Geometric Biplane Graphs I: Maximal Graphs*

Alfredo García[†] Ferran Hurtado[‡] Matias Korman[‡] Inês Matos[§] Maria Saumell[¶]
Rodrigo I. Silveira[‖,‡] Javier Tejel[†] Csaba D. Tóth[**]

Abstract

We study biplane graphs drawn on a finite point set S in the plane in general position. This is the family of geometric graphs whose vertex set is S and can be decomposed into two plane graphs. We show that two maximal biplane graphs—in the sense that no edge can be added while staying biplane—may differ in the number of edges, and we provide an efficient algorithm for adding edges to a biplane graph to make it maximal. We also study extremal properties of biplane graphs such as the maximum number of edges and the largest minimum degree of biplane graphs over n-element point sets. In a companion paper we study how to draw a biplane graph on a given point set S, or how to augment a given biplane graph on S, in such a way that the resulting graph is biplane and has good connectivity properties.

1 Introduction

In a *geometric graph* $G = (V, E)$, the vertices are distinct points in the plane in general position, and the edges are straight line segments between pairs of vertices. A *plane graph* is a geometric graph in which no two edges cross. Every (abstract) graph has a realization as a geometric graph (by simply mapping the vertices into distinct points in the plane, no three of which are collinear), and every planar graph can be realized as a plane graph by Fáry's theorem. The number of n-vertex labeled planar graphs is at least $27.22^n \cdot n!$ [18]. However, there are only $2^{O(n)}$ plane graphs on any given set of n points in the plane [2, 25].

We consider a generalization of plane graphs. A geometric graph $G = (V, E)$ is *k-plane* for some $k \in \mathbb{N}$ if it admits a partition of its edges $E = E_1 \dot\cup \ldots \dot\cup E_k$ such that $G_1 = (V, E_1), \ldots, G_k = (V, E_k)$ are each plane graphs. Let S be a point set in the plane in general position, that is, no three points in S are collinear. Denote by $\mathcal{G}_k(S)$ the family of k-plane graphs with vertex set S. With this terminology, $\mathcal{G}_1(S)$ is the family of plane graphs with vertex set S, and $\mathcal{G}_2(S)$ is the family of 2-plane graphs (also known as *biplane graphs*) with vertex set S.

In this and a companion paper, we study $\mathcal{G}_2(S)$ and contrast combinatorial properties of plane graphs $\mathcal{G}_1(S)$ and biplane graphs $\mathcal{G}_2(S)$. If only plane graphs drawn on S are considered, there are limitations

*A. G., F. H., M. K., R.I. S. and J. T. were partially supported by ESF EUROCORES programme EuroGIGA, CRP ComPoSe: grant EUI-EURC-2011-4306, and by project MINECO MTM2012-30951. F. H., and R.I. S. were also supported by project Gen. Cat. DGR 2009SGR1040. M. K. was supported by the Secretary for Universities and Research of the Ministry of Economy and Knowledge of the Government of Catalonia and the European Union. I. M. was supported by FEDER funds through COMPETE–Operational Programme Factors of Competitiveness, CIDMA and FCT within project PEst-C/MAT/UI4106/2011 with COMPETE number FCOMP-01-0124-FEDER-022690. M. S. was supported by ESF EuroGIGA project ComPoSe as F.R.S.-FNRS - EUROGIGA NR 13604. R.I. S. was also partially funded by the FP7 Marie Curie Actions Individual Fellowship PIEF-GA-2009-251235 and by FCT through grant SFRH/BPD/88455/2012. C. T. was supported in part by NSERC (RGPIN 35586) and NSF (CCF-0830734). A. G. and J. T. were also supported by project E58-DGA.

[†]Dept. de Métodos Estadísticos, Universidad de Zaragoza, Spain. {olaverri,jtejel}@unizar.es
[‡]Dept. de Matemàtica Aplicada II, Universitat Politècnica de Catalunya, Spain. {ferran.hurtado, matias.korman, rodrigo.silveira}@upc.edu
[§]Dept. de Matemática & CIDMA, Universidade de Aveiro, Portugal. ipmatos@ua.pt
[¶]Dept. d'Informatique, Université Libre de Bruxelles, Belgium. maria.saumell.m@gmail.com
[‖]Dept. de Matemática, Universidade de Aveiro, Portugal, rodrigo.silveira@ua.pt.
[**]Dept. of Mathematics, California State University Northridge, USA, and Dept. of Mathematics and Statistics, University of Calgary, Canada. cdtoth@acm.org

on achieving some desirable properties, such as high connectivity, as it is known that every plane graph H has a vertex with degree at most 5, hence $\kappa(H) \leq \lambda(H) \leq \delta(H) \leq 5$ (we use standard graph theory notation as in [8]). It is natural to expect that significantly better values can be obtained if the larger family $\mathcal{G}_2(S)$ is used. This is precisely the topic we explore in these papers, mostly focusing on graph size and graph connectivity.

Related concepts. Note that the above generalization of plane graphs is reminiscent to, although more restrictive, than the notion of thickness, geometric thickness, and book thickness, which are defined for abstract graphs. We recall their definitions for ease of comparison. The *thickness* of an (abstract) graph $G = (V, E)$ is the smallest $k \in \mathbb{N}$ such that G admits an edge partition $E = E_1 \dot{\cup} \ldots \dot{\cup} E_k$ with the property that $G_1 = (V, E_1), \ldots, G_k = (V, E_k)$ are each planar graphs. The *geometric thickness* of an (abstract) graph $G = (V, E)$ is the smallest $k \in \mathbb{N}$ such that G admits an edge partition $E = E_1 \dot{\cup} \ldots \dot{\cup} E_k$ satisfying that $G_1 = (V, E_1), \ldots, G_k = (V, E_k)$ can be simultaneously embedded as plane graphs where the vertex set is mapped to a common labeled point set. The *book thickness* is a restricted version of the geometric thickness where $G_1, \ldots, G_k$ are simultaneously embedded on a point set in convex position.

Notice that every k-plane graph, if interpreted as an abstract graph, has geometric thickness at most k, but in addition we are given a specific embedding in the plane in which the decomposition into k plane layers is possible. In other words, the term k-plane graph refers to a geometric object, a drawing, while having geometric thickness k is a property of the underlying abstract graph.

For disambiguation, we also mention two additional notions, which are commonly used in the graph drawing community, but have little to do with our subject matter. An (abstract) graph is called *k-planar* if it has a drawing in the plane (where the edges are Jordan arcs) such that each edge crosses at most k other edges. It is NP-hard to recognize already 1-planar graphs [22]. The other notion worth mentioning is *1-plane* that is used for a specific geometric drawing of a 1-planar graph [11] in which edges are crossed at most once. Therefore, both these notions have a different meaning from the definition of k-plane graphs introduced above.

Prior work and organization of the paper. We call *maximum graphs* in $\mathcal{G}_2(S)$ those that have the largest number of edges; this concept should not be confused with *maximal graphs* in $\mathcal{G}_2(S)$, which are those in which no additional edges can be added because the resulting graph would not be biplane. In contrast to plane graphs, two maximal biplane graphs on the same point set do not have necessarily the same number of edges. In Section 2 we give several interesting properties of biplane graphs.

Algorithmic issues are studied in Section 3. First, we give an algorithm for determining whether a given geometric graph is biplane or not. We then show how to augment a biplane graph to a maximal biplane supergraph. This result belongs to the family of problems called *graph augmentation*, in which one would like to add new edges, ideally as few as possible, to a given graph in such a way that some desired property is achieved. There has been extensive work on augmenting a disconnected plane graph to a connected one (see [20] for a recent survey) or on achieving good connectivity properties [1, 3, 4, 24, 26].

In Section 4, we study the highest values of connectivity that can be attained for the graphs in $\mathcal{G}_2(S)$ over all n-element points set S. Similar extremal problems have been considered for thickness-two graphs [6, 21], including the geometric thickness version. Due to lack of space, some of the proofs have been omitted in this reduced version.

In a companion paper [16] we consider several problems on augmenting plane graphs to biplane supergraphs with higher connectivity, including the case in which the input is only a point set S and the goal is to construct a *good* biplane graph on S. These problems are closely related to the results we present here, and have also received substantial attention for the case of plane graphs [9, 15, 17].

2 Fundamental Properties of Maximal Biplane Graphs

A (geometric) graph $G = (V, E)$ is *maximal* (a.k.a. *edge-maximal*) in a family of graphs $\mathcal{F}$ if there is no graph $G' = (V, E')$ in $\mathcal{F}$ such that $E \subset E'$. Some of our main results in this and the companion paper concern finding a biplane graph of high connectivity or high vertex degree in $\mathcal{G}_2(S)$ for point sets S.

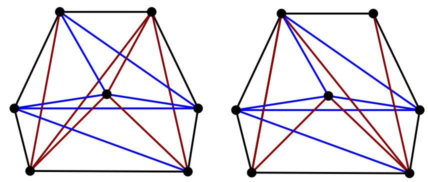

Figure 1: Two maximal biplane graphs on the same point set. The left graph is a maximum biplane graph with 18 edges, whereas the right one is maximal with 17 edges.

Since the addition of new edges does not decrease the vertex connectivity, we can restrict our attention to maximal graphs in $\mathcal{G}_2(S)$. In this section, we summarize some of the basic properties of maximal biplane graphs on a given point set S.

Recall that a maximal plane graph in $\mathcal{G}_1(S)$ is a *triangulation*, that is, a plane graph where all bounded faces are triangles, and the boundary of the outer face is the same as the boundary of the convex hull $\mathrm{ch}(S)$. We highlight the fact that two maximal graphs in $\mathcal{G}_2(S)$ may have a different number of edges, that is, not every maximal biplane graph is maximum (see Figure 1). Our first observation is that a maximal biplane graph is always the union of two maximal plane graphs.

Lemma 1 *If $G = (S,E)$ is a maximal biplane graph in $\mathcal{G}_2(S)$, then there are two triangulations $T' = (S,E')$ and $T'' = (S,E'')$ such that $E = E' \cup E''$.*

Proof. By definition, $G = (S,E)$ has an edge partition $E = E_1 \dot{\cup} E_2$, where $G_1 = (S,E_1)$ and $G_2 = (S,E_2)$ are plane graphs. Augment G_1 and G_2, independently, to maximal plane graphs $T' = (S,E')$ and $T'' = (S,E'')$; thus T' and T'' are triangulations. By construction, we have $E \subset E' \cup E''$. The geometric graph $(S, E' \cup E'')$ is biplane by definition. Hence $E = E' \cup E''$, otherwise G would not be maximal. $\square$

The two triangulations, $T' = (S,E')$ and $T'' = (S,E'')$, share some edges. The edges of the convex hull $\mathrm{ch}(S)$ are always part of both triangulations, but they may also share some interior edges. There is a simple characterization of shared edges in terms of edge flips. An edge e in a triangulation is *flippable* if the union of the two adjacent faces (triangles) is a convex quadrilateral.

Lemma 2 *Let $G = (S,E)$ be a maximal biplane graph in $\mathcal{G}_2(S)$ such that $E = E' \cup E''$, where $T' = (S,E')$ and $T'' = (S,E'')$ are two triangulations. If $e \in E' \cap E''$, then e is flippable in neither T' nor T''. In particular, any maximal biplane graph is 3-connected*

Proof. Suppose, on the contrary, that $e \in E' \cap E''$ is flippable in T' (the case that e is flippable in T'' is analogous). We can modify T' by flipping edge e. Specifically, let f be the other diagonal of the convex quadrilateral formed by the two faces of T' adjacent to e, and define a new triangulation $T''' = (S,E''')$ with $E''' = (E' \setminus \{e\}) \cup \{f\}$. It is clear that $(S, E'' \cup E''')$ is biplane, and it contains edge f and all edges in E (including $e \in E''$). Hence it is a biplane graph strictly larger than G, contradicting the maximality of G. Three-connectivity follows from the fact that a separating chord in a triangulation is always flippable. $\square$

We now study the minimum number of edges that a biplane graph must have in order to be maximal. It is known that every triangulation in $\mathcal{G}_1(S)$ has $3n-h-3$ edges, where $n = |S| \geq 3$ and $h \geq 3$ is the number of vertices of the convex hull $\mathrm{ch}(S)$. Since a maximal biplane graph $G \in \mathcal{G}_2(S)$ is the union of two triangulations, T' and T'', that share the convex hull edges, G has at most $6n-3h-6 \leq 6n-15$ edges. Hutchinson et al. [21] proved that every biplane graph in $\mathcal{G}_2(S)$ has at most $6n-18$ edges for $n \geq 8$. That is, when $h = 3$, the triangulations T' and T'' have to share at least 3 interior edges. In the remainder of this section we establish lower bounds for the number of edges in a maximum graph in $\mathcal{G}_2(S)$ in terms of $n = |S|$ and h, the number of vertices of $\mathrm{ch}(S)$.

Theorem 3 *For every set S of $n \geq 3$ points in the plane, every maximal graph in $\mathcal{G}_2(S)$ has at least $3n-6$ edges. Moreover, this bound is the best possible.*

Proof. Let S be a set of $n \geq 3$ points in the plane, and let $h \geq 3$ be the number of vertices of $\text{ch}(S)$. Let $G_1 = (S, E_1)$ be an arbitrary triangulation of S, with $3n-h-3$ edges. Hoffmann et al. [19] proved that every triangulation of S contains at least $\max(\frac{n}{2}-2, h-3)$ flippable edges. (In fact they prove a stronger statement, that there is a set of at least $\max(\frac{n}{2}-2, h-3)$ *pseudo-flippable* edges, but each edge in that set is also flippable.)

Now, let $G = (S, E)$ be a maximal biplane graph in $\mathcal{G}_2(S)$, and suppose that $T_1 = (S, E_1)$ and $T_2 = (S, E_2)$ are two triangulations such that $E = E_1 \cup E_2$. By Lemma 2, an edge in $E_1 \cap E_2$ is flippable in neither T_1 nor T_2. Thus all flippable edges that are in E_1 must be in $E_1 \setminus E_2$. Hence, $|E_1 \setminus E_2| \geq \max(\frac{n}{2}-2, h-3)$, and so $|E| = |E_2| + |E_1 \setminus E_2| \geq 3n-h-3+\max(\frac{n}{2}-2, h-3) \geq \max(\frac{7n}{2}-h-5, 3n-6) \geq 3n-6$.

The bound $3n-6$ is the best possible whenever S is a set of $n \geq 3$ points in convex position. In this case, every biplane graph is planar (by Lemma 1(i) in the companion paper [16]) and has at most $3n-6$ edges for $n \geq 3$. □

Theorem 4 *For every set S of $n \geq 3$ points in the plane whose convex hull has h vertices, the largest graph in $\mathcal{G}_2(S)$ has at least $4n-h-6$ edges if $h \geq 4$ or $n = 3$; and at least $4n-h-7$ edges if $h = 3$ and $n > 3$. Moreover, these bounds are the best possible.*

Proof. We proceed by induction on n, the number of points. The base case is $n = h$, where the union of two triangulations of a convex n-gon gives a biplane graph with $n + 2(n-3) = 3n-6 = 4n-h-6$ edges.

Suppose now that $n > h$, and the claim holds for every set of $n-1$ points with a convex hull of h vertices. Let $s \in S$ be a rightmost point in the interior of $\text{ch}(S)$, and let $S' = S \setminus \{s\}$. Let $G' = (S', E')$ be a maximum biplane graph on S'. By induction, G' has at least $4(n-1)-h-6$ edges if $h \geq 4$ or $n-1 = 3$; and at least $4(n-1)-h-7$ edges if $h = 3$ and $n-1 > 3$. By Lemma 1, it is the union of two triangulations $T'_1 = (S', E'_1)$ and $T'_2 = (S', E'_2)$.

We construct a biplane graph $G = (S, E)$ by augmenting G' with the new vertex s and some incident edges. If $h = 3$ and $n = 4$, then G' is a triangle, and s can only be joined to the 3 vertices of G'. Hence G has $6 = 4n-h-7$ edges, as required.

If $h \geq 4$ or $n \geq 5$, we join s to at least 4 vertices of G'. Point s lies in the interior of a triangle Δ' of T'_1, and a triangle Δ'' of T'_2. We can augment T'_1 and T'_2 each with 3 new edges that join s to the corners of Δ' and Δ'', respectively, to obtain two new triangulations T''_1 and T''_2. If Δ' and Δ'' together have at least 4 distinct vertices, then the induction step is complete.

It remains to consider the case in which Δ' and Δ'' together have only 3 distinct vertices, i.e., $\Delta' = \Delta''$. Since $h \geq 4$ or $n \geq 5$, T''_1 is not a wheel, and since s is the rightmost point in the interior of $\text{ch}(S)$, the rightmost vertex of Δ' is a vertex of $\text{ch}(S)$. It is easy to see that one of the edges of Δ' is flippable in T''_1 (proof of this claim is given in Property 3 of [16]). As this flippable edge is common to Δ' and Δ'', after flipping it in T''_1, we augment G' with 4 new edges.

Finally, we note that the bounds are tight in the sense that there are configurations of points for which the largest biplane graph has exactly $4n-h-7$ (resp., $4n-h-6$) edges. The construction consists of $n-1$ (resp., $n-2$) points on a parabola $x \to x^2$, and one point far below on line $x = 0$ so that the convex hull of S has 3 vertices (resp., two points on line $x = 0$ one far below, and the other is far above so that the convex hull has 4 vertices). □

3 Constructing Maximal Biplane Graphs

We now consider computational aspects related to biplane graphs. The most fundamental algorithmic question is how to recognize one, thus we start by showing how to determine if a given graph is biplane.

Lemma 5 *Given a geometric graph $G = (S, E)$ with n vertices and m edges, there is an $O(n \log n)$-time algorithm that tests whether G is biplane and produces, if possible, a partition $E = E_1 \dot\cup E_2$ such that both (S, E_1) and (S, E_2) are plane graphs.*

Proof. Let $G = (S, E)$ be a geometric graph with n vertices and m edges. We show that if $m > 2(3n - 6) - 3 = 6n - 15$ and $n \geq 3$, then G cannot be biplane. Indeed, every biplane graph is the union of two triangulations that share at least three hull edges by Lemma 1, and by Euler's formula, we have $m \leq 2(3n - 6) - 3 = 6n - 15$ for $n \geq 3$. Assume now that $m \leq 6n - 15$ or $n \leq 3$. Let G_X be the intersection graph of the open line segments in E, that is, the nodes of G_X correspond to the edges of G, and two nodes are adjacent in G_X if and only if the corresponding edges cross. An edge partition $E = E_1 \dot{\cup} E_2$, where $G_1 = (S, E_1)$ and $G_2 = (S, E_2)$ are plane graphs, corresponds to a bipartition of G_X. Given a intersection graph G_X of m line segments in the plane, an algorithm by Eppstein [12] can test whether G_X is bipartite and produce a bipartition in $O(m \log m)$ time. Recall that $m \in O(n)$, thus the algorithm runs in $O(n \log n)$ time. □

Another natural algorithmic question is how to augment a given biplane graph to a maximal one. That is, given $G \in \mathcal{G}_2(S)$, can we find a maximal graph $G' \in \mathcal{G}_2(S)$ such that $G \subseteq G'$? It is easy to augment a plane graph $G \in \mathcal{G}_1(S)$ to a triangulation: we can augment G with all edges of the convex hull $\text{ch}(S)$, and then triangulate each bounded face independently. However, it is not obvious how to augment two layers $E = E_1 \cup E_2$ into maximal plane graphs independently. In particular, the converse of Lemma 1 is not true: if $T' = (S, E')$ and $T'' = (S, E'')$ are triangulations, then $G = (S, E' \cup E'')$ is not necessarily maximal biplane.

If we are only interested in a polynomial-time algorithm, we can use Lemma 5 to obtain one: consider all $O(n^2)$ edges of the complement of G successively, and augment the graph with the new edge as long as the augmented graph remains biplane. Since each test takes $O(n \log n)$ time, we obtain an algorithm that runs in $O(n^3 \log n)$-time. In the following we show that, by exploiting several geometric properties, we can reduce the running time to $O(n \log n)$.

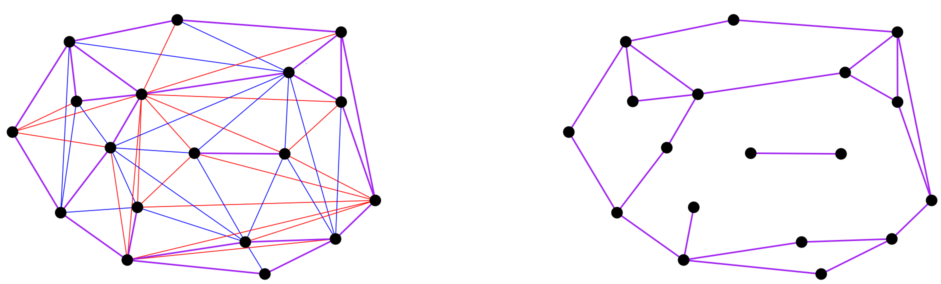

(a) (b)

Figure 2: (a): A biplane graph with three types of edges: only in red triangulation, only in blue triangulation, and in both (shown purple). (b): Subdivision created by purple edges.

We are given a biplane graph $G = (S, E)$ that we would like to augment to a maximal biplane graph. By Lemma 5, G decomposes into two plane graphs, $T_R = (S, R)$ and $T_B = (S, B)$, such that $E = R \cup B$. Without loss of generality, we can assume that T_R and T_B are triangulations. Now we classify the edges in E as *red* if they appear only in R, *blue* if they appear only in B, or *purple* if they appear in both R and B. See Figure 2(a). Let $P = R \cap B$ denote the set of purple edges. Intuitively, we would like to minimize the number of purple edges.

If a purple edge is flippable in T_R or T_B, then we can flip it in one triangulation and retain it in the other, thereby increasing the total number of edges by one (and decreasing the number of purple edges by one). A natural approach would be to flip purple edges whenever they are flippable in either T_R or T_B. However, the decomposition $E = B \cup R$ is not unique: it is possible that a purple edge is not flippable in either triangulation, but there is a different decomposition $E = R' \cup B'$ that admits a flippable edge in $R' \cap B'$. To overcome this difficulty, we introduce the concept of a *color-blind flippable* edge.

Consider the plane graph (S, P) formed by all purple edges. The purple graph (S, P) is a subgraph of both triangulations, T_R and T_B, and contains all convex hull edges. Each bounded face of (S, P) is a weakly simple polygon (possibly with holes), see Figure 2(b). Denote by $\mathcal{F}_1, \ldots, \mathcal{F}_k$ the bounded faces of the purple graph (S, P), where $k \in \mathbb{N}$.

Let $e \in P$ be a purple edge that is not an edge of the convex hull. We say that e is *color-blind flippable* (with respect to R and B) if e is adjacent to two different faces of (S, P), and it is adjacent to a red triangle in T_R and a blue triangle in T_B that form a convex quadrilateral. A purple edge $e \in P$ is called *flippable* (with respect to R and B) if it is flippable in the triangulation T_R or T_B, or it is color-blind flippable with respect to R and B. With this definition we can obtain a local characterization of maximal biplane graphs.

Theorem 6 *Let $G = (S, E)$ be a biplane graph, and let $E = R \cup B$ such that $T_R = (S, R)$ and $T_B = (S, B)$ are two triangulations of S. Then G is a maximal biplane graph if and only if it contains no edge $e \in R \cap B$ flippable with respect to R and B.*

Before proving this result we first establish a helpful lemma. Recall that a bounded face of the purple plane graph (S, P), $P = R \cap B$, is a weakly simple polygon $\mathcal{F}$, possibly with holes. A *chord* of $\mathcal{F}$ is an internal diagonal of $\mathcal{F}$ (connecting two vertices through the interior of $\mathcal{F}$). Denote by $R(\mathcal{F})$ (resp., $B(\mathcal{F})$) the set of chords of $\mathcal{F}$ in R (resp., in B). Note that $R(\mathcal{F})$ and $B(\mathcal{F})$ are two different triangulations of $\mathcal{F}$, and by definition, $R(\mathcal{F}) \cap B(\mathcal{F}) = \emptyset$. By exchanging the triangulations $R(\mathcal{F})$ and $B(\mathcal{F})$, we obtain a new decomposition $E = R' \cup B'$ into $R' = (R \setminus R(\mathcal{F})) \cup B(\mathcal{F})$ and $B' = (B \setminus B(\mathcal{F})) \cup R(\mathcal{F})$. We show that, essentially, those are the only possible changes that can be done to a decomposition of E into two triangulations.

Lemma 7 *Let $\mathcal{F}$ be a weakly simple polygon, and let $R(\mathcal{F})$ and $B(\mathcal{F})$ be two disjoint sets of chords, each of which forms a triangulation of the interior of $\mathcal{F}$. Then the intersection graph $G_{\mathcal{F}}$ of the open line segments $R(\mathcal{F}) \cup B(\mathcal{F})$ is connected.*

Proof. For an edge $e \in R(\mathcal{F}) \cup B(\mathcal{F})$, denote by $v(e)$ the corresponding node in the intersection graph $G_{\mathcal{F}}$. The chords $R(\mathcal{F})$ and $B(\mathcal{F})$ form two distinct triangulations of $\mathcal{F}$, which we call the red and the blue triangulation of $\mathcal{F}$. We prove that if $e, e' \in B(\mathcal{F})$ are edges of a triangle Δ in the blue triangulation, then $G_{\mathcal{F}}$ contains a path between the nodes $v(e)$ and $v(e')$. It follows that all blue chords in $B(\mathcal{F})$ must be in the same connected component of $G_{\mathcal{F}}$, since the dual graph of the blue triangulation is connected. Analogously, all red chords in $R(\mathcal{F})$ are in the same connected component of $G_{\mathcal{F}}$. Since every red edge crosses a blue edge (and vice versa), $G_{\mathcal{F}}$ is connected.

Let $e, e' \in B(\mathcal{F})$ be two blue edges adjacent to a triangle Δ in the blue triangulation of $\mathcal{F}$; see Figure 3. We show that there is a path in $G_{\mathcal{F}}$ between the corresponding nodes $v(e)$ and $v(e')$. There are two situations to consider. If a red chord $e_r \in R(\mathcal{F})$ intersects both e and e', then there is a path of length 2 from $v(e)$ to $v(e')$ in $G_{\mathcal{F}}$. Otherwise, no red edge in $R(\mathcal{F})$ intersects both e and e'. There must be a red edge $e_r \in R(\mathcal{F})$ that crosses e and a red edge $e'_r \in R(\mathcal{F})$ that crosses e'. Since e_r and e'_r do not cross each other, neither of them can be incident to any vertex of Δ. Hence both e_r and e'_r cross the third edge of Δ, which must be a blue chord in $B(\mathcal{F})$ (rather than an edge of $\mathcal{F}$). Therefore $v(e)$ and $v(e')$ are connected in $G_{\mathcal{F}}$ by a path of length 4. $\square$

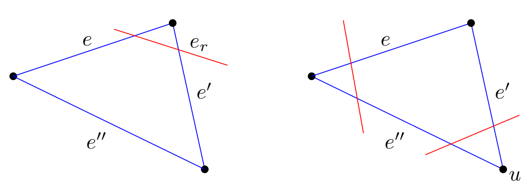

Figure 3: The blue edges $e, e' \in B(\mathcal{F})$ are adjacent to a common blue triangle Δ. Left: a red edge e_r crosses both e and e'. Right: no red edge crosses both e and e', but there is a red edge that crosses e and e'', and another red edge that crosses e' and e''.

By Lemma 7, the edges in $R(\mathcal{F}) \cup B(\mathcal{F})$ have a unique bipartition. Therefore, if any edge in $R(\mathcal{F})$ changes its color, then all edges in $R(\mathcal{F}) \cup B(\mathcal{F})$ must also change their colors. We are now ready to prove Theorem 6.

Proof. [Proof of Theorem 6.] Let $G = (S, E)$ be a biplane graph and let $E = R \cup B$ such that (S, R) and (S, B) are two triangulations of S.

Suppose that there is an edge $e \in E$ flippable with respect to R and B. We show that G cannot be a maximal biplane graph. If e is a flippable in R or B, then we can flip it in one triangulation and retain it in the other, thereby increasing the total number of edges by one. Let $e \in P = R \cap B$ be a color-blind flippable edge. Denote by $\mathcal{F}_i$ and $\mathcal{F}_j$ the faces of the purple graph (S, P) adjacent to e, and assume without loss of generality that e is adjacent to a red triangle in $R(\mathcal{F}_i)$ and a blue triangle in $B(\mathcal{F}_j)$ such that they form a convex quadrilateral. Then we can obtain a new decomposition $E = R' \cup B'$ with $R' = (R \setminus R(\mathcal{F}_i)) \cup B(\mathcal{F}_i)$ and $B' = (B \setminus B(\mathcal{F}_i)) \cup R(\mathcal{F}_i)$. By flipping e in the triangulation R', and retaining it in B', the total number of edges increases by one.

Suppose now that $G = (S, E)$ is not a maximal biplane graph. We show that there is a flippable purple edge in E (with respect to R and B). Since G is not maximal, G can be augmented to a larger biplane graph $G_{\max} = (S, E_{\max})$ such that $E \subset E_{\max}$. Let $(S, R_{\max})$, $(S, B_{\max})$ be two triangulations such that $E_{\max} = R_{\max} \cup B_{\max}$. (Possibly, $R \not\subseteq R_{\max}$ or $B \not\subseteq B_{\max}$.)

Denote the faces of the purple graph (S, P) by $\mathcal{F}_1, \ldots, \mathcal{F}_k$. Let R_i (resp. B_i) be the set of red (resp. blue) chords in face $\mathcal{F}_i$. Recall that $R = R_1 \cup \ldots \cup R_k \cup P$ and $B = B_1 \cup \ldots \cup B_k \cup P$, where P is the set of purple edges. By Lemma 7, each R_i and B_i must be completely contained in either $R_{\max}$ or $B_{\max}$. That is, we have $R_{\max} = C_1 \cup \ldots \cup C_k \cup R_{k+1}$ and $B_{\max} = \overline{C}_1 \cup \ldots \cup \overline{C}_k \cup B_{k+1}$, where $\{C_i, \overline{C}_i\} = \{R_i, B_i\}$, and R_{k+1} and B_{k+1} are two additional sets that complete $R_{\max}$ and $B_{\max}$ to a triangulation, respectively.

Consider now the triangulations $R' = C_1 \cup \ldots \cup C_k \cup P$, and $B' = \overline{C}_1 \cup \ldots \cup \overline{C}_k \cup P$. By construction, we have $E = R' \cup B'$, hence this is another decomposition of E into two triangulations. Note that R' and $R_{\max}$ share all edges in $C_1 \cup \ldots \cup C_k$.

Consider all triangulations on S that contain the edges $C_1 \cup \ldots \cup C_k$ (i.e., these edges are *constrained*), including R' and $R_{\max}$. It is known [23] that between any two constrained triangulations on the same point set, there is a sequence of edge flips (of unconstrained edges) that transform one into the other. In particular, there is a sequence of flips that transforms R' into $R_{\max}$, flipping only *unconstrained* edges. The first edge flipped in the sequence is in P, implying that P contains at least one flippable edge with respect to R'. With respect to the original decomposition $E = R \cup B$, this edge is flippable in the triangulation T_R or T_B, or color-blind flippable with respect to R and B, as required. ☐

Before presenting our algorithm, we show that the fact that an edge is flippable is mostly a local property. In the following we use the term *flippability* to refer to the condition of an edge as being flippable or color-blind flippable.

Lemma 8 *The flip of an edge e can only change the* flippability *of edges in the triangles that contain e.*

Proof. The fact that an edge e' is flippable in R or B only depends on the two triangles that are adjacent to e' in R and B (more precisely, on the up to four different combinations of adjacent triangles). Thus, flipping e cannot affect the flippability of e' if e' is not part of one of the four triangles containing e. Moreover, by flipping e we reduce the number of purple edges and purple faces by one. Thus, the color-blind flippability of e' cannot be affected either if e' is not in a triangle that also contains e. ☐

We now describe our algorithm to augment a given biplane graph $G = (S, E)$ to a maximal biplane graph. We begin with the biplanarity test (Lemma 5) that in $O(n \log n)$ time returns a decomposition $E = R \cup B$ such that (S, R) and (S, B) are plane graphs. In a preprocessing phase, the plane graphs (S, R) and (S, B) are triangulated. The edges are classified into red, blue, and purple; the faces of the purple graph (S, P) are computed.

Put all purple edges in a priority queue Q. Check for each $e \in Q$ whether it is flippable. If e is not flippable, it is simply discarded. Otherwise, we "flip" e, that is, we insert a flipped counterpart of e as a new edge, and update the face decomposition as follows. If e is flippable in T_R or T_B, then the new edge is part of one triangulation, and all other edges keep their original color; if e is color-blind flippable, then all chords in one of the adjacent faces of the purple graph change colors. After each flip, up to four other purple edges affected by the flip are reinserted into the priority queue Q (since they may become flippable). The algorithm terminates when the priority queue Q is empty. Theorem 6 guarantees that the resulting biplane graph is maximal.

Theorem 9 *Given a biplane graph $(S, E) \in \mathcal{G}_2(S)$ we can compute a maximal graph $(S, E_{\max}) \in \mathcal{G}_2(S)$ such that $E \subseteq E_{\max}$ in $O(n \log n)$ time.*

Proof. Correctness of the algorithm follows directly from Theorem 6 and Lemma 8. Thus, it remains to show that the algorithm runs in $O(n \log n)$ time. By Lemma 5, it takes $O(n \log n)$ time to produce the initial decomposition $E = R \cup B$. We can triangulate each of them in $O(n \log n)$ time. We assume the triangulations are represented in a data structure allowing constant-time navigation between edges and adjacent faces (such as a doubly connected edge list). Classification of the edges into red, blue, and purple, as well as creating the face-decomposition of the purple graph, can be done in $O(n \log n)$ time. For each purple edge, we store its two adjacent red triangles and two adjacent blue triangles. Hence, we can check whether a purple edge $e \in P$ is flippable (with respect to R and B) in constant time.

The second phase consists in checking all purple edges and trying to flip them. The algorithm maintains all purple edges in a priority queue. Note that when a purple edge is flipped, the two adjacent faces of the purple graph merge. We maintain the set of red and blue chords of the purple graph in a standard union-find data structure, so we can find which face a chord belongs to and merge two faces in $O(\log n)$ time. In this way, processing each purple edge takes $O(\log n)$ time. In addition to updating the face structure, we must check the flippability of up to four more purple edges each time an edge is added into G. We charge this extra cost to the added edge.

Since the number of edges in a biplane graph is bounded by $6n$, by Lemma 1, the number of edges we will check is also bounded by $O(n)$. That is, after an $O(n \log n)$-time preprocessing, our algorithm will check the flippability of $O(n)$ edges. Each purple edge can be processed in $O(\log n)$ time: in constant time if the edge is not flippable and in $O(\log n)$ time otherwise. Thus, we conclude that the algorithm runs in $O(n \log n)$ total time. ☐

Remark The algorithm in Theorem 9 augments a biplane graph (S, E) drawn on a point set S into a maximal biplane graph incrementally, adding edges one-by-one. If we need an arbitrary maximal biplane graph on S, then we can start with the empty graph $(S, \emptyset)$; if we would like to generate another maximal biplane graph, it suffices to execute the algorithm again where the initial graph consists of a single edge not present in the previously obtained maximal graph.

4 Connectivity of Maximal Biplane Graphs

In this section we consider the following question. What is the maximum possible connectivity of a graph in $\mathcal{G}_2(S)$ over all n-point sets S? In other words, this sections studies the problem of finding the value

$$\kappa_{2P}(n) = \max_{|S|=n} \max_{G \in \mathcal{G}_2(S)} \kappa(G).$$

If the points in S are in convex position, then every graph in $\mathcal{G}_2(S)$ is planar (by Lemma 1(i) in our companion paper [16]), and thus cannot be 6-connected. However, a biplane graph in $\mathcal{G}_2(S)$ may achieve higher connectivity for certain sets S. As noted above, Hutchinson et al. [21] proved that every biplane graph in $\mathcal{G}_2(S)$ has at most $6n - 18$ edges for $n \geq 8$. Therefore the sum of the vertex degrees is at most $12n - 36$, and there is always a vertex of degree at most 11 (for any $n \in \mathbb{N}$). Consequently, 11 is an upper bound for vertex-connectivity. In the following we show how to construct a biplane graph with minimum vertex degree 10, and then we modify this construction to obtain an 11-connected biplane graph. The construction combines elements of a construction by Huntchinson et al. [21] with fullerenes.

Huntchinson et al. [21] constructed a biplane graph with n vertices and $6n - 20$ edges (for sufficiently large values of n). The core of their construction is a set of k^2 points placed on a $k \times k$ section of the integer grid with coordinates (i, j) for $1 \leq i \leq k$ and $1 \leq j \leq k$. Essentially, a vertex (i, j) is connected to vertices $(i \pm 1, j)$, $(i, j \pm 1)$, $(i \pm 1, j + 1)$, $(i \pm 1, j - 1)$, $(i + 2, j + 1)$, $(i - 2, j - 1)$, $(i + 1, j - 2)$,

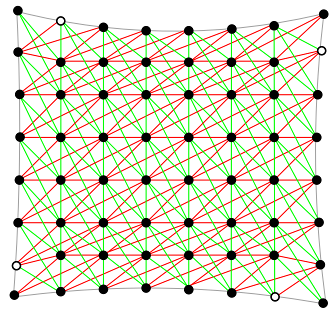

Figure 4: Scheme of the 10-connected biplane graph formed by two planar triangulations. Empty dots correspond to the vertices of degree 6 in the grid. Each one must have also degree 6 in the 5-connected graph attached to its side.

and $(i-1, j+2)$ whenever they exist (see Figure 4). We say that a vertex (i, j) is a *boundary vertex* if $i \in \{1, k\}$ or $j \in \{1, k\}$, and an *interior vertex* otherwise. The graph is the union of two lattice triangulations, as shown in Figure 4. Observe that all interior vertices have degree 10 or higher, but the boundary vertices have lower degree: the degree is 4 at the four corners, 6 at four neighbors of the corners and 7 at the other boundary vertices.

Fullerenes [5] are planar 3-regular graphs with exactly 12 pentagonal faces, and whose all other faces are hexagonal. It is known that a fullerene of $2k$ vertices can be constructed for $k = 10$ or $k \geq 12$. Moreover, there exist fullerenes in which pentagonal faces are sufficiently far from each other: for example, for $k \geq 36$ (and for $k = 30$), there is a fullerene with $2k$ vertices in which there are no two adjacent pentagonal faces [7]. The dual of such a fullerene is a planar graph with $k+2$ vertices and triangular faces, where all vertices have degree 6 except for twelve pairwise nonadjacent vertices of degree 5. Since every fullerene is cyclically 5-edge connected [10], its dual graph is 5-connected [5]. Thus the dual graph can be represented as a biplane graph drawn on points in convex position (c.f., Lemma 1(ii) in [16]).

We modify the $k \times k$ grid construction of Huntchinson et al. [21] as follows. Deform slightly the bounding box of the $k \times k$ integer grid in a way that each side becomes a reflex curve. Note that these curves need to be sufficiently flat to maintain the intersection pattern of the non-boundary edges (see Figure 4). Attach a 5-connected biplane graph to each side of the grid. Align the 5-connected biplane graphs along each side of the grid such that a vertex of degree 6 in the 5-connected biplane graph is identified with a vertex of degree 6 in the grid (marked with empty dot in Figure 4). Denote by $G_{10}(k)$ the resulting graph: it has k^2 vertices, and is clearly biplane. Moreover, all vertices of $G_{10}(k)$ have degree at least 10. In fact, each side of the grid only has thirteen vertices of degree 10 and there are four more interior vertices of degree 10; all the other vertices have degree 11 or higher. That is, for any $k > 0$, there is a constant number of vertices whose degree is exactly 10, whereas the others have higher degree.

Proposition 1 *For any $k \geq 40$ it holds that $\kappa(G_{10}(k)) = 10$.*

The proof of this claim is deferred to a forthcoming full version of this paper.

We highlight that $G_{10}(k)$ is not 11-connected because it has a constant number of vertices of degree 10 (i.e., the only way to disconnect $G_{10}(k)$ by removing only 10 vertices is by removing the neighbors of a vertex of low degree). By making some local flips around those vertices, the minimum vertex degree (and connectivity) can be increased to 11.

For this purpose we first characterize the points that have degree 10 in $G_{10}(k)$. Consider first those in the boundary of the grid and recall that we used a fullerene construction in each boundary side. In particular, all but twelve points will have six adjacencies with other vertices in the same boundary side (regardless of the value of k). The remaining twelve points will only have 5 adjacencies within the boundary side. Note that it is possible to choose where to place these vertices so that, for a sufficiently large k, the points with 5 adjacencies satisfy: (i) they are not adjacent in $G_{10}(k)$, (ii) they are sufficiently far apart (say, with at least 12 vertices in the grid in between any two of them), and (*iii*) they are also far from the grid corners. Note that property (*iii*) guarantees that these vertices also have 5 adjacencies towards the interior of the grid (thus their degree is 10).

The other situation in which a boundary vertex can have degree 10 is if it has degree 6 within the boundary, but only 4 adjacencies towards the interior of the grid. This only happens to the vertices at positions $(1, 2)$, $(2, k)$, $(k, k-1)$ and $(k-1, 1)$. Note that corner vertices have 6 adjacencies on each boundary side (and 2 towards the grid), thus their degree is 14. That is, regardless of the value of k, each boundary side will have $12 + 1 = 13$ vertices of degree 10.

Finally, we must consider points in the interior of the grid that have degree 10. Note that only those at locations $(2, 2)$, $(2, k-1)$, $(k-1, 2)$, and $(k-1, k-1)$ will have low degree. It is easy to see that, by the construction of Huntchinson et al. [21], any other vertex in the grid has at least 11 other neighbors (and in most cases, it will have 12 adjacencies). Thus, in total we have $(12 + 1) \times 4 + 4 = 56$ vertices of degree 10. Note that four couples of such vertices are located around the grid corners, while the others are spaced along the boundary of the grid.

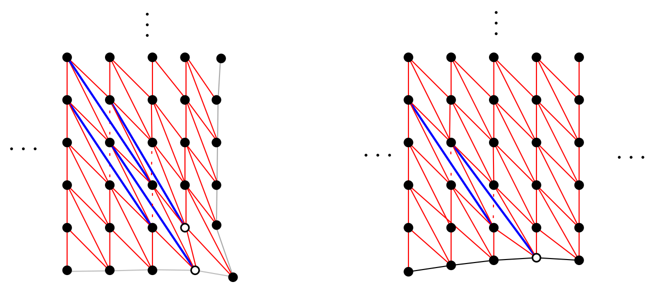

Figure 5: Four and two flips are sufficient to increase the minimum vertex degree around the corner (left) or along the boundary side (right), respectively. In both cases, blue (fat) edges must be added while the red (dashed) edges are removed. White dots represent the vertices that originally had degree 10.

In Figure 5 (left) we show the changes needed to increase the vertex degree for the lower right corner (for clarity purposes, only one of the triangulations is shown). The construction for the other corners is analogous, although the flips might happen in the other triangulation. The case in which the vertex of degree 10 is along the grid only needs 2 flips (see Fig 5 (right)). Observe that, these transformations are local, thus for a sufficiently large k, they can be done without one transformation affecting the others. Let $G_{11}(k)$ be the resulting biplane graph. By construction all vertices have degree 11 or higher.

We summarize the results of this Section with the following theorem.

Theorem 10 *There exist infinitely many* 11*-connected biplane graphs, and no biplane graph is 12-connected.*

The proof of this claim is also deferred to a forthcoming full version.

References

[1] M. Abellanas, A. García, F. Hurtado, J. Tejel, and J. Urrutia, Augmenting the connectivity of geometric graphs, *Comput. Geom. Theory Appl.* **40** (3) (2008), 220–230.

[2] M. Ajtai, V. Chvátal, M.M. Newborn, and E. Szemerédi, Crossing-free subgraphs, *Annals of Discrete Mathematics* **12** (1982), 9–12.

[3] M. Al-Jubeh, G. Barequet, M. Ishaque, D. L. Souvaine, C. D. Tóth, and A. Winslow, Constrained tri-connected planar straight line graphs, in *Thirty Essays on Geometric Graph Theory (J. Pach, ed.)*, Springer, 2013, pp. 49–70.

[4] M. Al-Jubeh, M. Ishaque, K. Rédei, D. L. Souvaine, C. D. Tóth, and P. Valtr, Augmenting the edge connectivity of planar straight line graphs to three, *Algorithmica* **61** (4) (2011), 971–999.

[5] D. Barnette, On generating planar graphs, *Discrete Mathematics* **7** (1974), 199–208.

[6] L. W. Beineke, Biplanar Graphs: A Survey, *Computers & Mathematics with Applications* **34** (11) (1997), 1–8.

[7] G. Brinkmann, J. Goedgebeur, B. D. McKay, The generation of fullerenes, *Journal of Chemical Information and Modeling* **52** (11) (2012), 2910–2918.

[8] G. Chartrand, L. Lesniak, Graphs and digraphs, Chapman and Hall/CRC, 2005.

[9] T. K. Dey, M. B. Dillencourt, S. K. Ghosh and J. M. Cahill, Triangulating with high connectivity, *Computational Geometry: Theory and Applications* **8** (1997), 39–56.

[10] T. Doslic, Cyclical edge-connectivity of fullerene graphs and $(k, 6)$-cages, *Journal of Mathematical Chemistry* **33** (2003), 103–112.

[11] P. Eades, S.-H. Hong, G. Liotta, and S.-H. Poon, Fáry's theorem for 1-planar graphs, in *Proc. 18th COCOON*, vol. 7434 of LNCS, Springer, 2012, pp. 335–346.

[12] D. Eppstein, Testing bipartiteness of geometric intersection graphs, *ACM Transactions on Algorithms* **5** (2) (2009), article 15.

[13] K.P. Eswaran and R.E. Tarjan, Augmentation problems, *SIAM Journal of Computing* **5** (1976), 653–665.

[14] I. Fáry, On straight-line representation of planar graphs, *Acta Scientiarum Mathematicarum* (Szeged) **11** (1948), 229–233.

[15] A. García, C. Huemer, F. Hurtado, J. Tejel, and P. Valtr, On triconnected and cubic plane graphs on given point sets, *Comput. Geom. Theory Appl.* **42** (9) (2009), 913–922.

[16] A. García, F. Hurtado, M. Korman, I. Matos, M. Saumell, R. I. Silveira, J. Tejel, and C. D. Tóth, Geometric Biplane Graphs II: Graph Augmentation, manuscript, 2013.

[17] A. García, C. Huemer, J. Tejel, and P. Valtr, On 4-connected geometric graphs, Presentation at XV Spanish Meeting on Computational Geometry, June 26-28, Seville, Spain, 2013.

[18] O. Giménez and M. Noy, Asymptotic enumeration and limit laws of planar graphs, *Journal of the AMS* **22** (2) (2009), 309–329.

[19] M. Hoffmann, A. Schulz, M. Sharir, A. Sheffer, C. D. Tóth, and E. Welzl, Counting plane graphs: flippability and its applications, in *Thirty Essays on Geometric Graph Theory (J. Pach, ed.)*, Springer, 2013, pp. 303–326.

[20] F. Hurtado and C. D. Tóth, Plane geometric graph augmentation: a generic perspective, in *Thirty Essays on Geometric Graph Theory (J. Pach, ed.)*, Springer, 2013, pp. 327–354.

[21] J. P. Hutchinson, T. C. Shermer and A. Vince, On representations of some thickness-two graphs, *Computational Geometry: Theory and Applications* **13** (1999), 161–171.

[22] V. P. Korzhik and B. Mohar, Minimal obstructions for 1-immersions and hardness of 1-planarity testing, in *Proc. 16th Graph Drawing*, vol. 5417 of LNCS, Springer, 2009, pp. 302–312.

[23] D. T. Lee and A. K. Lin, Generalized Delaunay triangulation for planar graphs, *Discrete and Computational Geometry* **1** (1986), 201–217.

[24] I. Rutter and A. Wolff, Augmenting the connectivity of planar and geometric graphs, *Journal of Graph Algorithms and Applications* **16** (2) (2012), 599–628.

[25] M. Sharir and A. Sheffer, Counting triangulations of planar point sets, *Electr. J. Comb.* **18** (1) (2011).

[26] C. D. Tóth, Connectivity augmentation in planar straight line graphs, *European J. of Combinatorics*, **33** (3) (2012), 408–425.