Tri-Edge-Connectivity Augmentation for Planar Straight Line Graphs

Marwan Al-Jubeh[1], Mashhood Ishaque[1], Kristóf Rédei[1], Diane L. Souvaine[1], and Csaba D. Tóth[1,2]

[1] Department of Computer Science, Tufts University, Medford, MA, USA ⋆
{maljub01,mishaque,kredei01,dls}@cs.tufts.edu
[2] Department of Mathematics, University of Calgary, AB, Canada ⋆⋆
cdtoth@ucalgary.ca

Abstract. It is shown that if a planar straight line graph (PSLG) with n vertices in general position in the plane can be augmented to a 3-edge-connected PSLG, then $2n-2$ new edges are enough for the augmentation. This bound is tight: there are PSLGs with $n \geq 4$ vertices such that any augmentation to a 3-edge-connected PSLG requires $2n-2$ new edges.

1 Introduction

Connectivity augmentation is a classical problem in graph theory with application in network design. Given a graph $G(V, E)$, with vertex set V and edge set E, and a constant $k \in \mathbb{N}$, find a minimum set E' of *new edges* such that $G'(V, E \cup E')$ is k-connected (respectively, k-edge-connected). For $k = 2$, Eswaran and Tarjan [2] and Plesník [10] showed independently that both edge- and vertex-connectivity augmentation problems can be solved in linear time. Watanabe and Nakamura [16] proved that the edge-connectivity augmentation problem can be solved in polynomial time for any $k \in \mathbb{N}$. The runtime was later improved (using the edge-splitting technique of Lovász [6] and Mader [7]) by Frank [3] and by Nagamochi and Ibaraki [8]. Jackson and Jordán [4] proved that the vertex-connectivity augmentation problem can be solved in polynomial time for any $k \in \mathbb{N}$. For related problems, refer to surveys by Nagamochi and Ibaraki [9], and by Kortsarz and Nutov [5].

The results on connectivity augmentation of *abstract* graphs do not apply if the input is given with a planar embedding, which has to be respected by the new edges (e.g., in case of physical communication or transportation networks). A *planar straight line graph* (PSLG) is a graph $G = (V, E)$, where V is a set of distinct points in the plane, and E is a set of straight line segments between the points in V such that two segments may intersect at their endpoints only. Given a PSLG $G = (V, E)$ and an integer $k \in \mathbb{N}$, the *embedding preserving k-connectivity* (resp., *k-edge-connectivity*) *augmentation* problem asks for a set of new edges E' of minimal cardinality such that $G = (V, E \cup E')$ is a k-connected (resp., k-edge-connected) PSLG. Since every planar graph has at

⋆ Partially supported by NSF grant CCF-0830734.
⋆⋆ Supported by NSERC grant RGPIN 35586. Research done at Tufts University.

least one vertex of degree at most 5, the embedding preserving k-connectivity (k-edge-connectivity) augmentation problems make sense for $1 \le k \le 5$ only.

Rutter and Wolff [11] showed that both the embedding preserving vertex- and edge-connectivity augmentation problems are NP-hard for $k = 2, \ldots, 5$. They reduce PLANAR3SAT to a decision problem whether a PSLG with m vertices of degree $k-1$ can be augmented to a k-edge-connected PSLG with $m/2$ new edges. The preservation of the input embedding imposes a severe restriction: for example, a path (as an abstract graph) can be augmented to a 2-connected graph by adding one new edge—however if a path is embedded as a zig-zag path with n vertices in convex position, then it takes $n-2$ (resp., $\lfloor n/2 \rfloor$) new edges to augment it to a 2-connected (2-edge-connected) PSLG [1]. If the vertices of a PSLG G are in convex position, then it *cannot* be augmented to a 3-connected (or 3-edge-connected) PSLG, since any triangulation (which is a *maximal* augmentation) on $n \ge 3$ vertices in convex position has a vertex of degree 2.

A few combinatorial bounds are known on the maximum number of new edges needed for the embedding preserving augmentation of PSLGs. It is easy to see that every PSLG with $n \ge 2$ vertices and p connected components can be augmented to a *connected* PSLG by adding at most $p-1 \le n-1$ new edges. Every connected PSLG G with $n \ge 3$ vertices and $b \ge 2$ distinct 2-connected blocks can be augmented to a *2-connected* PSLG by adding at most $b-1 \le n-2$ new edges [1]. Every connected PSLG with $n \ge 3$ vertices can be augmented to a *2-edge-connected* PSLG by adding at most $\lfloor (2n-2)/3 \rfloor$ new edges [13]. These bounds are tight in the worst case [1].

Valtr and Tóth [14] recently characterized the PSLGs that can be augmented to a 3-connected or a 3-edge-connected PSLG, respectively. In particular, a PSLG $G = (V, E)$ can be augmented to a *3-connected* PSLG if and only if V is not in convex position and E does not contain a chord of the convex hull $\mathrm{ch}(V)$. It can be augmented to a *3-edge-connected* PSLG if and only if V is not in convex position and E does not contain a chord of the convex hull $\mathrm{ch}(V)$ such that all vertices on one side of the chord are incident to $\mathrm{ch}(V)$. A PSLG with these properties is called *3-augmentable* and *3-edge-augmentable*, respectively. They also showed that every 2-edge-connected 3-edge-augmentable PSLG can be augmented to a 3-edge-connected PSLG with at most $n-2$ new edges, and this bound is the best possible [14].

Results. We show that every 3-edge-augmentable PSLG with n vertices can be augmented to 3-edge-connected PSLG with at most $2n-2$ new edges. This bound is the best possible. Our proof is algorithmic, and the new edges can be computed in $O(n \log^2 n)$ time in the real RAM model.

Lower bounds. We present two lower bound constructions. First, consider the empty graph with $n-1$ vertices in convex position, $n \ge 4$, and one vertex in the interior of their convex hull (Fig. 1, left). The only 3-connected or 3-edge-connected augmentation is a wheel graph with $2n-2$ edges. Second, consider a triangulation with m vertices, $2m-5$ bounded faces and the outer face. Put a singleton vertex in each bounded face, and 2 singletons behind each edge in the outer face as in Fig. 1, right. The only 3-edge connected augmentation is obtained by adding 3 new edges for each singleton in a bounded face, and 5 new edges for a pair of singletons behind each edge. A graph with $n = m + (2m-5) + 6 = 3m + 1$ vertices requires $3(2m-5) + 5 \cdot 3 = 2n - 2$ new edges.

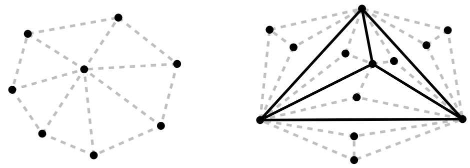

Fig. 1. The two lower bound constructions.

2 Preliminaries

In the next section (Section 3), we present an algorithm for augmenting a 3-edge-augmentable PSLG with n vertices to a 3-edge-connected PSLG with at most $2n - 2$ new edges. In this section, we introduce some notation for the number of bridges, components, reflex vertices, and singletons. They will play a key role in tracking the number of new edges. We also present a few simple inequalities used for verifying that at most $2n - 2$ edges have been added.

A vertex in a PSLG G is *reflex* if it is the apex of an angle greater than $180°$ in one of the incident faces of G. For example, a vertex of degree 1 or 2 is always reflex, but a singleton is not reflex. An edge in a graph G is a *bridge* if the deletion of the edge disconnects one of the connected components of G. Similarly a pair of edges in a graph G is a *2-bridge* (or 2-edge-cut) if the deletion of both edges disconnects one of the connected components of G.

By Euler's formula, a PSLG with n vertices has at most $2n - 5$ bounded faces. We use a stronger inequality that includes the number of reflex vertices.

Lemma 1 ([12]). *Let G be a* PSLG *with f bounded faces and n vertices, r of which are reflex. Then we have $f + r \leq 2n - 2$.*

Applying Lemma 1 for each 2-edge-connected component of a PSLG G, independently, we can conclude the following.

Corollary 1. *Let G be a* PSLG *with b bridges, c non-singleton components, f bounded faces, n vertices, r of which are reflex, and s singletons. Then*

$$b + c + f + r + 2s \leq 2n, \tag{1}$$

with equality if and only if G is a forest.

Proof. Let G_0 be the disjoint union of the 2-edge-connected components of G. Then G_0 has no bridges. Denote by c_0, f_0, and r_0, resp., the number of non-singleton components, bounded faces, and reflex vertices in G_0. Applying Lemma 1 for each non-singleton component, and charging 2 for each singleton, we have $f_0 + r_0 + 2s_0 \leq 2n - 2c_0$, or $c_0 + f_0 + r_0 + 2s_0 \leq 2n - c_0$. Whenever we add a new edge between two components, the number of components (including both singleton and non-singleton components) decreases by one, and the number of bridges increases by one. If one endpoint of a new edge is a singleton, then the singleton becomes a reflex vertex. Hence $b + c + f + r + 2s$ remains constant. Therefore, $b + c + f + r + 2s \leq 2n$, with equality if and only if $c_0 = 0$. □

After each step of our augmentation algorithm (in Section 3 below), we will derive an upper bound for the number of newly added edges in terms of b, c, f, r, and s. Inequality (1) will ensure that altogether at most $2n - 2$ edges are added.

k**-edge-connected components.** Given a graph $G = (V, E)$, two vertices $u, v \in V$ are k-edge-connected if they are connected by at least k edge-disjoint paths. This defines a binary relation on V, which is an equivalence relation [15]. The equivalence classes are the *k-edge-connected components* of G. By Menger's theorem, G is k-edge-connected if and only if there are k edge-disjoint paths between any two vertices, that is, if V is a single k-edge-connected component. For a graph G, let $\lambda(G)$ denote the number of 3-edge-connected components. For a PSLG G, where the boundary of the outer face is a simple polygon P, let $\lambda_h(G)$ denote the number of 3-edge-connected components that have at least one vertex incident on P.

Connecting singletons. To raise the edge-connectivity of G to 3, the degree of every singleton has to increase to at least 3. We can charge at most two new edges to each singleton (i.e., the term $2s$ in Inequality (1)). We will charge additional new edges at singletons to faces and reflex vertices (i.e., the terms $f + r + 2s$ in Inequality (1)). Every vertex of degree 2 is reflex, and we will charge one new edge per reflex vertex to the term r in Inequality (1), as well. The greatest challenge in designing our augmentation algorithm is to add r new edges at reflex vertices that serve two purposes: they (i) connect each reflex vertex to another vertex and (ii) connect a possible nearby singleton to the rest of the graph.

The following **technical lemma** is used in the analysis of Algorithm 1 below. A *wedge* (or angular domain) $\angle(\overrightarrow{a}, \overrightarrow{b})$, defined by rays $\overrightarrow{a}$ and $\overrightarrow{b}$ emanating from a point o, is the region swept by the ray rotated about o counterclockwise from position $\overrightarrow{a}$ to $\overrightarrow{b}$. For every reflex angular domain $\angle(\overrightarrow{a}, \overrightarrow{b})$, we define the *reverse wedge* to be $\angle(-\overrightarrow{b}, -\overrightarrow{a})$.

Lemma 2. *Let Q be a convex polygon, and let $\overrightarrow{r}_1$ and $\overrightarrow{r}_2$ be two rays emanating from Q. Then Q has a vertex u such that the reverse wedge of the exterior angle of Q at u is disjoint from both $\overrightarrow{r}_1$ and $\overrightarrow{r}_2$.*

3 Augmentation algorithm

Let $G = (V, E)$ be a 3-edge-augmentable PSLG with $n \geq 4$ vertices. We augment G to a 3-edge-connected PSLG with at most $2n - 2$ new edges in seven stages. At the end of stage $i = 1, 2, \ldots, 7$, the input G has been augmented to a PSLG G_i, where G_7 is 3-edge-connected. In stages 1-4 of the algorithm, we maintain a unique deformable edge $\tau(F)$ for each bounded face F of the current PSLG. In stage 5, the bounded faces will be partitioned into convex regions with the property (♡). In stage 6, each deformable edge $u_j v_j$ will be replaced by a path between u_j and v_j.

> (♡) The interior of $\mathrm{ch}(G_3)$ is decomposed into pairwise disjoint convex *regions*, $C_1, C_2, \ldots, C_\ell$. For every $j = 1, 2, \ldots, \ell$, there is a unique deformable edge e_j whose endpoints lie on the boundary of C_j; and the only edge of G_3 that may intersect the interior of C_j is e_j.

Notation for bridges, components, and vertices along the convex hull. We distinguish two types of bridges, edges, singletons, and non-singleton components. In the input graph G, let b_h (resp., g_h, r_h, and s_h) denote the number of bridges (resp., edges, reflex vertices, and singletons) along the convex hull $\mathrm{ch}(G)$. Clearly, we have $b_h \leq g_h$.

Let c_h be the number of non-singleton components with at least one vertex incident on the convex hull. Let $b_i = b - b_h$, $c_i = c - c_h$, $r_i = r - r_h$, and $s_i = s - s_h$.

Stage 1. Deformable edges for bounded faces. For every bounded face F in G, add a deformable edge $\tau(F)$ parallel to an arbitrary edge in E adjacent to F. We have created $\boxed{f}$ deformable edges, each of which is parallel to an existing edge of G.

Stage 2. Convex hull edges. Augment G_1 with the edges of the convex hull $\mathrm{ch}(G)$ (if they are not already present in G_1). The number of vertices along the convex hull is $r_h + s_h$, and so we have added at most $\boxed{r_h + s_h - g_h}$ new edges.

Lemma 3. *At the end of stage 2, we have $\lambda_h(G_2) \leq c_h + s_h + g_h$.*

Proof. Since all hull edges are already included in the PSLG G_2, the hull vertices are in one 2-edge-connected component of G_2. If any two hull vertices are connected by a path in the interior of $\mathrm{ch}(G)$, then they are also in the same 3-edge-connected component. Walk around the boundary of $\mathrm{ch}(G)$ in counterclockwise orientation, starting from an arbitrary vertex. We will assume that every vertex along the walk belongs to distinct 3-edge-connected components unless we can show that it is connected by a path in the interior of $\mathrm{ch}(G)$ to a previous vertex along the walk. The walk visits all $c_h + s_h$ (singleton or non-singleton) components of the input graph G incident on the convex hull. There are $c_h + s_h$ edges along which the walk arrives to a component of G for the first time, and the walk traverses g_h edges of G (staying in the same component of G). Let z denote the number of remaining hull edges: these are not edges of G and they lead to a component previously visited by the walk. The convex hull has a total of $c_h + s_h + g_h + z$ edges. Assume that the walk arrives to component $C \subseteq G$ for the kth time, $k \geq 2$, at some vertex v. Let u denote the last vertex of C visited by the walk. Then C contains a path from u to v whose edges lie in the interior of $\mathrm{ch}(G)$, hence v and u are 3-edge-connected in G_2. Hence, we have $\lambda_h(G_2) \leq c_h + s_h + g_h$. □

Stage 3. Connecting non-singleton components. For each non-singleton component of G_2 that lies in the interior of $\mathrm{ch}(G)$, we add two new edges to connect it to the component containing the convex hull.

Repeat the following procedure while there is a non-singleton component in the interior of $\mathrm{ch}(G)$. Refer to Fig. 2. Let G_h denote the connected component that contains the hull edges. Let H be the disjoint union of all non-singleton components in the interior of $\mathrm{ch}(G)$. Let U denote the set of vertices of $\mathrm{ch}(H)$. Each vertex $u \in U$ is a reflex vertex in G_2. Shoot a ray along the bisector of the reflex angle at each $u \in U$, until it hits an edge $v_u w_u$ outside of $\mathrm{ch}(H)$. Pick a vertex $u \in U$ for which the distance between u and the supporting line of $v_u w_u$ is minimal. Compute the shortest (i.e., geodesic) paths $\mathrm{path}(u, v_u)$ and $\mathrm{path}(u, w_u)$ with respect to the connected PSLG G_h. Any internal vertex of these paths is closer to the supporting line of $v_u w_u$ than u, and so it must be a vertex of G_h. Augment G_2 with the edges incident to u along each of $\mathrm{path}(u, v_u)$ and $\mathrm{path}(u, w_u)$. The new edges are not bridges, and they partition the reflex angle at u into three convex angles. We have

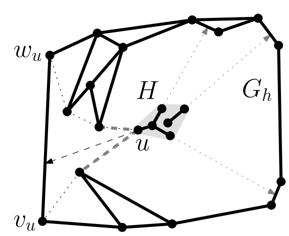

Fig. 2. Connecting a non-singleton component.

used $\boxed{2c_i}$ new edges. All c_i non-singleton components have been connected, and the number of reflex vertices in the interior of $\mathrm{ch}(G)$ has decreased by at least c_i. The resulting PSLG G_3 has one non-singleton component (which contains all hull edges) and all singletons lie in bounded faces.

Stage 4. Eliminate bridges in the interior of $\mathrm{ch}(G)$. It was shown by Abellanas *et al.* [1, Lemma 4] that if $G = (V, E)$ is a connected PSLG with a bridge $e \in E$, then it can be augmented with one new edge e' to a PSLG $(V, E \cup \{e'\})$ in which e and e' are not bridges. In other words, we can decrease the number of bridges by adding one new edge. We use at most $\boxed{b_i}$ new edges to eliminate bridges of G_3, which lie inside $\mathrm{ch}(G)$. We obtain a PSLG G_4, on which we execute Algorithm 1 during the stage 5.

Algorithm 1

Input: A PSLG $G_4 = (V, E_4)$ that consists of some singletons and a 2-edge-connected component such that the boundary of the outer face is a simple polygon P_4; and a function τ that maps a unique edge $\tau(F)$ to every bounded face F of G_4.
Output: $G_5 = (V, E_5)$.
For every bounded face F in G_4, set $C(F) := \mathrm{int}(F)$. Let R be the set of reflex vertices lying in the interior of P_4. Compute W_u and $\overrightarrow{a}_u$ with respect to G_4 for every vertex $u \in R$. Set $E_5 := E_4$.
while $R \neq \emptyset$ **do**
  if there is a vertex $u \in R$ that sees a non-singleton vertex v in W_u (Fig. 3(b)), **then**
    set $E_5 := E_5 \cup \{uv\}$ and $R := R \setminus \{u\}$. Edge uv splits face F_u into two faces, and it also splits region $C(F_u)$ into two regions. If $v \in R$ and uv splits the reflex angle at v into two convex angles, then set $R := R \setminus \{v\}$.
  else if there is a vertex $u \in R$ such that $\overrightarrow{a}_u$ does not hit edge $\tau(F_u)$ (Fig. 3(c)), **then**
    pick a vertex $u \in R$ such that $e_u \neq \tau(F_u)$ and the distance between u and the supporting line of e_u is maximal. Set $v_1v_2 = e_u$ such that v_1 is on the same side of $\overrightarrow{a}_u$ as $\tau(F_u)$. Compute the shortest path $\mathrm{path}(u, v_2)$ in F, and denote its vertices by $\mathrm{path}(u, v_2) = (u = w_0, w_1, w_2, \ldots, w_\ell = v_2)$. Set $E_5 := E_5 \cup \{w_j w_{j+1} : 0 \leq j \leq \ell - 1\}$ and $R := R \setminus \{w_j : 0 \leq j \leq \ell - 1\}$. The new edges split F into $\ell + 1$ faces. The rays $\overrightarrow{a}_{w_j}$, $j = 0, 1, \ldots, \ell - 1$ (in this order) split region $C(F_u)$ into $\ell + 1$ regions. Each new edge $w_j w_{j+1}$ lies between two bisectors $\overrightarrow{a}_{w_j}$ and $\overrightarrow{a}_{w_{j+1}}$, and hence in a unique new region, which contains a unique new face incident on $w_j w_{j+1}$.
  else
    for every $u \in R$, ray $\overrightarrow{a}_u$ hits edge $\tau(F_u)$ (Fig. 3(d)). Pick a vertex $u \in R$ such that the distance between u and the supporting line of the edge $\tau(F_u)$ is minimal. Let $vw = \tau(F_u)$. Set $E_5 := (E_5 \setminus \{\tau(F_u)\}) \cup \{uv, uw\}$ and $R := R \setminus \{u\}$. The removal of $\tau(F_u)$ merges face F_u with an adjacent face, and then the two new edges split it into three faces; ray $\overrightarrow{a}_u$ splits region $C(F_u)$ into two regions.
  end if
end while

Stage 5. Adding new edges at reflex vertices in the interior of $\mathrm{ch}(G)$. At the beginning of this stage, we have a PSLG G_4 that consists of a 2-edge-connected PSLG G_h (which contains all hull edges) and some singletons. There are at most $r_i - c_i$ reflex vertices in the interior of $\mathrm{ch}(G)$. We modify G_h (by adding some new edges and "deforming" some of the deformable edges) to a PSLG where every 3-edge-connected component is incident to the outer face, and increase the number of edges by at most

$r_i - c_i$. We also compute a decomposition of the interior of $\mathrm{ch}(G)$ into convex regions with property (♡).

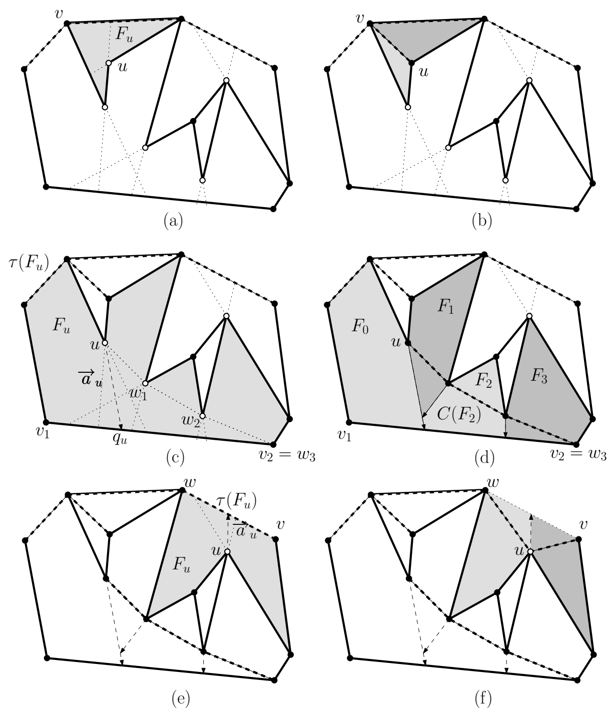

Fig. 3. (a) A PSLG G_4. Deformable edges are marked with dashed lines. Reflex vertices in R are marked with empty dots. The three rows show three consecutive *while* loops of Algorithm 1.

Let R be the set of reflex vertices in the interior of $\mathrm{ch}(G)$. For a reflex vertex u, let F_u denote the unique face adjacent to the reflex angle at u. Let W_u denote the reverse wedge of the reflex angle at u in G_4; let $\overrightarrow{a}_u$ be the bisector ray of the reflex angle at u, and let e_u be the first edge of the current PSLG hit by $\overrightarrow{a}$. Visibility is defined with respect to the current (augmented) PSLG, where all edges are *opaque*: a point p is *visible* to point q if the relative interior of segment pq is disjoint from edges of the PSLG.

Algorithm 1 will *process* the vertices in R, and increase the number of edges by one for each $u \in R$. In particular, it either adds a new edge uv, or replaces a deformable edge vw by two new deformable edges vu and uw. That is, at every reflex vertex, we add one or two new edge(s) that split(s) the reflex angle at u into smaller (but not necessarily convex) angles. In particular, some of the vertices in R remain reflex after they have been processed. Algorithm 1 will maintain a set of bounded faces $\mathcal{F}$ and a set of (not necessarily convex) regions $\mathcal{C}$. Each face $F \in \mathcal{F}$ corresponds to a region $C(F) \in \mathcal{C}$. All reflex vertices of a region $C(F) \in \mathcal{C}$ are reflex vertices of the corresponding face, and all *unprocessed* reflex vertices of a face $F \in \mathcal{F}$ are reflex vertices of $C(F)$. It follows that when all reflex vertices in the interior of $\mathrm{ch}(G)$ have been processed, $\mathcal{C}$ is a *convex* decomposition of the interior of $\mathrm{ch}(G)$. Initially, $\mathcal{F}$ is the set of all bounded faces in G_4, and $\mathcal{C} = \mathcal{F}$.

In the following lemma, we assume that the input of Algorithm 1, G_4, has no singletons. Note that singletons are not affected during Algorithm 1.

Lemma 4. *Let $G_4 = (V, E_4)$ be a 2-edge-connected* PSLG *such that the boundary of the outer face is a simple polygon P_4; and let τ map a unique edge $\tau(F)$ to every bounded face F of G_4. Algorithm 1 outputs a* PSLG *G_5 such that the boundary of the outer face is a simple polygon P_5, and every 3-edge-connected component is incident on P_5.*

Proof. Consider the output $G_5 = (V, E_5)$ of Algorithm 1. Note that Algorithm 1 does not necessarily *augment* G_4 to G_5, since it may replace a deformable edge $vw = \tau(F_u)$ by a path (vu, uw). In the course of several steps, an edge $vw = \tau(F_u)$ of G_4 may "evolve" into a simple path between v and w. In particular, the number of edges between two subsets of vertices of V cannot decrease. The boundary of the outer face is modified during the algorithm if a ray $\overrightarrow{a}_u$ of a reflex vertex in the interior of P_4 hits an edge $\tau(F_u) = vw$ along P_4, and the edge vw is replaced by the edges uv and uw. Every such step maintains the property that the boundary of the outer face is a simple polygon. Hence the boundary of the outer face in G_5 is a simple polygon P_5. Moreover, all vertices of P_4 remain on the boundary of the outer face, and so $\mathrm{int}(P_5) \subseteq \mathrm{int}(P_4)$.

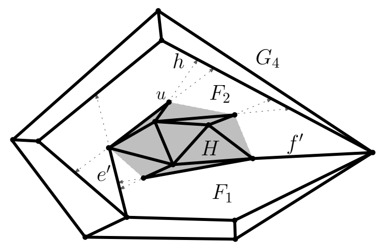

Fig. 4. A 2-bridge $\{e', f'\}$ in G_5.

Next we show that every 3-edge-connected component of G_5 is incident on the outer face. Assume that there is a 2-bridge $\{e', f'\}$ in G_5 such that both e' and f' are inside P_5. Denote the two connected components of $G_5 - \{e', f'\}$ by H and $G_5 - H$. We may assume w.l.o.g., that H lies in the interior of P_4 (and $G_5 - H$ contains all vertices of P_5). Hence H also lies in the interior of P_4. Refer to Fig. 4. Since G_4 is 2-edge-connected, there are exactly two edges, say, e and f, between vertices of H and $G_4 - H$. We may assume that either $e = e'$ or e has evolved to a path that contains e'; and similarly, either $f = f'$ or f has evolved to a path that contains f'. Since G_4 is 2-edge-connected, the minimum vertex degree in both G_4 and G_5 is at least 2. Every vertex of degree 2 is reflex. Algorithm 1 increased the degree of every vertex of degree 2 lying in the interior of P_4 to at least 3. It follows that H cannot be a singleton (which would be incident to e' and f' only), or a single edge (where each endpoint would be incident to this edge and either e' or f'). Therefore, H has at least three vertices. Hence, at least three vertices of H are incident on the convex hull $\mathrm{ch}(H)$. The vertices on the convex hull of H may be incident to exactly two bounded faces of G_4, say F_1 and F_2, which are adjacent to both e and f. Algorithm 1 modifies edges e or f only if they are the special edges $\tau(F_1)$ or $\tau(F_2)$, and the bisectors of all reflex angles in those faces hit these edges.

By Lemma 2, there is a vertex u on the convex hull $Q = \mathrm{ch}(H)$ such that the reverse wedge of the exterior angle at u does not intersect e or f. Vertex u is reflex in G_4, it lies in the interior of P_4, and so it is initially in R. Since the edges of G_4 incident on u lie inside $\mathrm{ch}(H)$, the reverse wedge W_u is part of the reverse wedge of $\mathrm{ch}(H)$ at u. Hence $\overrightarrow{a}_u$ always hits some edge h in $G_4 - H$, with $h \neq e, f$. It follows that Algorithm 1 does not modify e or f as long as $u \in R$. In the while loop where Algorithm 1 removes u from R, there are three possible cases: (1) some vertex v in W_u is visible from u and we add a new edge uv; (2) ray $\overrightarrow{a}_u$ hits an edge $vw \neq \tau(F_u)$ outside of H, and we add all edges along the geodesic $\mathrm{path}(u,v)$; (3) ray $\overrightarrow{a}_u$ hits an edge $vw = \tau(F_u)$ outside of H, and we add the edges uv and uw. In all three cases, we add new edges between H and $G_4 - H$. This contradicts the assumption that there are only two edges between H and $G_5 - H$. We conclude that $\{e', f'\}$ is not a 2-bridge in G_5. □

Lemma 5. *At the end of stage 5, we have* $\lambda(G_5) \leq c_h + g_h + s$.

Proof. At the end of stage 2, we have $\lambda_h(G_2) \leq c_h + s_h + g_h$ by Lemma 3. Stages 3-4 did not change the convex hull edges, so at the end of stage 4 we have $\lambda_h(G_4) \leq c_h + s_h + g_h$. Since every 3-edge-connected component of G_5 is either one of the s_i singletons or incident on the outer face P_5, we have $\lambda(G_5) = \lambda_h(G_5) + s_i$. It is enough to show that $\lambda_h(G_5) \leq \lambda_h(G_4)$. Assume that P_4 is the boundary of the outer face in G_4. Algorithm 1 may change P_4 by replacing an edge vw of P_4 by the edges vu and uw for some reflex vertex $u \in \mathrm{int}(P_4)$. In each such step, a new vertex appears along the outer face, however this vertex is connected to another vertex of the outer face by a path that lies in the interior of P_5. None of these steps increases the number of components incident on the outer face, and so we have $\lambda_h(G_5) \leq \lambda_h(G_4)$. □

Stage 6. Connecting singletons. There are s_i singletons in the interior of $\mathrm{ch}(G)$, which lie in convex regions $C_j \in \mathcal{C}$ with property (♡). In each convex region C_j, $j = 1, 2, \ldots, \ell$, we replace the deformable edge $e_j = u_j v_j$ by a path between u_j and v_j that lies entirely in C_j and passes through all singletons in C_j. Let m be the number of singletons in the interior of C_j. Label them as $w_1, w_2, \ldots w_m$ as follows. First label the singletons on the left of $\overrightarrow{u_j v_j}$ in the decreasing order of angles $\angle(v_j, u_j, w_i)$; then label the singletons on the right of $\overrightarrow{u_j v_j}$ in increasing order of angles $\angle(u_j v_j, w_i)$. See two examples in Fig. **??**. Replace edge $u_j v_j$ by the simple path $(u_j, w_1, w_2, \ldots, w_m, v_j)$. We obtain a 2-edge-connected PSLG G_6. The number of edges has increased by $\boxed{s_i}$. Each of the s_i singletons of G_5 becomes a 3-edge-connected component in G_6. Hence the number of 3-edge-connected components does not change, and we have $\lambda(G_6) = \lambda(G_5) \leq c_h + g_h + s$.

Stage 7. Eliminating 2-bridges. The input PSLG G was 3-edge-augmentable. In stages 1-6, we have not added any chords of $\mathrm{ch}(G)$, and so at the beginning of stage 6, we have a 2-edge-connected 3-edge-augmentable PSLG G_6. We to obtain a 3-edge-connected PSLG G_7 with at most $\lambda(G_6) - 1 = \boxed{c_h + g_h + s - 1}$ new edges by [14]. This completes our augmentation algorithm.

Theorem 1. *Every 3-edge-augmentable* PSLG G *with* $n \geq 4$ *vertices can be augmented to 3-edge-connected* PSLG *with at most* $2n - 2$ *new edges.*

Proof. In stages 1-7, we have added at most $b_i + c + f + r + 2s - 1$ new edges. If G is not a forest, this is at most $2n - 2$ by Corollary 1. If G is a forest and has an edge (bridge) along $\mathrm{ch}(G)$, then $b_h \geq 1$, and again $b_i + c + f + r + 2s - 1 \leq 2n - 2$.

Let G be a forest with no edges along the convex hull (i.e., $b_h = 0$). We would like to show that our augmentation algorithm used fewer than $b_i + c + f + r + 2s - 1$ new edges. We distinguish four cases.

• *Case 1. A non-singleton component of G has two vertices, u and v, along the convex hull.* By adding all hull edges in stage 2, we eliminate the bridges along the path between u and v, and so we add fewer than b_i new edges in stage 4.
• *Case 2. Every component of G has at most one vertex along the convex hull, and $c_i \geq 1$.* Then in stage 3 we add two new edges to a vertex of each non-singleton component in the interior of $\mathrm{ch}(G)$. The two edges form a circuit with G_h, and so it either decreases $\lambda_h(G_3)$ or eliminates a bridge in the interior of $\mathrm{ch}(G)$.
• *Case 3. Every component of G has at most one vertex along the convex hull, $c_i = 0$, but G_2 has a bridge.* Then in stage 4 we add a new edge for each bridge. The first such new edge creates a circuit, which either contains another bridge of G_3 (eliminating at least two bridges at once), or contains a hull edge, decreasing $\lambda_h(G_3)$ by at least one.
• *Case 4. G consists of n singletons.* Then one can augment G to a plane Hamiltonian circuit (e.g., the Euclidean TSP tour on n vertices), with n new edges. We can augment the edge-connectivity from 2 to 3 with at most $n-2$ new edges [14]. In this case, again, G can be augmented to a 3-edge-connected PSLG with at most $2n-2$ new edges. □

References

1. M. Abellanas, A. García, F. Hurtado, J. Tejel, J. Urrutia, Augmenting the connectivity of geometric graphs, *Comput. Geom. Theory Appl.* **40** (3) (2008), 220–230.
2. K. P. Eswaran and R. E. Tarjan, Augmentation problems, *SIAM J. Comput.* **5** (4) (1976), 653-665.
3. A. Frank, Augmenting graphs to meet edge-connectivity requirements, *SIAM J. Discrete Math.* **5** (1) (1992), 22–53.
4. B. Jackson and T. Jordán, Independence free graphs and vertex connectivity augmentation, *J. Comb. Theory Ser. B* **94** (2005), 31-77.
5. G. Kortsarz and Z. Nutov, Approximating minimum cost connectivity problems, Chap. 58 in *Handbook of Approximation Algorithms and Metaheuristics (T. F. Gonzalez, ed.)*, CRC Press, Boca Raton, 2007.
6. L. Lovász, *Combinatorial Problems and Exercises*, North-Holland, 1979.
7. W. Mader, A reduction method for edge-connectivity in graphs, *Ann. Discr. Math.* **3** (1978), 145–164.
8. H. Nagamochi and T. Ibaraki, Augmenting edge-connectivity over the entire range in $\tilde{O}(nm)$ time, *J. Algorithms* **30** (1999), 253-301.
9. H. Nagamochi and T. Ibaraki, Graph connectivity and its augmentation: applications of MA orderings, *Discrete Appl. Math.* **123** (2002), 447-472.
10. J. Plesník, Minimum block containing a given graph, *Arch. Math.* **27** (6) (1976), 668–672.
11. I. Rutter and A. Wolff, Augmenting the connectivity of planar and geometric graphs, in *Proc. Conf. Topological Geom. Graph Theory (Paris, 2008)*, pp. 55–58.
12. D. L. Souvaine and Cs. D. Tóth, A vertex-face assignment for plane graphs, *Comput. Geom. Theory Appl.* **42** (5) (2009), 388–394.
13. Cs. D. Tóth, Connectivity augmentation in planar straight line graphs, in: *Proc. Intl. Conf. on Topological and Geometric Graph Theory (Paris, 2008)*, pp. 51–54.
14. Cs. D. Tóth and P. Valtr, Augmenting the edge connectivity of planar straight line graphs to three, in *Proc. 13th Spanish Meeting on Comp. Geom. (Zaragoza, 2009)*.
15. W. T. Tutte, Connectivity in Graphs, University of Toronto Press, 1966.
16. T. Watanabe and A. Nakamura, Edge-connectivity augmentation problems, *Comp. Sys. Sci.* **35** (1) (1987), 96–144.