-to-one mappings to the plane that satisfy the locality condition expressed above, i.e., that the mapping includes all the neighbors of u and v.

Acknowledgement

The authors would like to thank Professor C. Delorme for helpful electronic discussions and especially for providing the counterexample of section 4.1.

Appendix

Proof of proposition 1. Let $\mathcal{D} = \mathcal{D}_+ \cup \mathcal{D}_-$ where $\mathcal{D}_+$ is the half-line starting at u and containing v. By the definition of θ_w, we have $\theta_w = 0$ if and only if $w \in \mathcal{D}_+$, $\theta_w = \pi$ if and only if $u \in \mathcal{D}_-$ and otherwise $0 < |\theta_w| < \pi$. Thus $\sin\theta_w = 0$ if and only if $w \in \mathcal{D}$.

Now, if $w \neq u$, the following expression holds,

$$\sin\theta_w = \frac{\det(\overrightarrow{uv}, \overrightarrow{uw})}{||\overrightarrow{uv}|| \cdot ||\overrightarrow{uw}||}.$$

Let $(\vec{e}_1, \vec{e}_2)$ be an orthonormal basis such that $\overrightarrow{uv} = k\vec{e}_1$. If $\overrightarrow{uw} = \alpha\vec{e}_1 + \beta\vec{e}_2$, then the half-planes are characterized by $\beta > 0$ and $\beta < 0$. Also, one gets

$$\sin\theta_w = \frac{k\beta}{|k|\sqrt{\alpha^2 + \beta^2}}.$$

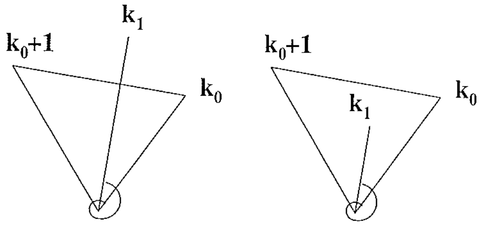

Figure 17. Local drawings which are not embeddings.

Thus the sign of $\sin\theta_w$ is constant over each half-plane delimited by $\mathcal{D}$, and the same holds for θ_w. □

Proof of proposition 2. First, we prove that θ_1 and θ_2 have the same sign, then the extension to all angles (θ_k) is straightforward by induction. Let us consider the labelling of the faces and neighbors. If u is an interior node, v_0 is chosen, and then T_1 is selected among two faces bordering (u, v_0). If u is a contour node, v_0 is chosen on the contour and T_1 is the only face bordering (u, v_0). In both cases, this defines in turn v_1 as $T_1 = \{u, v_0, v_1\}$. These nodes and faces are drawn in $\mathrm{Dr}(G)$. In order to get an embedding, T_2, the other face bordering (u, v_1) is on the other side of (u, v_1) with respect to T_1. Equivalently, this means that the point v_2 is not in the same half-plane as the point v_0 with respect to (u, v_1). By using proposition 1, this is equivalent to saying that $(\overrightarrow{uv_1}, \overrightarrow{uv_0})$ and $(\overrightarrow{uv_1}, \overrightarrow{uv_2})$ are not of the same sign, i.e., θ_1 and θ_2 have the same sign. Then the same holds for θ_2 and θ_3, etc.

Let us prove now that $|\sum_{k=\varepsilon}^{d-1}\theta_k| \leqslant 2\pi$. For instance, assume that for all k, $\theta_k > 0$. If $\sum_k \theta_k > 2\pi$, then there exist k_0 and k_1, $k_0 < k_1$, such that

(1) $\sum_{k\leqslant k_1}\theta_k > 2\pi$ and

(2) $\sum_{k\leqslant k_0}\theta_k \leqslant \sum_{k\leqslant k_1}\theta_k - 2\pi \leqslant \sum_{k\leqslant k_0+1}\theta_k$.

In other words, one of the two instances in figure 17 occurs. In both cases, the drawing is not an embedding. Thus $\sum_{k=\varepsilon}^{d-1}\theta_k \leqslant 2\pi$.

Now, if u is an interior node, $\sum_k \theta_k = (\overrightarrow{uv_0}, \overrightarrow{uv_0})$, i.e., $\sum_k \theta_k \equiv 0$, modulo 2π. Thus, as $0 < \sum_k \theta_k \leqslant 2\pi$, we have $\sum_k \theta_k = 2\pi$ in this case. If u is a contour node, $\sum_k \theta_k = (\overrightarrow{uv_0}, \overrightarrow{uv_{d-1}})$, i.e., $\sum_k \theta_k \not\equiv 0$, modulo 2π. Thus $\sum_k \theta_k < 2\pi$ in this case. This proves the proposition. □

Proof of proposition 3. C_u is not empty as $u \in C_u$ and convex by definition. From proposition 1, it is clear that $u' \in C_u$ iff $\theta'_k > 0$, $\forall k \in \{\varepsilon, d-1\}$.

For $u' \in C_u$, let $f(u') = \sum_{k=\varepsilon}^{d-1}\theta'_k$. As $u' \notin \{v_0, \ldots, v_{d-1}\}$, we have

$$\cos\theta'_k = \frac{(\overrightarrow{u'v_{k-1}} | \overrightarrow{u'v_k})}{\|\overrightarrow{u'v_{k-1}}\| \cdot \|\overrightarrow{u'v_k}\|}, \quad \forall k.$$

We also know that $0 < \theta'_k < \pi$, thus $\theta'_k = \arccos(\cos\theta'_k)$. This leads to the following expression for $f(u')$:

$$f(u') = \sum_{k=\varepsilon}^{d-1} \arccos\left\{ \frac{(\overrightarrow{u'v_{k-1}}|\overrightarrow{u'v_k})}{||\overrightarrow{u'v_{k-1}}|| \cdot ||\overrightarrow{u'v_k}||} \right\},$$

and in particular this shows that f is continuous in C_u.

If u is an interior point, then $f(u) = 2\pi$ and $f(u') = (\overrightarrow{u'v_0}, \overrightarrow{u'v_0})$, i.e., $f(u') \equiv 0$, modulo 2π. Now, as C_u is a connected region of the plane and f is continuous in C_u, this implies $f(u') = 2\pi$.

If u is a contour point, using the same arguments as before, we get $0 < f(u) < 2\pi$ and $f(u') \not\equiv 0$, modulo 2π. This implies in turn $0 < f(u') < 2\pi$. □

References

[1] K. Deljouie-Rakhshandeh, An approach to the generation of triangular grids possessing few obtuse triangles, Numer. Methods Engrg. 29 (1990) 1299–1321.

[2] L. R. Foulds, *Graph Theory Applications* (Springer, 1994).

[3] P.-L. George and J.-M. Nizard, Maillage aigu, Technical Report # 1414, INRIA, Le Chesnay, France (1991).

[4] J. L. Gross and T. W. Tucker, *Topological Graph Theory*, Wiley-Interscience Series in Discrete Mathematics and Optimization (Wiley, New York, 1987).

[5] F. Harary, *Graph Theory* (Addison-Wesley, 1969).

[6] G. R. Kampen, Orienting planar graphs, Discrete Math. 4 (1976) 337–341.

[7] M. S. Milgram, Does a point lie inside a polygon?, Comput. Phys. 84 (1989) 134–144.

[8] T. Nishizeki and N. Chiba, *Planar Graphs: Theory and Algorithms*, Annals of Discrete Mathematics 32 (North-Holland, 1988).

[9] R. C. Read, A new method for drawing a planar graph given the cyclic order of the edges at each vertex, Congr. Numer. 56 (1987) 31–44.

[10] W. Schnyder, Embedding planar graphs on the grid, in: *Proc. 1st Annual ACM–SIAM Symposium on Discrete Mathematics* (1990) pp. 138–148.

[11] C. Wooff, Checking whether a point lies inside a polygon, Comput. Phys. Comm. 36 (1985) 219–222.

[12] F. Yamaguchi, A unified approach to interference problems using a triangle processor, in: *Proc. ACM SIGGRAPH'85* (1985) pp. 141–149.