Constrained tri-connected planar straight line graphs*

Marwan Al-Jubeh† Gill Barequet‡† Mashhood Ishaque† Diane L. Souvaine†
Csaba D. Tóth§† Andrew Winslow†

Abstract

It is known that for any set V of $n \geq 4$ points in the plane, not in convex position, there is a 3-connected planar straight line graph $G = (V, E)$ with at most $2n - 2$ edges, and this bound is the best possible. We show that the upper bound $|E| \leq 2n$ continues to hold if $G = (V, E)$ is constrained to contain a given graph $G_0 = (V, E_0)$, which is either a 1-factor (*i.e.*, disjoint line segments) or a 2-factor (*i.e.*, a collection of simple polygons), but no edge in E_0 is a proper diagonal of the convex hull of V. Since there are 1- and 2-factors with n vertices for which any 3-connected augmentation has at least $2n - 2$ edges, our bound is nearly tight in these cases. We also examine possible conditions under which this bound may be improved, such as when G_0 is a collection of interior-disjoint convex polygons in a triangular container.

1 Introduction

A graph is *k-connected* if it remains connected upon deleting any $k - 1$ vertices along with all incident edges. Connectivity augmentation problems are an important area in optimization and network design. The *k-connectivity augmentation* problem asks for the minimum number of edges needed to augment an input graph $G_0 = (V, E_0)$ to a k-connected graph $G = (V, E)$, $E_0 \subseteq E$. In abstract graphs, the connectivity augmentation problem can be solved in $O(|V| + |E|)$ time for $k = 2$ [3, 6, 7, 12], and in polynomial time for any fixed k [9].

Researchers have considered the connectivity augmentation problems over planar graphs where both the input G_0 and the output G have to be planar (that is, they have no K_5 or $K_{3,3}$ minors). Kant and Bodlaender [10] proved that already the 2-connectivity augmentation over planar graphs is NP-hard, and they devised a 2-approximation algorithm that runs in $O(n \log n)$ time. We consider 3-connectivity augmentation over planar *geometric* graphs, where the given straight line embedding of the input graph has to be preserved.

A planar straight-line graph (for short, PSLG) is a graph with a straight-line embedding in the plane. That is, the vertices are distinct points in the plane and the edges are straight-line segments between the incident endpoints (that do not pass through any other vertices). The *k-connectivity augmentation for* PSLG*s* asks for the minimum number of edges needed to augment an input PSLG

*This material is based upon work supported by the National Science Foundation under Grant No. 0830734. Research by Tóth was also supported by NSERC grant RGPIN 35586. Preliminary results have been presented at the *26th European Workshop on Computational Geometry (2010, Dortmund)* and at the *20th Annual Fall Workshop on Computational Geometry (2010, Stony Brook, NY)*.

†Department of Computer Science, Tufts University, Medford, MA, USA. {maljub01, barequet, mishaque, dls, cdtoth, awinsl02}@cs.tufts.edu

‡Department of Computer Science, Technion, Haifa, Israel. barequet@cs.technion.ac.il

§Department of Mathematics and Statistics, University of Calgary, Calgary, AB, Canada. cdtoth@ucalgary.ca

$G_0 = (V, E_0)$ to a k-connected PSLG $G = (V, E)$, $E_0 \subseteq E$. Rutter and Wolff [13] showed that this problem is NP-hard for every integer k, $2 \le k \le 5$. Note that the problem is infeasible for $k \ge 6$, since every planar graph has a vertex of degree at most 5. There are two possible approaches to get around the NP-hardness of the augmentation problem: (i) approximation algorithms, as was done for planarity-preserving 2-connectivity augmentation; and (ii) proving extremal bounds for the minimum number of edges sufficient for the augmentation in terms of the number of vertices, which we do here.

It is easy to see that for every $n \ge 4$, there is a 3-connected planar graph with n vertices and $\lceil 3n/2 \rceil$ edges, where all but at most one vertex have degree 3. On a set V of $n \ge 4$ points in the plane, however, a 3-connected PSLG may require many more edges. García *et al.* [4] proved that if $3 \le h < n$ points lie on the convex hull of V, then it admits a 3-connected PSLG $G = (V, E)$ with at most $\max(\lceil 3n/2 \rceil, n + h - 1) \le 2n - 2$ edges, and this bound is the best possible. If the points in V are in convex position (that is, $h = n$), then V does not admit any 3-connected PSLG.

Tóth and Valtr [14] characterized the *3-augmentable* planar straight-line graphs, that is, graphs that can be augmented to 3-connected PSLGs. Specifically, a PSLG $G_0 = (V, E_0)$ is 3-augmentable if and only if E_0 does not contain any edge that is a proper diagonal of the convex hull of V. Every 3-augmentable PSLG on n vertices can be augmented to a 3-connected triangulation, which has up to $3n - 6$ edges, but in some cases significantly fewer edges are sufficient. As mentioned above, the 3-connectivity augmentation problem for PSLGs is NP-hard, and no approximation is known. It is also not known how many new edges are sufficient for augmenting *any* 3-augmentable PSLG with n vertices. Such a worst case bound is known only for edge-connectivity: Al-Jubeh *et al.* [2] proved recently that every 3-edge-augmentable PSLG with n vertices can be augmented to a 3-edge-connected PSLG by adding at most $2n - 2$ new edges, and this bound is the best possible.

Our results. In the 3-connectivity augmentation problem for PSLGs, we are given a PSLG $G_0 = (V, E_0)$, and asked to augment it to a 3-connected PSLG $G = (V, E)$, $E_0 \subseteq E$. Intuitively, the edges in E_0 are either "useful" for constructing a 3-connected graph or they are "obstacles" that prevent the addition of new edges which would cross them. In this note, we explore some classes of 3-augmentable PSLGs with $n \ge 4$ vertices that can be augmented to 3-connected PSLGs which have at most $2n$ edges. Recall that $2n - 2$ edges may be necessary even for a completely "unobstructed" input $G_0 = (V, \emptyset)$. We prove that if G_0 is 1-regular (that is, a crossing-free perfect matching) or 2-regular (a collection of pairwise noncrossing simple polygons), then it can be augmented to a 3-connected PSLG which has at most $2n - 2$ or $2n$ edges, respectively.

Theorem 1. *Every 1-regular 3-augmentable* PSLG $G_0 = (V, E_0)$ *with* $n \ge 4$ *vertices can be augmented to a 3-connected* PSLG $G = (V, E)$, $E_0 \subseteq E$, *with* $|E| \le 2n - 2$ *edges.*

Theorem 2. *Every 2-regular 3-augmentable* PSLG $G_0 = (V, E)$ *with* $n \ge 4$ *vertices can be augmented to a 3-connected* PSLG $G = (V, E)$, $E_0 \subseteq E$, *with* $|E| \le 2n$ *edges.*

Figures 1(a)-1(b) depict 1- and 2-regular PSLGs, respectively, where all but one of the vertices are on the boundary of the convex hull. Clearly, the only 3-connected augmentation is the wheel graph, which has $2n - 2$ edges. We conjecture that Theorems 1 and 2 can be generalized to PSLGs with maximum degree at most 2.

Conjecture 1.1. *Every 3-augmentable* PSLG $G_0 = (V, E)$ *with* $n \ge 4$ *vertices and maximum degree at most 2 can be augmented to a 3-connected* PSLG $G = (V, E)$, $E_0 \subseteq E$, *with* $|E| \le 2n - 2$ *edges.*

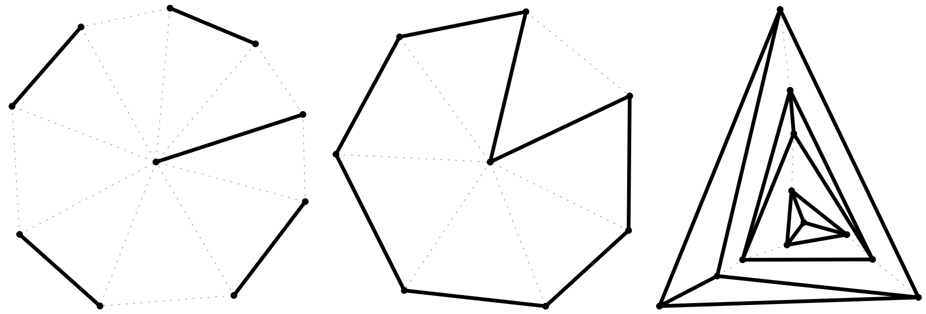

(a) (b) (c)

Figure 1: (a-b) 1- and 2-regular PSLGs whose only 3-connected augmentation is the wheel graph. (c). Nested copies of K_4, for which every 3-connected augmentation has at least $\frac{9}{4}n-3$ edges.

It is not possible to extend Theorems 1 and 2 to 3-regular PSLGs. For example, if $G_0=(V,E_0)$ is a collection of nested 4-cliques as in Fig. 1(c), then every 3-connected augmentation requires $3(\frac{n}{4}-1)$ new edges, which gives a total of $\frac{9}{4}n-3$ edges.

As mentioned above, every set of $n \geq 4$ points in the plane, $h \leq n$ of which lie on the boundary of the convex hull, admits a 3-connected PSLG with at most $\max(\lceil 3n/2 \rceil, n+h-1) \leq 2n-2$ edges [4]. We could not strengthen our Theorems 1 and 2 to be sensitive to the number of hull vertices. Some improvement may be possible for 1-regular PSLGs with fewer than $n-1$ vertices on the convex hull; the best lower bound construction we found with a triangular convex hull requires only $\frac{7}{4}(n-2)$ edges in total (Fig. 2(a)). For 2-regular PSLGs, however, one cannot expect significant improvement even if $h=3$. If $G_0=(V,E_0)$ consists of $\frac{n}{3}$ nested triangles (Fig. 2(b)), then any augmentation to a 3-connected PSLG has at least $2n-3$ edges.

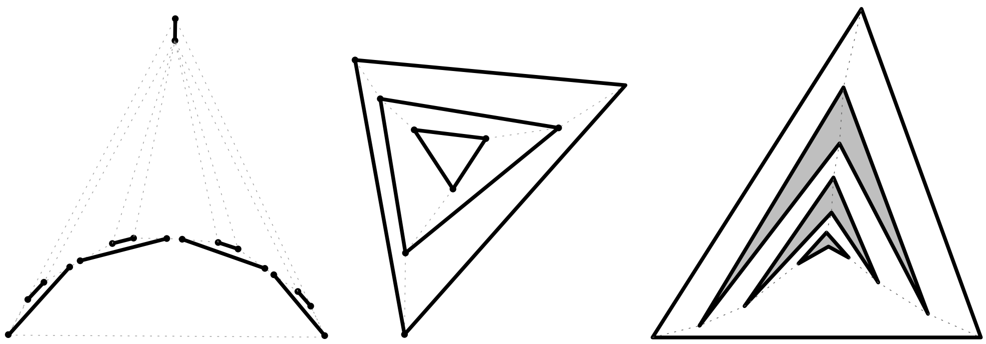

(a) (b) (c)

Figure 2: (a) A 1-regular PSLG on n vertices with a triangular convex hull whose 3-connected augmentations have at least $\frac{7}{4}(n-2)$ edges. (b) A 2-regular PSLG on n vertices with a triangular convex hull such that every 3-connected augmentation has at least $2n-3$ edges. (c) Interior-disjoint simple polygons in a triangular container, for which every 3-connected augmentation has $2n-3$ edges.

Obstacles in a container. We have considered whether Theorem 2 can be improved for collections of simple polygons, where the convex hull is a triangle, and there is no nesting among the remaining polygons. We model such 2-regular PSLGs as a collection of *interior-disjoint simple polygons in a triangular container.* Figure 2(c) shows a construction where every 3-connected augmentation still requires $2n-3$ edges. In this example the polygons are nonconvex, and they are "nested" in the sense that each polygon is visible from at most one larger polygon.

In Section 5, we derive lower bounds for the 3-connectivity augmentation of 2-regular PSLGs G_0, where G_0 is a collection of interior-disjoint *convex* polygons (called *obstacles*) lying in a triangular container. All our lower bounds in this section are below $2n-2$, which suggests that Theorem 2 may be improved in this special case.

Organization. In Section 2, we introduce a general framework for 3-connectivity augmentation, and prove that every nonconvex simple polygon with n vertices can be augmented to a 3-connected PSLG which has at most $2n-2$ edges. We prove Theorems 1 and 2 in Section 3 and 4, respectively. Lower bounds for the model of disjoint convex obstacles in a triangular container are presented in Section 5. We conclude with open problems in Section 6.

2 Preliminaries

In this section, we prove two preliminary results about *abstract* graphs, which are directly applicable to the 3-connectivity augmentation of simple polygons. In an (abstract) graph $G=(V,E)$, a subset $U \subseteq V$ is called *3-linked* if G contains at least three independent paths between any two vertices of U. (Two paths are *independent* if they do not share any edges or vertices apart from their endpoints.) By Menger's theorem, a graph $G=(V,E)$ is 3-connected if and only if V is 3-linked in G. The following lemma gives a criterion for incrementing a 3-linked set of vertices with one new vertex.

Lemma 2.1. *Let $G=(V,E)$ be a graph such that $U \subset V$ is 3-linked. If G contains three independent paths from $v \in V \setminus U$ to three distinct vertices in U, then $U \cup \{v\}$ is also 3-linked.*

Proof. Assume that G contains three independent paths from $v \in V \setminus U$ to distinct vertices $u_1, u_2, u_3 \in U$. We need to show that for every $u \in U$, there are three independent paths between v and u. By Menger's theorem, it is enough to show that if we delete any two vertices $w_1, w_2 \in V \setminus \{u,v\}$, the remaining graph $G \setminus \{w_1,w_2\}$ still contains a path between v and u. Since there are three independent paths from v to u_1, u_2, and u_3, the graph $G \setminus \{w_1,w_2\}$ contains a path from v to u_i for some $i \in \{1,2,3\}$. If $u_i = u$, then we are done. Otherwise, $G \setminus \{w_1,w_2\}$ contains a path from u_i to u, since U is 3-linked. The union of these two paths (from v to u_i and from u_i to u) contains a path from v to u. □

Lemma 2.2. *Let $G_A=(V,A)$ be a 2-connected graph, and let $G_C=(V,C)$ be a 3-connected graph with $A \subseteq C$. Let $U_A \subseteq V$ be the set of vertices that have degree 3 or higher in G_A, and assume that U_A is 3-linked in G_A. Then $G_A=(V,A)$ can be augmented to a 3-connected graph $G_B=(V,B)$ with $A \subseteq B \subseteq C$, by adding at most $|V \setminus U_A|$ new edges. Furthermore, if $U_A=\emptyset$, then $|V|-2$ new edges are sufficient for the augmentation.*

Proof. We describe an algorithm that augments $G_A=(V,A)$ to a 3-connected graph $G_B=(V,B)$, $A \subseteq B \subseteq C$. We maintain a graph $G_i=(V,E_i)$ with $A \subseteq E_i \subseteq C$. Initially, we start with $i=0$ and $E_0=A$. We augment G_i incrementally by adding new edges from C until G_i becomes 3-connected, and then output $G_B=G_i$. We also increment the set $U_i \subseteq V$ of vertices that have degree 3 or higher in G_i, and maintain the property that U_i is 3-linked in G_i. In each step, we will increment G_i with one new edge such that U_i increases with at least one new vertex. If $U_i=\emptyset$ or $|U_i|=2$, then U_i will increase with *two* new vertices in a single step. Our algorithm terminates with $U_i = V$, and the above properties guarantee that altogether at most $|V \setminus U_0|$ new edges are added, and if $U_0 = \emptyset$, then at most $|V| - 2$ new edges are added.

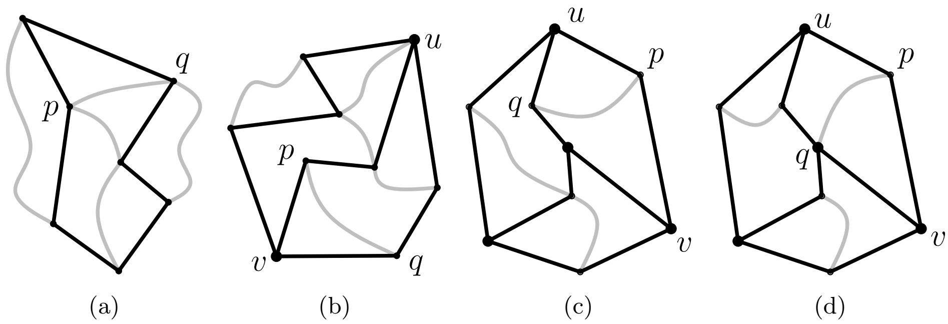

Figure 3: Illustration for the proof of Lemma 2.2. Edges of G_A are black, additional edges of G_C are gray, vertices in U_i are marked with large dots. (a) G_A is a Hamiltonian cycle. (b) G_A has two vertices of degree 3. (c) Both p and q lie in the interior of some paths between vertices of U_i. (d) Vertex q is in U_i.

It remains to describe one step of the augmentation, in which we increment E_i with one new edge from C. We distinguish between three cases.

Case 1: $U_i = \emptyset$. Since G_i is 2-connected, all vertices have degree 2 in G_i, and so it is a Hamiltonian cycle (Fig. 3(a)). Pick an arbitrary edge $pq \in C \setminus E_i$, and set $E_{i+1} = E_i \cup \{pq\}$. Let $U_{i+1} = \{p, q\}$ be the set of the two vertices of degree 3. Note that U_{i+1} is indeed 3-linked in G_{i+1}, as required.

Case 2: $|U_i| = 2$. Denote the vertices in U_i by u and v. Every edge in E_i is part of a path between u and v (Fig. 3(b)). Let $\mathcal{P}_i$ denote the set of all (at least three) paths of G_i between u and v. Note that every vertex in $V \setminus U_i$ lies in the interior of a path in $\mathcal{P}_i$. Since G_i is a simple graph, at least two paths in $\mathcal{P}_i$ have interior vertices. Let $P \in \mathcal{P}_i$ be a path with at least one interior vertex. Graph G_C contains some edge $pq \in C$ between an interior vertex p of P and a vertex q outside of P, otherwise the deletion of u and v would disconnect G_C. Set $E_{i+1} = E_i \cup \{pq\}$ and $U_{i+1} = U_i \cup \{p, q\}$. Note that G_{i+1} now contains three independent paths between any two vertices of $U_{i+1} = \{u, v, p, q\}$.

Case 3: $|U_i| \geq 3$. In this case, every edge in E_i is part of a path between two vertices in U_i. The vertices in U_i decompose G_i into a set $\mathcal{P}_i$ of paths, each of which connects two vertices in U_i. Note that every vertex in $V \setminus U_i$ lies in the interior of a path in $\mathcal{P}_i$. Pick two vertices $u, v \in U_i$ connected by a path in $\mathcal{P}_i$ with at least one interior vertex, and let V_{uv} be the set of interior vertices of all paths in $\mathcal{P}_i$ between u and v (Fig. 3(c)–3(d)). Graph G_C contains some edge $pq \in C$ between a vertex $p \in V_{uv}$ and a vertex q outside $V_{uv} \cup \{u, v\}$, otherwise the deletion of u and v would disconnect G_C. Set $E_{i+1} = E_i \cup \{pq\}$. Now G_{i+1} contains three independent paths from p to three vertices of U_i: independent paths to u and v along a path in $\mathcal{P}_i$, and a third path starting with edge pq and, if $q \notin U_i$, then continuing along a path containing q to a third vertex in $U_i \setminus \{u, v\}$. Similarly, if $q \notin U_i$, then G_{i+1} contains three independent paths from q to three vertices of $U_i \cup \{p\}$. By Lemma 2.1, $U_i \cup \{p, q\}$ is 3-linked in G_{i+1}. So we can set $U_{i+1} = U_i \cup \{p, q\}$. □

We show next that every nonconvex simple polygon in the plane with $n \geq 3$ vertices can be augmented to a 3-connected PSLG which has at most $2n - 2$ edges.

Corollary 2.3. *Every simple polygon with $n \geq 4$ vertices, not all in convex position, can be augmented to a 3-connected* PSLG *which has at most $2n-2$ edges.*

Proof. The edges and vertices of a simple polygon form a Hamiltonian cycle $G_A = (V, A)$. By the results of Valtr and Tóth [14], if the polygon is nonconvex, then it is 3-augmentable, so there is a 3-connected PSLG $G_C = (V, C)$, $A \subset C$. Lemma 2.2 completes the proof. □

3 Disjoint Line Segments

In this section, we prove Theorem 1. Let $G_A = (V, A)$ be a straight-line embedding of a perfect matching with $n \geq 4$ vertices, not all in convex position. We show that if no edge in A is a proper chord of the convex hull of the vertices, then G_A can be augmented to a 3-connected PSLG which has at most $2n-2$ edges. We use the result by Hoffmann and Tóth [5] that G_A can be augmented to a Hamiltonian PSLG G_H. If G_A is 3-augmentable, then G_H is also 3-augmentable and can be augmented to a 3-connected Hamiltonian PSLG G_C. In the following lemma, we use such a graph G_C, but we no longer rely on its *straight-line* embedding, our argument works for any plane drawing (where the edges are represented by Jordan arcs).

Lemma 3.1. *Let $G_A = (V, A)$ be a perfect matching with $n \geq 4$ vertices, and let $G_C = (V, C)$ be a 3-connected Hamiltonian planar graph with $A \subseteq C$. Then there is a 3-connected graph $G_B = (V, B)$ such that $A \subseteq B \subseteq C$ and $|B| \leq 2n-2$.*

Proof. Fix an arbitrary plane embedding of G_C. Let (V, H) be an arbitrary Hamiltonian cycle in G_C. If $A \subset H$, then the result follows from Lemma 2.2. Assume that $A \not\subset H$.

We construct a 3-connected graph $G_B = (V, B)$, $A \subseteq B \subseteq C$, incrementally. We maintain a 2-connected graph $G_i = (V, E_i)$ with $E_i \subseteq C$ (but E_i does not necessarily contain A). Let $U_i \subseteq V$ denote the set of *all* vertices that have degree at least 3 in G_i, which are called *hubs*. The hubs decompose G_i into a set $\mathcal{P}_i$ of paths between hubs. We maintain the following invariants for G_i.

- $\mathbf{I}_1$ $U_i \subseteq V$ is 3-linked in G_i;
- $\mathbf{I}_2$ between any two hubs in U_i, there are at most two paths in $\mathcal{P}_i$, at most one of which has interior vertices;
- $\mathbf{I}_3$ if there is an edge $uv \in A \setminus E_i$ between two hubs $u, v \in U_i$, then there is another path in $\mathcal{P}_i$ between u and v (such a path is called a lens);
- $\mathbf{I}_4$ $|E_i| \leq (n-2) + |U_i| - b_i$, where b_i is the number of bad paths in $\mathcal{P}_i$ (defined below).

In the next paragraphs, we define three families of so-called bad paths in $\mathcal{P}_i$ (lenses, diamonds, and monsters). Some of the definitions are formulated for subpaths of a path $P \in \mathcal{P}_i$ in order to keep track of subpaths that may become a bad path in $\mathcal{P}_{i+1}$ when some interior vertices of P become hubs. We begin with introducing some notation for paths. For two vertices, p and q, of a path $P \in \mathcal{P}_i$, let $P[p, q]$ denote the subpath of P between p and q (if p and q are the two endpoints of P, then $P = P[p, q]$). We say that the edges between interior vertices of $P[p, q]$ and vertices outside of $P[p, q]$ *go out of* $P[p, q]$.

Let P be a path in $\mathcal{P}_i$. Refer to Fig. 4. A subpath $P[u, v] \subseteq P$ is a *lens* if $uv \in A \setminus E_i$. A subpath $P[u_1, u_4] \subseteq P$ is a *diamond* if there are vertices u_1, u_2, u_3, u_4 along P in this order such that

(1) $u_1u_3, u_2u_4 \in A \setminus E_i$, and
(2) every edge going out of $P[u_1, u_4]$ is incident to u_2 or u_3.

The third family of subpaths, called monsters, is defined recursively. A subpath $P[v_1, v_4] \subseteq P$ is a *monster* if there are vertices v_1, v_2, v_3, v_4 along P in this order such that

(1) each of $P[v_1, v_2]$ and $P[v_3, v_4]$ is a lens, a diamond, or a smaller monster;
(2) $P[v_2, v_3]$ has at least one interior vertex;
(3) every edge going out of $P[v_1, v_4]$ is incident to v_2;
(4) every edge going out of $P[v_1, v_2]$ is incident to v_3, every edge going out of $P[v_2, v_3]$ is incident to v_4, and every edge going out of $P[v_3, v_4]$ is incident to v_1.

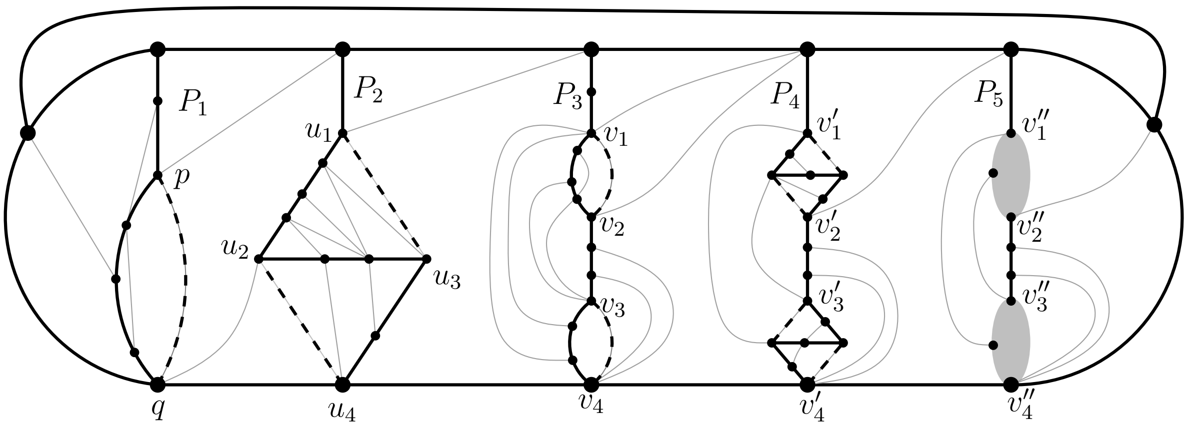

Figure 4: A lens $P_1[p, q]$. A diamond path $P_2[u_1, u_4]$. Monsters $P_3[v_1, v_4]$, $P_4[v'_1, v'_4]$, and $P_5[v''_1, v''_4]$. The vertices p, u_1, v_1, v'_1 and v''_1 are dangerous. Large dots are hubs in U_i, solid edges are in E_i, dashed edges are in $A \setminus E_i$, and gray edges are in $C \setminus E_i$, respectively. Gray ovals represent a lens, a diamond or a monster.

In a minimal monster, each of $P[v_1, v_2]$ and $P[v_3, v_4]$ is either a lens or a diamond. We say that a (sub)path $P' \subseteq P \in \mathcal{P}_i$ is *dangerous* if it is a lens, a diamond, or a monster. Note that every dangerous path has at least one interior vertex. A key property of a dangerous path P' is that each endpoint of P' is incident to some edge in $A \setminus E_i$ which goes to some other vertex of P'. For example, this implies that the middle portion $P[v_2, v_3]$ of a monster is *not* dangerous. A vertex p is called *dangerous* if p is an interior vertex of a path $P \in \mathcal{P}_i$ with endpoints u and v, and either $P[p, u]$ or $P[p, v]$ is a dangerous path. Note that at most one of $P[p, u]$ and $P[p, v]$ can be dangerous, since p is incident to at most one edge of $A \setminus E_i$.

We are now ready to define bad and and good paths in $\mathcal{P}_i$. A path $P \in \mathcal{P}_i$ is *bad* if it is dangerous, and *good* otherwise. We denote by b_i the number of bad paths in $\mathcal{P}_i$.

In each step of our algorithm, we modify E_i so that the set of hubs U_i strictly increases and invariants $\mathbf{I}_1$–$\mathbf{I}_4$ are maintained. The algorithm terminates when $U_i = V$. At that time, G_i is a 3-connected subgraph of G_C by invariant $\mathbf{I}_1$, every path in $\mathcal{P}_i$ is a single edge (hence there is no bad path in $\mathcal{P}_i$), all edges of A are contained in E_i by invariant $\mathbf{I}_3$, and $|E_i| \leq 2n - 2$ by invariant $\mathbf{I}_4$. So we can output $G_B = G_i$. We note here that the set of edges, E_i, does not always increase. Sometimes we may delete an edge from E_i (and add several edges from $C \setminus E_i$) to obtain E_{i+1}.

Initialization. Recall that (V, H) is a Hamiltonian cycle in G_C, with n edges, such that $A \not\subseteq H$. Let $pq \in A$ be an arbitrary chord of H. Vertices p and q decompose the Hamiltonian cycle into two paths, each of which has some interior vertices. Since G_C is 3-connected, it contains an edge $st \in C$ between two interior vertices of the two paths. Let $G_0 = (V, E_0)$ with $E_0 = H \cup \{pq, st\}$. Then the set of hubs is $U_0 = \{p, q, s, t\}$, which is 3-linked in G_0. The matching A contains pq and possibly st, so E_0 contains all edges of A induced by U_0. There are 6 paths in $\mathcal{P}_0$ between hubs. Two of these paths, pq and st, have no interior vertices, hence they are good. The other four paths are incident to p or q, where the incident edge of the matching is $pq \in A$, and so these paths are good, as well. We have $|E_0| = n + 2$, $|U_0| = 4$, and $b_0 = 0$, which gives $|E_0| = (n-2) + |U_0| - b_0$. The initial graph G_0 satisfies invariants $\mathbf{I}_1$–$\mathbf{I}_4$.

General Step i. We are given a graph $G_i = (V, E_i)$ satisfying invariants $\mathbf{I}_1$–$\mathbf{I}_4$ and $U_i \neq V$. We construct a graph G_{i+1} maintaining invariants $\mathbf{I}_1$–$\mathbf{I}_4$ so that the set of hubs strictly increases. Let X_i be the set of vertices $x \in V \setminus U_i$ such that x is an interior vertex of some path $P_x \in \mathcal{P}_i$, and it is adjacent to a vertex outside of P_x.

Remark 3.2. In every augmentation step, we augment G_i with an edge xy where $x \in X_i$ is an interior vertex of a path $P_x \in \mathcal{P}_i$ and y is outside of path P_x. We will distinguish several cases based on x and P_x, while vertex y is treated uniformly as outlined here. In all cases below, we describe the augmentation step assuming that y is already a hub in U_i, unless stated otherwise. Suppose now that y is an interior vertex of some path $P_y \in \mathcal{P}_i$. Then y becomes a hub in G_{i+1}. It has three independent paths to three distinct hubs in U_i, and so the extra hub at y does not violate $\mathbf{I}_1$. The hub at y decomposes P_y into two paths in $\mathcal{P}_{i+1}$, which satisfy $\mathbf{I}_2$. At most one of the two paths is bad (in this case an edge in $A \setminus E_i$ joins y to a vertex in P_y). Alternatively, if there is an edge $yz \in A \setminus E_i$ such that z is a hub outside of P_y, then we add edge yz to E_{i+1} in order to maintain invariant $\mathbf{I}_3$ for this edge. Inequality $|E_i| \leq (n-2) + |U_i| - b_i$ in Invariant $\mathbf{I}_4$ remains valid after adding one extra hub at y and either a possible bad path in $\mathcal{P}_{i+1}$ or a possible new edge yz in E_{i+1}. Furthermore, if P_y is a lens, then we add the edge between the endpoints of P_y to E_{i+1} to maintain invariant $\mathbf{I}_3$ for this edge. Inequality $|E_i| \leq (n-2) + |U_i| - b_i$ is again balanced by one additional edge and one fewer bad path. For verifying that invariants $\mathbf{I}_1$–$\mathbf{I}_4$ are maintained, we may assume that y is already a hub in U_i, unless stated otherwise.

Dispersing dangerous paths. We first show that if there is a bad path $P \in \mathcal{P}_i$, then we can always augment G_i so that the new hubs break P into good paths. The augmentation operation `Disperse` below is formulated in a more general setting, since it is a basic building block in several cases below. In certain cases, it is applied for a graph that satisfies invariants $\mathbf{I}_1$–$\mathbf{I}_3$ only.

> `Disperse` (G_i, P', xy).
>
> Input: $G_i = (V, E_i)$ is a graph satisfying invariants $\mathbf{I}_1$–$\mathbf{I}_3$, P' is a dangerous subpath of some path $P \in \mathcal{P}_i$ (possibly, $P' = P_x$), and $xy \in C$ is an edge between a vertex $x \in X_i$ in the interior of P' and a vertex outside of P, such that either $xy \in A \setminus E_i$ or none of the edges incident to x and going out of P' is in $A \setminus E_i$.
>
> 1. If P' is a lens with endpoints u and v, then $E_i := E_i \cup \{xy, uv\}$ (see Fig. 5(a)).
> 2. If P' is a diamond defined by vertices (u_1, u_2, u_3, u_4) such that $x = u_2$, then $E_i := E_{i+1} \cup \{xy, u_1u_3, u_2u_4\}$ (see Fig. 5(b)).
> 3. If P' is a monster defined by vertices (v_1, v_2, v_3, v_4), then augment G_i in three steps (Fig. 5(c)): set $E_i := E_i \cup \{xy\}$; call `Disperse`$(G_i, P[v_1, v_2], sv_3)$ for an edge sv_3 going out of $P[v_1, v_2]$; and call `Disperse`$(G_i, P[v_3, v_4], tv_1)$ for an edge tv_1 going out of $P[v_3, v_4]$.
>
> Let $E_{i+1} = E_i$ and return $G_{i+1} = (V, E_{i+1})$.

We show that operation `Disperse` maintains invariants $\mathbf{I}_1$–$\mathbf{I}_3$. Instead of maintaining Invariant $\mathbf{I}_4$, we show that the value of $|E_i| - |U_i| + b_i$ does not increase.

Proposition 3.3. *Operation* `Disperse`(G_i, P', xy) *augments* G_i *to* G_{i+1} *such that*

- *Invariants* $\mathbf{I}_1$–$\mathbf{I}_3$ *are maintained;*
- $|E_{i+1}| - |U_{i+1}| + b_{i+1} \leq |E_i| - |U_i| + b_i$*;*
- *the subgraph of* G_{i+1} *induced by the vertices of* P' *contain a simple cycle* σ_x *passing through vertex* x *and the two endpoints of* P'*.*

The second condition implies, in particular, that invariant $\mathbf{I}_4$ is maintained if G_i satisfies $\mathbf{I}_4$.

Proof. We proceed by induction on the length of P'. We distinguish three cases, depending upon whether P' is a lens, a diamond, or a monster. We may assume $y \in U_i$ by Remark 3.2.

Case (α): P' is a lens. In this case, $U_{i+1} = U_i \cup \{x, u, v\}$. The vertices of P' induce a cycle $\sigma_x = P' \cup \{uv\}$ in G_{i+1}. The vertices u, v, and x each have three independent paths to the two endpoints of P and to y in G_{i+1}. By invariants $\mathbf{I}_1$, $\mathbf{I}_2$, and Lemma 2.1, U_{i+1} is 3-linked in G_{i+1}. It is easy to verify that invariants $\mathbf{I}_2$–$\mathbf{I}_3$ are also maintained. All new paths in $\mathcal{P}_{i+1}$ are good. We have added 2 new edges. Vertex x is always a new hub. If x is the only new hub, then $P' = P_x$ and so $b_{i+1} = b_i - 1$. At any rate, we have $|E_{i+1}| - |U_{i+1}| + b_{i+1} \leq |E_i| - |U_i| + b_i$.

Case (β): P' is a diamond. In this case, $U_{i+1} = U_i \cup \{u_1, u_2, u_3, u_3\}$. The vertices u_1, u_2, u_3, u_4 are in the simple cycle $\sigma_x = P[u_1, u_2] \cup \{u_2u_4\} \cup P[u_4, u_3] \cup \{u_3u_1\}$ in G_{i+1}. Each new hub has three independent paths to the two endpoints of P and to y. By $\mathbf{I}_1$, $\mathbf{I}_2$, and Lemma 2.1, U_{i+1} is 3-linked in G_{i+1}. It is easy to verify that invariants $\mathbf{I}_2$–$\mathbf{I}_3$ are also maintained. All new paths in $\mathcal{P}_{i+1}$ are good. We have added 3 new edges. Vertices u_2, u_3 are always new hubs. If u_2 and u_3 are the only new hubs, then $P' = P$ and $b_{i+1} = b_i - 1$. Hence, we have $|E_{i+1}| - |U_{i+1}| + b_{i+1} \leq |E_i| - |U_i| + b_i$.

Case (γ): P' is a monster. It is easy to verify that the first step maintains $\mathbf{I}_1$–$\mathbf{I}_3$, and the last two steps maintain $\mathbf{I}_1$–$\mathbf{I}_3$ by induction. We construct a simple cycle σ_x in the subgraph of G_{i+1} induced by the vertices of P' that passes through v_1, v_2, and v_4. We construct σ_x explicitly as a union of two independent arcs between v_1 and v_3, one of which passes through $v_2 = x$, the other one through v_4. Refer to Fig. 5(c). In the second step, we added an edge sv_3 for some interior vertex s of path $P[v_1, v_2]$, and there is a simple cycle σ_s through v_1, v_2 and s by induction. Let one arc of σ_x be the union of v_3s and the part of σ_s from s to v_1 passing through v_2. In the third step, we added an edge tv_1 for some interior vertex t of path $P[v_3, v_4]$, and there is a simple cycle σ_t through v_3, v_4 and t by induction. Let the second arc of σ_x be the union of v_1t and the part of C_t from t to v_3 passing through v_4.

To verify $|E_{i+1}| - |U_{i+1}| + b_{i+1} \leq |E_i| - |U_i| + b_i$, we distinguish two subcases.

Case ($\gamma1$): $P' = P$. If $P' = P$, then the first two steps preserve the number of bad paths (P is bad initially, $P[v_1, v_2]$ is bad after the first step, and $P_x[v_3, v_4]$ is bad after the second step), however all new paths in $\mathcal{P}_{i+1}$ are good. Altogether, we have $b_{i+1} < b_i$. The quantity $|E_i| + b_i - |U_i|$ is unchanged in the first step, and it can only decrease in the last two steps by induction.

Case ($\gamma2$): $P' \neq P$. If $P' \neq P$, then at least one endpoint of P' is an interior vertex of P, and becomes a hub in G_{i+1}. Path $P \in \mathcal{P}_i$ may be either good or bad, and all new paths in $\mathcal{P}_{i+1}$ are good. Compared to case ($\gamma1$), the number of hubs increases by one more; while the first step may increase the number of bad paths, that is, $b_{i+1} \leq b_i + 1$. Altogether, invariant $\mathbf{I}_4$ is maintained. □

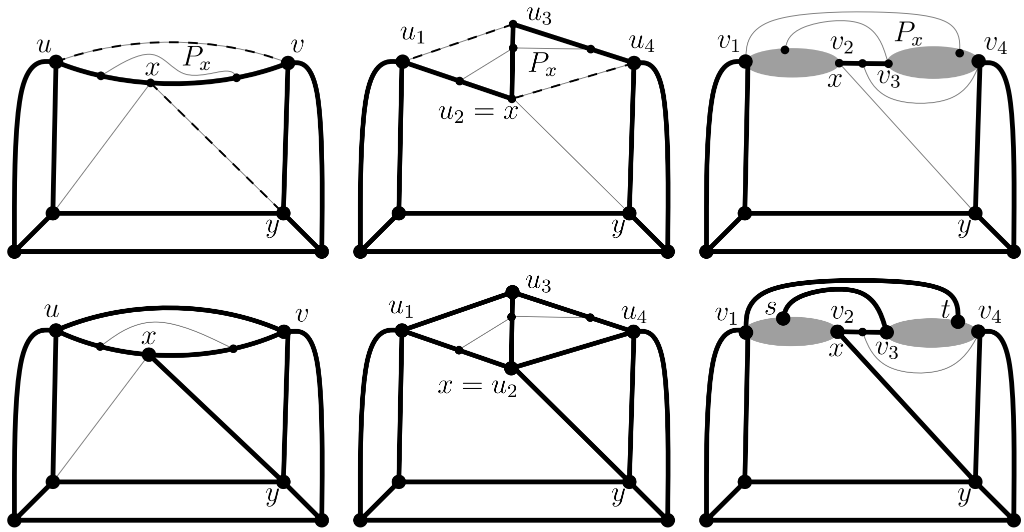

(a) (b) (c)

Figure 5: Operation $\texttt{Disperse}(G_i, P', xy)$ where $P' = P$ and P is a lens (a), a diamond (b), and a monster (c). The top row shows G_i, and the bottom row G_{i+1}. Vertex x is in the interior of a dangerous path $P' \subseteq P$ and y is a hub outside of path P. Solid edges are in E_i, dashed edges are in $A \setminus E_i$, and dotted edges are in $C \setminus (A \cup E_i)$.

Case analysis. We are given a graph $G_i = (V, E_i)$ satisfying invariants $\mathbf{I}_1$–$\mathbf{I}_4$ and $U_i \neq V$, and we would like to construct a graph G_{i+1} maintaining invariants $\mathbf{I}_1$–$\mathbf{I}_4$ such that $|U_i| < |U_{i+1}|$. We proceed with a case analysis, distinguishing between three main cases.

Case 1. There is a vertex $x \in X_i$ which is not dangerous. Refer to Fig. 6(a). If there is an edge in $A \setminus E_i$ between x and a vertex outside of P_x, then let this edge be xy, otherwise let xy be an arbitrary edge between x and a vertex outside of P_x. If $P_x \in \mathcal{P}_i$ is a bad path, then we can apply $\texttt{Disperse}(G_i, P_x, xy)$, and invariants $\mathbf{I}_1$–$\mathbf{I}_4$ are maintained by Proposition 3.3. Assume now that P_x is a good path. Let $E_{i+1} = E_i \cup \{xy\}$. For verifying that invariants $\mathbf{I}_1$–$\mathbf{I}_4$ are maintained, we may assume $y \in U_i$ by Remark 3.2. Then we have $U_{i+1} = U_i \cup \{x\}$, and U_{i+1} is 3-linked in G_{i+1} by $\mathbf{I}_1$, $\mathbf{I}_2$, and Lemma 2.1. The new hub x subdivides P_x into two good paths in $\mathcal{P}_{i+1}$ since x is not a dangerous vertex in P_x. It is easily checked that invariants $\mathbf{I}_1$–$\mathbf{I}_4$ are maintained.

Case 2. There is a dangerous vertex $x \in X_i$ such that x is in the interior of a dangerous path P'_x, $P'_x \subseteq P_x \in \mathcal{P}_i$. Let $\widehat{X}_i \subseteq X_i$ be the set of all vertices $x \in X_i$ that are dangerous and lie in the interior of some dangerous subpath $P'_x \subseteq P_x \in \mathcal{P}_i$.

For every $x \in \widehat{X}_i$, let y_x be an arbitrary adjacent vertex outside of P_x. Let $P'_x \subseteq P_x$ be a maximal dangerous subpath of P_x that contains x in its interior. If there is a vertex $x \in \widehat{X}_i$ such that $A \setminus E_i$ contains no edge incident to x going out of P'_x, then we apply operation $\texttt{Disperse}(G_i, P'_x, xy_x)$, and invariants $\mathbf{I}_1$–$\mathbf{I}_4$ are maintained by Proposition 3.3.

Assume now that for every $x \in \widehat{X}_i$, there is an edge $xz_x \in A \setminus E_i$ such that z_x is outside of P'_x. Note that z_x is a vertex of P_x, otherwise x would not be a dangerous vertex. Note that P'_x must be a lens since the edges going out of a diamond or a monster are not in $A \setminus E_i$. Denote the endpoints of P'_x by u_x and v_x, where $u_x v_x \in A \setminus E_i$. The vertices u_x, v_x, x, z_x are pairwise distinct (they are the endpoints of two edges in the matching A). Refer to Fig. 6(b).

We show that there is an $x \in \widehat{X}_i$ such that at least 3 vertices in $\{u_x, v_x, x, y_x\}$ are in the interior of P_x. Suppose, to the contrary, that for every $x \in \widehat{X}_i$, the two endpoints of P_x are in $\{u_x, v_x, x, y_x\}$. Without loss of generality, the endpoints of P_x are u_x and z_x. Then each interior vertex of P_x lies in the interior of lens $P_x[u_x, v_x]$ or lens $P_x[x, z_x]$. So every edge going out of P_x is incident to a vertex in $\widehat{X}_i$. Moreover, every $x \in \widehat{X}_i$ is joined to an endpoint of P_x. Therefore, P_x is a diamond for every $x \in \widehat{X}_i$, and the maximum dangerous subpath containing x in its interior is $P'_x = P_x$, contradicting our assumption that $P'_x = P_x[u_x, v_x]$.

Let $x \in \widehat{X}_i$ such that at least 3 vertices in $\{u_x, v_x, x, y_x\}$ are in the interior of P_x. Set $E_{i+1} = E_i \cup \{u_x v_x, x y_x, x z_x\}$. For the verifying invariants $\mathbf{I}_1$–$\mathbf{I}_4$ for G_{i+1}, we may assume $y \in U_i$ by Remark 3.2. Then $U_{i+1} = U_i \cup \{u_x, v_x, x, z_x\}$. We have added 3 new edges and at least 3 new hubs. The new hubs subdivide P_x into good paths, and so we have $b_{i+1} \leq b_i$. It is easily checked that invariants $\mathbf{I}_1$–$\mathbf{I}_4$ are maintained.

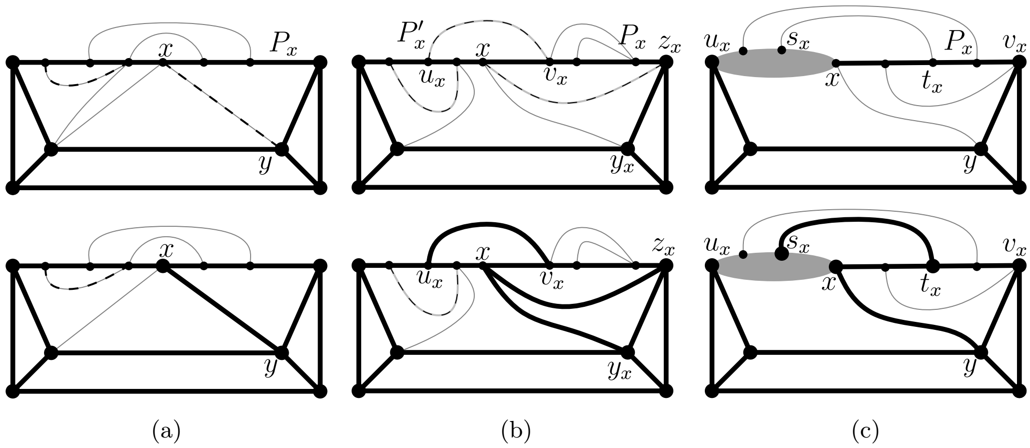

Figure 6: Step i of the algorithm. The top row shows G_i, and the bottom row G_{i+1}. (a) Case 1: vertex $x \in X_i$ is not dangerous. (b) Case 2: $x \in X_i$ is an interior vertex of a dangerous subpath $P'_x \subseteq P_x$, and it is incident to an edge $x z_x \in A \setminus E_i$. (c) Case 3a: $x \in X_i$ is a dangerous vertex of P_x because subpath $P_x[u_x, x] \subset P_x$ is dangerous, but t_x is not a dangerous vertex.

Case 3. Every $x \in X_i$ is a dangerous vertex of P_x but not an interior vertex of any dangerous subpath of P_x. We introduce some notation and then distinguish between four subcases. For every $x \in X_i$, denote the endpoints of P_x by $u_x, v_x \in U_i$. Assume, without loss of generality, that the subpath $P_x[u_x, x]$ of P_x is dangerous (hence $P_x[v_x, x]$ is not dangerous).

There is no vertex $x' \in X_i$ in the interior of $P_x[u_x, x]$, since $P_x[u_x, x]$ is a dangerous subpath of P_x. It follows that every path $P \in \mathcal{P}_i$ contains at most two vertices from X_i. Moreover, if $x, x' \in X_i$ are two distinct interior vertices of some path $P \in \mathcal{P}_i$, then the two endpoints of P are u_x and $u_{x'}$, and the subpaths $P[u_x, x]$ and $P[u_{x'}, x']$ are disjoint.

For every $x \in X_i$, the subpath $P_x[u_x, x]$ has at least one interior vertex (because $P_x[u_x, x]$ is dangerous). Since G_C is 3-connected, there must be at least one edge going out of $P_x[u_x, x]$. Let C_x denote the set of edges in C going out of $P_x[u_x, x]$. As noted above, all edges in C_x are incident to some vertex of P_x outside of $P_x[u_x, x]$. For every edge $s_x t_x \in C_x$, we use the convention that s_x denotes an interior vertex of $P_x[u_x, x]$ and t_x denotes a vertex of P_x outside of $P_x[u_x, x]$.

Note that if an edge $s_x t_x \in C_x$ is in $A \setminus E_i$, then t_x is not a dangerous vertex. Indeed, if t_x is dangerous, then either $P_x[u_x, t_x]$ or $P_x[v_x, t_x]$ is a dangerous subpath. However, $P_x[u_x, t_x]$ cannot be dangerous, since it contains x in its interior, and we assumed that x is not in the interior of any dangerous subpath of P_x. Hence $P_x[v_x, t_x]$ is dangerous, and so an edge in $A \setminus E_i$ joins t_x to some other vertex in $P_x[v_x, t_x]$.

Case 3a: There is a vertex $x \in X_i$ and an edge $s_x t_x \in C_x$ such that t_x is not a dangerous vertex. If C_x contains an edge in $A \setminus E_i$, then let $s_x t_x$ be such an edge, otherwise let $s_x t_x \in C_x$ be an arbitrary edge such that t_x is not dangerous. Refer to Fig. 6(c). Construct E_{i+1} from E_i in two steps as follows: set $E_i := E_i \cup \{xy\}$, and call $\texttt{Disperse}(G_i, P_x[u_x, x], s_x t_x)$. It is easily checked that the first step maintains $\mathbf{I}_1$–$\mathbf{I}_3$, and the second step maintains $\mathbf{I}_1$–$\mathbf{I}_3$ by Proposition 3.3. For invariant $\mathbf{I}_4$, we need to show that $|E_i| - |U_i| + b_i$ does not increase. In the first step, $|E_i| - |U_i| + b_i$ increases by one (each of $|E_i|$, $|U_i|$, and b_i is incremented by one). In the second step, operation $\texttt{Disperse}(G_i, P_x[u_x, x], s_x t_x)$ does not increase $|E_i| - |U_i| + b_i$ by Proposition 3.3. However, t_x emerges as a new hub (*c.f.*, Remark 3.2), but there is no edge $t_x z_x \in A \setminus E_i$ between t_x and any hub z_x. With an additional hub at t_x, the value of $|E_i| - |U_i| + b_i$ strictly decreases in the second step.

In the remaining cases (Cases 3b–3d), we assume that for every $x \in X_i$, the edges $s_x t_x \in C_x$ are not in $A \setminus E_i$, otherwise Case 3a would apply.

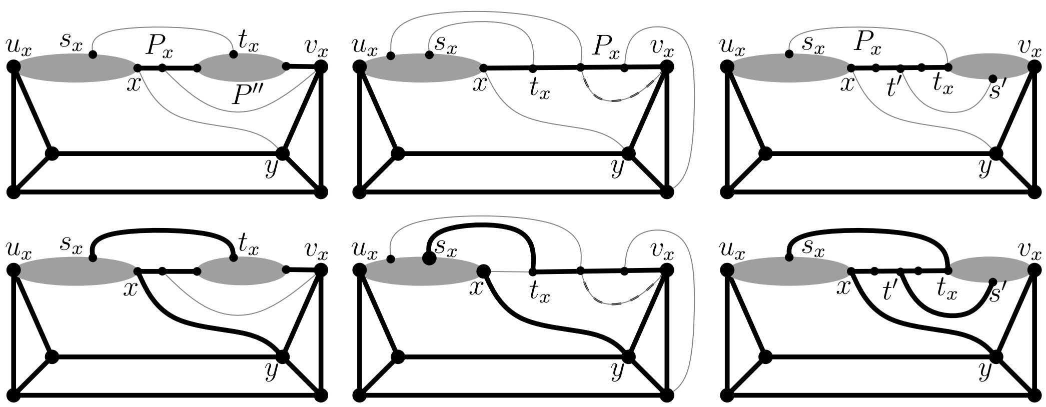

(a) (b) (c)

Figure 7: Step i of the algorithm. The top row shows G_i, and the bottom row G_{i+1}. (a) Case 3b: subpaths $P_x[u_x, x] \subset P_x$ and $P_x[v_x, t_x]$ are dangerous and t_x is in the interior of $P_x[v_x, x]$. (b) Case 3c: subpath $P_x[u_x, x] \subset P_x$ is dangerous but $P_x[x, t_x]$ has no interior vertices. (c) Case 3d: t_x is not adjacent to any vertex outside of P_x. There is an edge $s't'$ where s' is an interior vertex of $P_x[v_x, t_x]$ and t' is an interior vertex of $P_x[x, t_x]$.

Case 3b: There is a vertex $x \in X_i$ and an edge $s_x t_x \in C_x$ such that t_x is in the interior of a dangerous subpath P'' of $P_x[x, v_x]$. Refer to Fig. 7(a). Construct E_{i+1} from E_i in two steps as follows: set $E_i := E_i \cup \{xy\}$, and then apply independently $\texttt{Disperse}(G_i, P_x[u_x, x], s_x t_x)$, and $\texttt{Disperse}(G_i, P'', s_x t_x)$. The first step increases $|E_i| - |U_i| + b_i$. However, the two independent calls to $\texttt{Disperse}$ add the same edge $s_x t_x$, and so $|E_i| - |U_i| + b_i$ decreases by at least one. Altogether, invariants $\mathbf{I}_1$–$\mathbf{I}_4$ are maintained.

Case 3c: There is a vertex $x \in X_i$ and an edge $s_x t_x \in C_x$ such that t_x is a dangerous vertex but $P_x[x, t_x]$ has no interior vertices. Refer to Fig. 7(b). Construct E_{i+1} from E_i in three steps as follows: set $E_i := E_i \cup \{xy\}$, call $\texttt{Disperse}(G_i, P_x[u_x, x], s_x t_x)$, and then delete edge $x t_x$. The first two steps clearly maintain invariants $\mathbf{I}_1$–$\mathbf{I}_3$, but we need to be careful about the edge deletion. We show that after the deletion of $x t_x$, the original hubs in U_i remain 3-linked, vertex x becomes a hub with 3 independent paths to hubs in U_i; and t_x will be a hub in G_{i+1} if and only if it was already a hub in G_i. Indeed, the two endpoints of P_x are connected by $P_x[u_x, s_x] \cup \{s_x t_x\}$, so U_i remains 3-linked. In the second step, we create a simple cycle σ_{s_x} passing through x, s_x, and

u_x by Proposition 3.3. Hence the degree of x is at least 3 in G_{i+1}, and it has 3 independent paths to u_x, v_x, and y. Finally the degree of t_x does not change, and so it is a hub if and only if $t_x \in U_i$. So the third step also maintains invariants $\mathbf{I}_1$–$\mathbf{I}_3$,

For invariant $\mathbf{I}_4$, notice that the number of dangerous paths can only decrease by Proposition 3.3 (that is $b_{i+1} \le b_i$), and the effect of inserting edge xy and deleting edge xt_x cancel each other. Hence, invariant $\mathbf{I}_4$ is maintained.

Case 3d: For every $x \in X_i$ and $s_x t_x \in C_x$, vertex t_x is dangerous, it is not in the interior of any dangerous subpath of $P_x[x, v_x]$, and $P_x[x, t_x]$ has some interior vertices. For every $x \in X_i$, we choose an edge $s_x t_x \in C_x$ as follows. Let $s_x t_x \in C_x$ be an edge such that t_x is the closest vertex to x along P_x. Fix a vertex $x \in X_i$ such that the length of $P_x[x, t_x]$ is minimal.

Path P_x with vertices u_x, x, t_x, v_x satisfies conditions (1)-(2) in the definition of monsters. The vertices u_x, x, t_x, v_x appear in this order along P_x such that $P_x[u_x, x]$ and $P_x[t_x, v_x]$ are dangerous and $P_x[x, t_x]$ has at least one interior vertex. We show next that it satisfies condition (3), as well.

We show that every vertex going out of P_x is incident to x. Suppose to the contrary that there is a vertex $x' \in X_i$, $x' \ne x$, in P_x. Then $P_x[x', v_x]$ is a dangerous subpath, which is disjoint from $P_x[u_x, x]$. Vertex $x' \in X_i$ is not in the interior of any dangerous subpath of P_x. So x' is in $P_x[x, t_x]$ (possibly, $x' = t_x$). Note that x' is not in the interior of $P_x[x, t_x]$, otherwise t_x would be in the interior of the dangerous subpath $P_x[x', v_x]$ of $P_x[x, v_x]$. The only remaining possibility is $x' = t_x$. For vertex $x' \in X_i$, we have defined an edge $s_{x'} t_{x'}$ where $s_{x'}$ is an interior vertex of $P_x[x', v_x] = P_x[t_x, v_x]$ and $t_{x'}$ is a vertex of P_x outside of $P_x[v_x, x']$. However, $t_{x'}$ cannot be at u_x or in the interior of the dangerous path $P_x[u_x, x]$, otherwise Case 3a or 3b would apply for x'. Also, $t_{x'}$ cannot be in the interior of $P_x[x, x']$ because then subpath $P_{x'}[x', t_{x'}]$ would be strictly shorter than $P_x[x, t_x]$, contradicting the choice of $x \in X_i$ (Fig. 8, left). Therefore, we have $t_{x'} = x$ (Fig. 8, middle). Now consider the interior vertices of $P_x[x, x']$. These vertices are adjacent to the interior of neither $P_x[u_x, x]$ nor $P_x[x', v_x]$ by the choice of the edges $s_x t_x$ and $s_{x'} t_{x'}$. They are separated from all other vertices outside of $P_x[x, x']$ by the cycle $\{s_x x'\} \cup P_x[x', s_{x'}] \cup \{s_{x'} x\} \cup P_x[x s_x]$. There are no edges going out of $P_x[x, x']$, contradicting the 3-connectivity of G_C. We conclude that every vertex going out of P_x is incident to x.

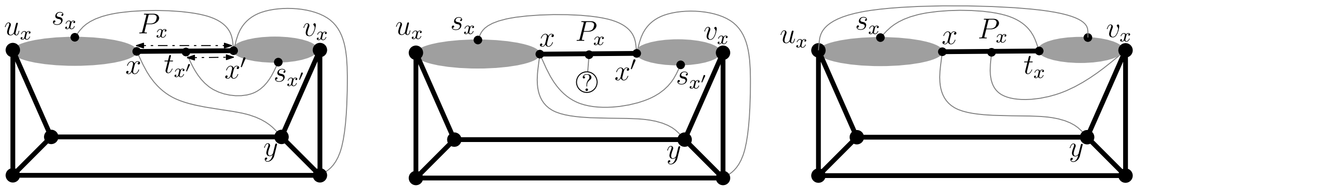

Figure 8: Case 3d: subpaths $P_x[u_x, x] \subset P_x$ and $P_x[v_x, t_x]$ are dangerous. Suppose that $t_x = x'$ is adjacent to a vertex outside of P_x. Left: If $t_{x'}$ is in the interior of $P_x[x, x']$, then $P_x[x', t_{t'}]$ is shorter than $P_x[x, t_x]$. Middle: If $t_{x'} = x$, then the are no edges out of $P_x[x, x']$. Right: If there is no edge between the interiors of $P_x[x, t_x]$ and $P_x[t_x, v_x]$, then all edges out of $P_x[t_x, v_x]$ go to v_x.

We also show that every edge going out of $P_x[u_x, x]$ is incident to t_x. Indeed, as noted above all edges in C_x are incident to a vertex in P_x outside of $P_x[u_x, x]$. They cannot be incident to interior vertices of $P_x[x, t_x]$ by the choice of $s_x t_x$. They also cannot be incident to v_x or any interior vertex of $P_x[t_x, v_x]$, otherwise Case 3a or 3b would apply.

Next we show that there is an edge between the interior of $P_x[x, t_x]$ and the interior of $P_x[t_x, v_x]$. Suppose to the contrary that there is no such edge. Consider the edges going out of $P_x[x, t_x]$. They are not incident to any vertex outside of P_x. By our assumption, they are not incident to any interior vertex of $P_x[t_x, v_x]$. They are not incident to any interior vertex of $P_x[u_x, x]$ by the choice of $s_x t_x$. They are also not incident to u_x because of planarity: path $P_x[v_x, t_x] \cup \{s_x t_x\} \cup P_x[s_x, x] \cup \{xy\}$ separates them from u_x. Hence all edges out of $P_x[x, t_x]$ are incident to v_x. Note that $s_x t_x$ and the edges out of $P_x[x, t_x]$ lie on opposite sides of path P_x by planarity. Now consider the edges out of $P_x[t_x, v_x]$. They are not incident to any interior vertex of $P_x[u_x, x]$ as noted above. They are not incident to interior vertices of $P_x[x, t_x]$ by our assumption. They are also not incident to x because $s_x t_x$ and an edge going out of $P_x[x, t_x]$ to v_x separate them from x. Therefore, all edges going out of $P_x[t_x, v_x]$ are incident to u_x (Fig. 8, right). This means that P_x is a monster, contradicting the conditions in Case 3. We conclude that there is an edge going out of $P_x[t_x, v_x]$ to an interior vertex of $P_x[x, t_x]$.

Let $s't'$ be an edge between an interior vertex s' of $P_x[t_x, v_x]$ and an interior vertex t' of $P_x[x, t_x]$. If any edge going out of $P_x[x, t_x]$ is in $A \setminus E_i$, then let one of them be $s't'$, otherwise we can choose $s't'$ arbitrarily. Refer to Fig. 7(c). Construct E_{i+1} from E_i in three steps as follows: set $E_i := E_i \cup \{xy\}$, call operation `Disperse` for path $P_x[u_x, x]$ and edge $s_x t_x$, and then call `Disperse` for path $P_x[t_x, v_x]$ and edge $s't'$ (which introduces a new hub at t'). Invariants $\mathbf{I}_1$–$\mathbf{I}_3$ are maintained in all three steps by Proposition 3.3. For Invariant $\mathbf{I}_4$, the first step increases $|E_i| - |U_i| + b_i$ by one, the second step maintains it by Proposition 3.3, and the third step decreases it because of the extra hub at t' (*c.f.*, Remark 3.2). Altogether invariant $\mathbf{I}_4$ is maintained. This completes the description of Case 3d.

While $U_i \neq V$, we can apply Case 1, 2, or 3 and increase the number of hubs. If $U_i = V$, then $G_i = (V, G_i)$ is a 3-connected graph with $A \subseteq E_i \subseteq C$ and $|E_i| \leq 2n - 2$, as required. This completes the proof of Lemma 3.1. □

Corollary 3.4. *Every 3-augmentable planar straight-line matching with $n \geq 4$ vertices can be augmented to a 3-connected* PSLG *which has at most $2n - 2$ edges.*

Proof. Let $G_A = (V, A)$ be a 3-augmentable planar straight-line matching with $n \geq 4$ vertices. By the results of Hoffmann and Tóth [5], there is a PSLG Hamiltonian cycle H on the vertex set V that does not cross any edge in A. Since the Hamiltonian cycle H is crossing-free, none of its edges is a chord of the convex hull of vertices (otherwise the removal of this edge would disconnect H). Hence both (V, H) and $(V, A \cup H)$ are 3-augmentable [14]. That is, there is a 3-connected PSLG $G_C = (V, C)$ such that $A \cup H \subset C$. Lemma 3.1 completes the proof. □

4 A Collection of Simple Polygons

In this section, we prove Theorem 2. We are given a 2-regular PSLG $G_A = (V, A)$ with $n \geq 4$ vertices and n edges. If G_A is 3-augmentable, then it is contained in some 3-connected PSLG $G_C = (V, C)$, say a triangulation of G_A, which may have up to $3n - 6$ edges. We will construct an augmentation $G_B = (V, B)$, $A \subseteq B \subseteq C$, with $|B| \leq 2n$ edges.

Lemma 4.1. *Let $G_A = (V, A)$ be a 2-regular graph with $n \geq 4$ vertices, and let $G_C = (V, C)$ be a 3-connected* PSLG *with $A \subseteq C$ such that all bounded faces are triangles. Then $G_A = (V, A)$ can be augmented to a 3-connected graph $G_B = (V, B)$ with $A \subseteq B \subseteq C$, such that $|B| \leq 2n$.*

Proof. Since G_C is 3-connected, its outer face is a simple polygon, which we denote by Q_C, and at least one vertex of V lies in the interior of Q_C. We construct a 3-connected graph G_B, $A \subseteq B \subseteq C$, incrementally. We maintain a 2-connected graph $G_i = (V_i, E_i)$ with $V_i \subseteq V$ and $E_i \subseteq C$. We also maintain a set $U_i \subseteq V$ of vertices, called *hubs*, which is the set of *all* vertices in V_i with degree 3 or higher in G_i. The hubs naturally decompose G_i into a set $\mathcal{P}_i$ of paths in G_i between hubs. We maintain the following invariants for G_i.

- $\mathbf{J}_1$ $Q_C \subseteq E_i \subseteq C$;
- $\mathbf{J}_2$ U_i is 3-linked in G_i;
- $\mathbf{J}_3$ every bounded face of G_i is incident to at least three vertices in U_i;
- $\mathbf{J}_4$ no edge of $C \setminus E_i$ joins a pair of vertices of any path in $\mathcal{P}_i$.

Initially G_0 will have 4 vertices, and we incrementally augment it with new edges *and* vertices, until we have $U_i = V$. The vertex sets V_i, U_i, and the edge set E_i will monotonically increase during this algorithm, and we gradually add all edges of A to E_i. When our algorithm terminates and $U_i = V$, the graph G_i is a 3-connected subgraph of G_C, which contains all edges of A. Whenever we add an edge $e \in C \setminus A$ to E_i, we charge e to one of the endpoints of e so that every vertex is charged at most once. This charging scheme ensures that we add at most n edges from $C \setminus A$. Together with the n edges of G_A, we obtain a 3-connected augmentation with at most $2n$ edges.

Initialization. We construct the initial graph G_0 with $|U_0| = 4$ hubs. Consider the 3-connected PSLG G_C where Q_C is the boundary of the outer face. Let $v \in V$ be a vertex in the interior of Q_C, and let u and w be its two neighbors in the 2-regular graph G_A.

Construct an auxiliary graph $G_C^* = (V \cup \{a, b\}, C^*)$, with $C \subset C^*$, as follows. The edges of G_C^* are all edges in C, edges au, av, and aw, and edges connecting the auxiliary vertex b to all vertices of the outer face Q_C (Fig. 9(b)). By Lemma 2.1, G_C^* is 3-connected (albeit not necessarily planar). Hence G_C^* contains three independent paths between a and b. Fix three independent paths of minimal total length. The minimality implies that no two nonconsecutive vertices in any of the three paths are joined by an edge of C. Replace the edges au and aw with vu and vw, respectively, to obtain three independent paths *in* C from $v \in V$ to three distinct vertices of the outer face Q_C, such that two of these paths leave v along edges of A. Denote by P_1, P_2, P_3 the three paths, with endpoints p_1, p_2, p_3 along Q_C, respectively.

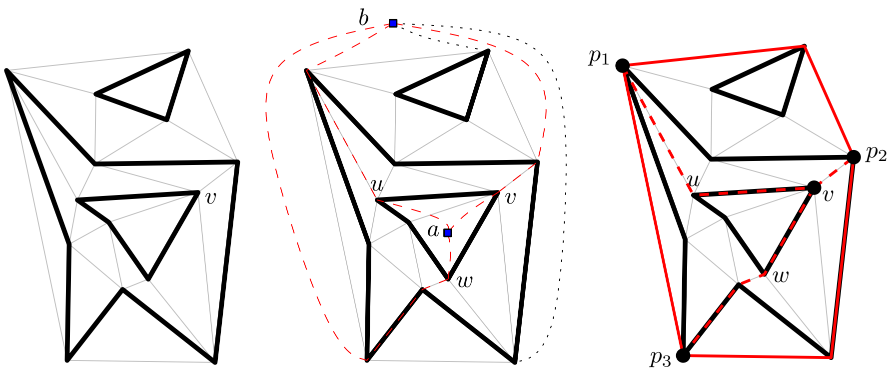

(a) (b) (c)

Figure 9: (a) A 2-regular PSLG G_A (black) in a 3-connected triangulation G_C (gray). (b) Graph G_C^* with two auxiliary vertices, a and b, is also 3-connected. (c) Three independent paths from v to three boundary points p_1, p_2, and p_3.

Let our initial graph $G_0 = (V_0, E_0)$ consist of all edges and vertices of $Q_C \cup P_1 \cup P_2 \cup P_3$. There are exactly four vertices of degree 3, namely $U_0 = \{v, p_1, p_2, p_3\}$, which are 3-linked in G_0.

Each of the three bounded faces of G_0 is incident to 3 hubs. So G_0 satisfies invariants $\mathbf{J}_1$–$\mathbf{J}_3$. For invariant $\mathbf{J}_4$, note also that no edge in C joins nonadjacent vertices of Q_C, otherwise G_C would not be 3-connected.

Let us estimate how many edges of G_0 are from $C \setminus A$. Orient Q_C counterclockwise, and charge every edge $e \in C \setminus A$ along Q_C to its origin. Clearly, every vertex in Q_C is charged at most once. Direct the paths P_1, P_2, and P_3 from v to p_1, p_2, and p_3; and charge each edge $e \in C \setminus A$ along the paths to its origin. Since two paths leave v along edges of A, vertex v is charged exactly once. All interior vertices of the three paths are charged at most once, because the paths are independent.

Our algorithm proceeds in three phases.

Phase 1. In the first phase of our algorithm we augment $G_i = (V_i, E_i)$ until $V_i = V$, but at the end of this phase some edges of A may still not be contained in E_i. We augment $G_i = (V_i, E_i)$ with new edges and vertices incrementally. It is enough to describe a general step of this phase.

Pick an arbitrary vertex $v \in V \setminus V_i$. We will augment G_i to include v (and possibly other vertices). Our argument is similar to the initialization. Let Q_v denote the boundary of the face of G_i that contains v. Let U_v denote the set of hubs along Q_v. We have $|U_v| \geq 3$ by $\mathbf{J}_3$, and Q_v consists of at least three paths from $\mathcal{P}_i$ between consecutive hubs along Q_v. Let G_v be the subgraph of C that contains all edges and vertices of G_C lying in the closed polygonal domain bounded by Q_v. Note that G_v is a (Steiner) triangulation of the simple polygon Q_v. This implies that G_v is 2-connected and every 2-cut of G_v consists of a pair of vertices joined by a chord of Q_v. By invariant $\mathbf{J}_4$, however, there is a hub in U_v on each side of every chord of Q_v.

Let u and w be the neighbors of v in the 2-regular graph G_A. Note that both u and w must be vertices of G_v. Construct an auxiliary graph G_v^* as follows. The vertices of G_v^* are the vertices of G_v and two auxiliary vertices, a and b. The edges of G_v^* are the edges of G_v; the edges au, av, and aw; and edges between b and every hub in U_v. We show that G_v^* is 3-connected. First note that none of the 2-cuts of G_v is a 2-cut in G_v^*, since the hubs on the two sides of a chord of Q_v are now joined to b. This implies that the vertices of G_v are 3-linked in G_v^*. Vertices a and b are each joined to three vertices of G_v, and by Lemma 2.1 there are three independent paths between any two vertices of G_v^*.

Choose three independent paths in G_v^* between a and b of minimal total length. The minimality implies that each path goes from a to a vertex along Q_v, then follows Q_v to a hub in U_v, and reaches b along a single edge from Q_v. In particular, no two nonconsecutive vertices of any of the three paths are joined by an edge of G_v^* (*i.e.*, no shortcuts). Replace the edges au and aw with edges vu and vw, respectively, to obtain three independent paths from v to three distinct hubs along Q_v, such that two of these paths leave v along edges of A. Denote by P_1, P_2, and P_3 the initial portions of the paths between a and Q_v; and let p_1, p_2, and p_3 be their endpoints on Q_v (these endpoints are not necessarily in U_i).

We construct G_{i+1} by augmenting G_i with all vertices and edges of the paths P_1, P_2, and P_3. The new vertices of degree 3 are v and, if they were not hubs already, p_1, p_2, and p_3. In G_{i+1}, three independent paths connect v to three hubs in U_i, so $U_i \cup \{v\}$ is 3-linked in G_{i+1}. Similarly, p_1, p_2, and p_3 are each connected to three hubs in $U_i \cup \{v\}$ along three independent paths. We conclude that $U_{i+1} = U_i \cup \{u, p_1, p_2, p_3\}$ is 3-linked in G_{i+1}. We can construct $\mathcal{P}_{i+1}$ from $\mathcal{P}_i$ by adding the three new paths P_1, P_2, and P_3; and decomposing the paths in $\mathcal{P}_i$ containing p_1, p_2, and p_3 if necessary.

Paths P_1, P_2, and P_3 decompose a face of G_i into three faces, each of which is incident to at least three hubs of U_{i+1}. So invariants $\mathbf{J}_1$–$\mathbf{J}_4$ hold for G_{i+1}. It remains to charge the new edges of E_{i+1} taken from $C \setminus A$ to some new vertices in V_{i+1}. Direct the paths P_1, P_2, and P_3 from v to p_1, p_2, and p_3, respectively; and charge any new edge $e \in C \setminus A$ to its origin. Each new vertex of V_{i+1} is charged at most once: v is charged at most once because two incident new edges are contained in A; and any other new vertex is charged at most once because the paths P_1, P_2, and P_3 are independent.

Phase 2. In the second phase, we augment $G_i = (V, E_i)$ with edges of $A \setminus E_i$ successively until $A \subseteq E_i$. We can add all edges of A at no charge, we only need to check that that invariants $\mathbf{J}_1$–$\mathbf{J}_4$ are maintained. We describe a single step of the augmentation. Consider an edge $pq \in A \setminus E_i$. Let $G_{i+1} = (V, E_{i+1})$ with $E_{i+1} = E_i \cup \{pq\}$ and $U_{i+1} = U_i \cup \{p, q\}$. By $\mathbf{J}_2$–$\mathbf{J}_4$, and Lemma 2.1, U_{i+1} is 3-linked in G_{i+1}. The edge pq subdivides a bounded face of G_i into two faces of G_{i+1}. Since pq does not join two vertices of the same path in $\mathcal{P}_i$ by invariant $\mathbf{J}_4$, both new faces are incident to at least three hubs in U_{i+1} (including p and q). The paths in $\mathcal{P}_{i+1}$ are obtained from $\mathcal{P}_i$ by adding the 1-edge path pq, and possibly decomposing the paths containing p and q into two. Since $\mathcal{P}_i$ satisfies $\mathbf{J}_4$, no edge in C joins two nonconsecutive vertices of any path in $\mathcal{P}_{i+1}$, either. So invariants $\mathbf{J}_1$–$\mathbf{J}_4$ hold for G_{i+1}.

Phase 3. We have a graph $G_i = (V, E_i)$ with $A \subseteq E_i \subseteq C$, where U_i is the set of vertices of degree 3 or higher. Let $W = V \setminus U_i$ be the set of vertices that have degree 2 or less than 2 in G_i. Since G_A is 2-regular, and $A \subseteq E_i$, every vertex in W has degree 2 and is incident to two edges in A. Since we charged every edge $E_i \cap (C \setminus A)$ to an incident vertex, no vertex in W has been charged so far. Apply Lemma 2.2 to augment $G_i = (V, E_i)$ to a 3-connected graph G_B with $|W|$ additional edges. Charge the new edges to the vertices in W. At the end of phase 3, every vertex is charged at most once, and we obtain a 3-connected PSLG with at most $|A| + n = 2n$ edges. □

5 Obstacles in a Container

In this section, we consider augmenting a PSLGs $G_0 = (V, E_0)$ with $n \geq 6$ vertices that consists of a set of interior-disjoint convex polygons (obstacles) in the interior of a triangular container. Since no edge is a proper chord of the convex hull, every such PSLG is 3-augmentable [14], and by Theorem 2 it can be augmented to a 3-connected PSLG with at most $2n$ edges. We believe, however, that significantly fewer edges are sufficient for 3-connectivity augmentation. The best lower bounds we were able to construct require fewer than $2n - 2$ edges.

When there is only one convex obstacle, three edges are obviously required for connecting it to the container. However, for $k \in \mathbb{N}$ convex obstacles at least $3k - 1$ edges are necessary in the worst case. Our lower bound construction is depicted in Figure 10(a). It includes one large convex obstacle which hides one small obstacle behind each side (except the base) such that each small obstacle can "see" only three different vertices (the top vertex of the container and two adjacent vertices of the large obstacle). Thus, we need three edges for each small obstacle and only two edges for the larger obstacle, connecting its two bottom vertices to the two endpoints of the base of the container.

The large obstacle in the above construction is a convex k-gon, and so the lower bound $3k - 1$ does not hold if every obstacle has at most s sides, for some fixed $3 \leq s < k$. In that case we use a similar construction, in which a big s-sided obstacle hides $s - 1$ smaller obstacles behind all its

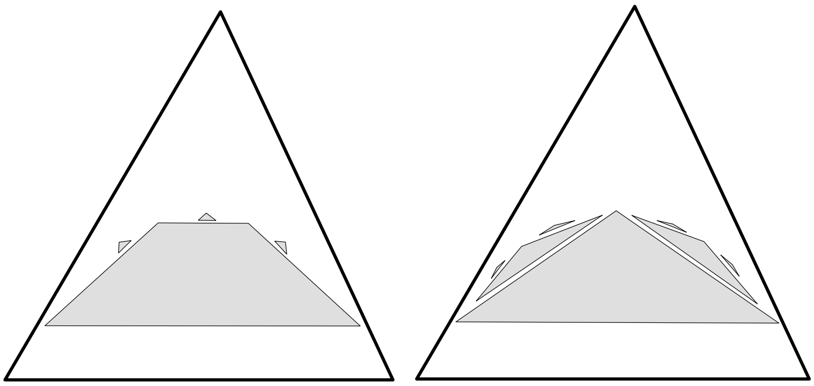

(a) (b)

Figure 10: (a) 3-connectivity augmentation for k interior-disjoint convex obstacles in a triangular container requires $3k-1$ new edges. (b) For k interior-disjoint triangular obstacles in a triangular container, we need $(5k+1)/2$ new edges.

sides except one, and the construction is repeated recursively. This construction corresponds to a complete tree with branching factor $s-1$, in which the smaller obstacles are the children of a larger obstacle. For a fixed value of s, we set h as the height of the complete $(s-1)$-ary tree. Thus, the number of obstacles,

$$k = \frac{(s-1)^h - 1}{s-2},$$

can be as high as we desire. The number of leaves in the tree is $(s-1)^{h-1}$. A simple manipulation of this equation shows that this number equals $k - \frac{k-1}{s-1}$. Hence, the number of internal nodes in the tree is $\frac{k-1}{s-1}$. For the 3-connectivity augmentation, each leaf obstacle needs at least s new edges and each nonleaf obstacle needs at least two new edges. The total number of edges required is at least

$$s\left(k - \frac{k-1}{s-1}\right) + 2\left(\frac{k-1}{s-1}\right) = sk - \frac{s-2}{s-1}(k-1) = (n-3) - \frac{s-2}{s-1}\cdot\left(\frac{n-3}{s} - 1\right),$$

which ranges from $\frac{5}{6}n - \frac{5}{2}$ to $n - O(\sqrt{n})$ for $3 \leq s \leq \sqrt{n-3}$. Figure 10(b) depicts this lower bound construction for $s = 3$.

6 Discussion

We have shown that a 1-regular (resp., 2-regular) PSLG with n vertices, where no edge is a chord of the convex hull, can be augmented to a 3-connected PSLG which has at most $2n-2$ (resp., $2n$) edges. We conjecture that our results generalize to PSLGs with maximum degree at most 2 (Conjecture 1.1).

The bound of $2n-2$ for the number of edges is the best possible in general, but it may be improved if few vertices lie on the convex hull, and the components of the input graph are interior-disjoint convex obstacles, possibly with a container. It remains an open problem to derive tight extremal bounds for 3-connectivity augmentation for (i) 1-regular PSLGs with n vertices, h of which lie on the convex hull; and (ii) 2-regular PSLGs formed by $\frac{n}{s}$ interior-disjoint convex polygons, each with s vertices for $s \geq 3$.

The 3-connectivity augmentation problem (finding the *minimum* number of new edges for a given PSLG) is known to be NP-hard [13]. However, the hardness proof does not apply to 1- or 2-regular polygons. It is an open problem whether the connectivity augmentation remains NP-hard restricted to these cases.

We have compared the number of edges in the resulting 3-connected PSLGs with the benchmark $2n-2$, which is the best possible bound for 0-, 1-, and 2-regular PSLGs. More generally, for a 3-augmentable PSLG $G_0=(V,E_0)$ with $n\geq 4$ vertices, let $f(G_0)=|E_1|$ be the minimum number of edges in a 3-connected augmentation (V,E_1) of the *empty* PSLG $(V,\emptyset)$; and let $g(G_0)=|E_2|$ be the minimum number of edges in a 3-connected augmentation (V,E_2), $E_0\subseteq E_2$, of the PSLG G_0. It is clear that $f(G_0)\leq g(G_0)$. With this notation, we can characterize the PSLGs G_0 where *all* edges in E_0 are "useful" for 3-connectivity: these are the PSLGs for which $f(G_0)=g(G_0)$ is possible. In general, it would be interesting to study the behavior of the difference $g(G_0)-f(G_0)$.

Finally, we note here that augmenting a 1-regular PSLG to a 2-regular PSLG has been considered by Aichholzer *et al.* [1]. The problem is always feasible if the input PSLG (perfect matching) has an even number of edges [8]. For the odd case, however, neither combinatorial characterizations nor hardness results are known for the corresponding decision problem.

Acknowledgement. We are grateful to the anonymous referee who pointed out several errors and omissions in an earlier version of the proof of Lemma 3.1.

References

[1] O. Aichholzer, S. Bereg, A. Dumitrescu, A. García, C. Huemer, F. Hurtado, M. Kano, A. Márquez, D. Rappaport, S. Smorodinsky, D. L. Souvaine, J. Urrutia, and D. R. Wood, Compatible geometric matchings, *Computational Geometry: Theory and Applications* **42** (2009), 617–626.

[2] M. Al-Jubeh, M. Ishaque, K. Rédei, D. L. Souvaine, and C. D. Tóth, Augmenting the edge connectivity of planar straight line graphs to three, *Algorithmica* **61** (2011), 971–999.

[3] K. P. Eswaran and R. E. Tarjan, Augmentation problems, *SIAM J. on Computing* **5** (1976), 653–665.

[4] A. García, F. Hurtado, C. Huemer, J. Tejel, and P. Valtr, On triconnected and cubic plane graphs on given point sets, *Computational Geometry: Theory and Applications* **42** (2009), 913–922.

[5] M. Hoffmann and C. D. Tóth, Segment endpoint visibility graphs are Hamiltonian, *Computational Geometry: Theory and Applications* **26** (2003), 47–68.

[6] T.-S. Hsu and V. Ramachandran, On finding a minimum augmentation to biconnect a graph, *SIAM J. on Computing* **22** (5) (1993), 889–912.

[7] T.-S. Hsu, Simpler and faster biconnectivity augmentation, *J. of Algorithms* **45** (1) (2002), 55–71.

[8] M. Ishaque, D. L. Souvaine, and C. D. Tóth, Disjoint compatible geometric matchings, in *Proc. 27th Sympos. on Comput. Geom.*, ACM Press, 2011, pp. 125–134.

[9] B. Jackson and T. Jordán, Independence free graphs and vertex connectivity augmentation, *J. of Combinatorial Theory, Ser. B* **94** (2005), 31–77.

[10] G. Kant and H. L. Bodlaender, Planar graph augmentation problems, *Proc. 2nd Workshop on Algorithms and Data Structures*, vol. 519 of *Lecture Notes in Computer Science*, Springer, 1991, pp. 286–298.

[11] D. L. Millman, M. O'Meara, J. Snoeyink, and V. Verma, Maximum geodesic routing in the plane with obstacles, *Proc. 22nd Canadian Conf. on Computational Geometry*, Winnipeg, MB, 2010, pp. 107–108.

[12] J. Plesník, Minimum block containing a given graph, *Archiv der Mathematik* **21** (1976), 668–672.

[13] I. Rutter and A. Wolff, Augmenting the connectivity of planar and geometric graphs, *P*roc. Int. Conf. on Topological and Geometric Graph Theory, Paris, France, *Electronic Notes in Discrete Mathematics*, **31** (2008), 53–56.

[14] C. D. Tóth and P. Valtr, Augmenting the edge connectivity of planar straight line graphs to three, in: *13th Spanish Meeting on Computational Geometry*, Zaragoza, Spain, 2009.