

Every Collinear Set in a Planar Graph Is Free

Vida Dujmović, Fabrizio Frati, Daniel Gonçalves, Pat Morin, Günter Rote

► **To cite this version:**

Vida Dujmović, Fabrizio Frati, Daniel Gonçalves, Pat Morin, Günter Rote. Every Collinear Set in a Planar Graph Is Free. *SODA: Symposium on Discrete Algorithms*, Jan 2019, San Diego, CA, United States. pp.1521-1538, 10.1137/1.9781611975482.92 . lirmm-02046476

HAL Id: lirmm-02046476

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-02046476>

Submitted on 8 Mar 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

EVERY COLLINEAR SET IN A PLANAR GRAPH IS FREE[‡]*Vida Dujmović,[‡] Fabrizio Frati,[‡] Daniel Gonçalves,[‡] Pat Morin,[‡] and Günter Rote[§]*

November 9, 2018

ABSTRACT. We show that if a planar graph G has a plane straight-line drawing in which a subset S of its vertices are collinear, then for any set of points, X , in the plane with $|X| = |S|$, there is a plane straight-line drawing of G in which the vertices in S are mapped to the points in X . This solves an open problem posed by Ravsky and Verbitsky in 2008. In their terminology, we show that every collinear set is free.

This result has applications in graph drawing, including untangling, column planarity, universal point subsets, and partial simultaneous drawings.

[‡]The work of VD and PM was partly funded by NSERC. The work of FF was partially supported by MIUR Project MODE under PRIN 20157EFM5C and by H2020-MSCA-RISE project 734922, CONNECT. The work of DG was partly funded by the ANR project GATO, under contract ANR-16-CE40-0009.

[‡]Department of Computer Science and Electrical Engineering, University of Ottawa

[‡]Dipartimento di Ingegneria, Università Roma Tre

[‡]LIRMM, Université de Montpellier, CNRS

[‡]School of Computer Science, Carleton University

[§]Institut für Informatik, Freie Universität Berlin

Contents

1	Introduction	1
1.1	Applications and Related Work	3
1.2	Proof Outline for Theorem 1	5
2	Preliminaries	5
2.1	Characterization of Collinear Sets	6
3	A-Graphs	6
3.1	Setting the Proportionality Constraints	11
3.2	Ordering constraints	12
3.3	Strong Ordering Constraints	14
3.4	Uniqueness of Solutions Satisfying Ordering Constraints	14
3.5	A Parametric Family of Linear Systems	15
3.6	Existence (and uniqueness) of solutions to $A^t \cdot s = b^t$	16
4	Triangulations	21
5	Open Problems	26

1 Introduction

A *straight-line drawing* of a graph G maps each vertex to a point in the plane and each edge to a line segment between its endpoints. A straight-line drawing is *plane* if no pair of edges cross except at a common endpoint. A set of vertices $S \subseteq V(G)$ in a planar graph G is a *free set* if for any set of points X in the plane with $|X| = |S|$, G has a plane straight-line drawing in which the vertices of S are mapped to the points in X . Free sets are useful tools in graph drawing and related areas and have been used to settle problems in untangling [5, 9, 13, 22], column planarity [9, 13], universal point subsets [9, 13], and partial simultaneous geometric drawings [13].

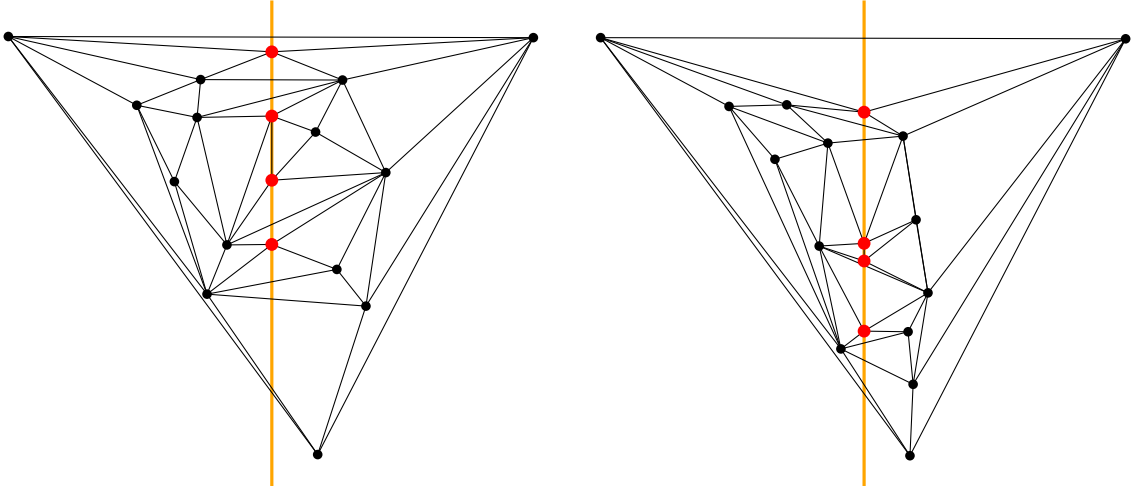


Figure 1: The 4 red vertices form a collinear set S . On the right, the graph is redrawn so that vertices of S lie at some other collinear locations.

A set of vertices $S \subseteq V(G)$ in a planar graph G is a *collinear set* if G has a plane straight-line drawing in which all vertices in S are mapped to a single line, see Figure 1. A collinear set S is a *free collinear set* if, for any collinear set of points in the plane X with $|X| = |S|$, G has a plane straight-line drawing in which the vertices of S are mapped to the points in X . Ravsky and Verbitsky [22] define $\bar{v}(G)$ and $\tilde{v}(G)$ as the respective sizes of the largest collinear set and largest free collinear set in G , and ask the following question:

How far or close are parameters $\tilde{v}(G)$ and $\bar{v}(G)$? It seems that *a priori* we even cannot exclude equality. To clarify this question, it would be helpful to (dis)prove that every collinear set in any straight-line drawing is free.

Here, we answer this question by proving that, for every planar graph G , $\tilde{v}(G) = \bar{v}(G)$, that is:

Theorem 1. *Every collinear set is a free collinear set.*

Let $v(G)$ denote the largest free set for a planar graph G . Clearly, we have $v(G) \leq \tilde{v}(G) \leq \bar{v}(G)$. Further, as discussed in detail below, it is well-known that $v(G) = \tilde{v}(G)$. However, prior to our work, the best known bound between $v(G)$, $\tilde{v}(G)$, and $\bar{v}(G)$ in the other

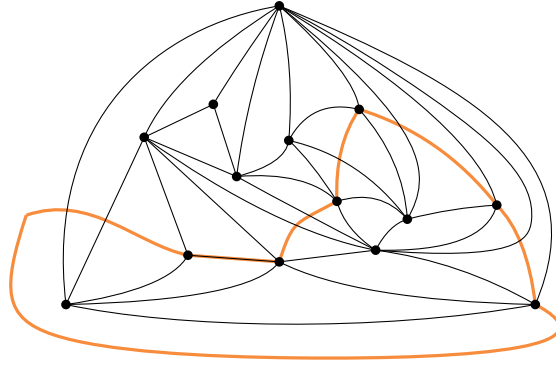


Figure 2: A proper good curve

direction was $v(G), \tilde{v}(G) \geq \sqrt{\bar{v}(G)}$, proved by Ravsky and Verbitsky [22]. Thanks to Theorem 1, we now know a stronger bound, in fact the ultimate $v(G) = \tilde{v}(G) = \bar{v}(G)$ relationship. This relationship was previously only known for planar 3-trees [9]. Theorem 1, in fact, implies a stronger result than $v(G) = \tilde{v}(G) = \bar{v}(G)$:

Corollary 1. *In a planar graph G , a set $S \subseteq V(G)$ is a free set if and only if it is a collinear set.*

That every free set is a collinear set is immediate. Theorem 1 then implies Corollary 1 since every free collinear set is also a free set. This fact, which implies that $v(G) = \tilde{v}(G)$, has been observed by several authors [5, 9, 13, 15]. To see it, let $X = \{(x_1, y_1), \dots, (x_{|S|}, y_{|S|})\}$ be the desired target locations on which S is supposed to be drawn. By rotation, we may assume that no two points have the same y -coordinate. Let $X_0 = \{(0, y_1), \dots, (0, y_{|S|})\}$. By the definition of free collinear set, G has a plane straight-line drawing Γ_0 in which S maps to X_0 . Since the set of plane straight-line drawings of G is an open set, we can arbitrarily perturb the vertices in some small neighborhood, resulting in some plane straight-line drawing Γ_ϵ in which S maps to $X_\epsilon = \{(\epsilon x_1, y_1), \dots, (\epsilon x_{|S|}, y_{|S|})\}$, for some $\epsilon > 0$. Dividing all the x -coordinates of Γ_ϵ by ϵ then yields a plane straight-line drawing in which S maps to X .

Thus, Theorem 1 is our main result and this paper is dedicated to proving it. The following characterization of collinear sets by Da Lozzo, Dujmović, Frati, Mchedlidze, and Roselli [9] is helpful in that goal.

Definition 1. Given a drawing G , a Jordan curve C is a *proper good curve* if it contains a point in the outer face of G and the intersection between C and each edge e of G is either empty, a single point, or the entire edge e . See Figure 2 for an example. (This is a conjunction of the two definitions of *proper* and of *good curves* from [9].)

Theorem 2. [9] *A set S of vertices of a graph G is a collinear set if and only if there is a plane drawing of G and a proper good curve C that contains every vertex in S .*

The *if* part of this theorem says, in other words, that the curve C and the edges of G can be simultaneously straightened (after cutting C open at some point in the outer face) while keeping the vertices of S on C . Theorem 2 is helpful because it reduces the problem of finding large collinear sets in a graph G to a topological game in which one only needs

to find a curve that contains many vertices of G . Da Lozzo et al. [9] used Theorem 2 to give tight lower bounds on the sizes of collinear sets in planar graphs of treewidth at most 3 and triconnected cubic planar graphs. Despite the conceptual simplification provided by Theorem 2, the identification of collinear sets is highly non-trivial: Mchedlidze, Radermacher, and Rutter [19] showed that it is NP-hard to determine if a given set of vertices in a planar graph is a collinear set. Nevertheless, Theorem 2 is a useful tool for finding large collinear sets. In combination with Corollary 1, it gives a characterization of free sets:

Corollary 2. *A set S of vertices of a graph G is a free set if and only if there is a plane drawing of G and a proper good curve C that contains every vertex of S .*

This is a useful tool for finding free sets, which have a wide variety of applications, as outlined in the next section.

1.1 Applications and Related Work

The applicability of Corollary 1 comes from the fact that a number of graph drawing applications require (large) free sets, whereas finding large collinear sets is an easier task. Indeed there are planar graphs for which large collinear sets were known to exist, however large free sets were not. Those include triconnected cubic planar graphs and planar graphs of treewidth at least k . We now review applications of our result.

Untangling. Given a straight-line drawing of a planar graph G , possibly with crossings, to *untangle* it means to assign new locations to some of the vertices of G so that the resulting straight-line drawing of G becomes noncrossing. The goal is to do so while *keeping fixed* the location of as many vertices as possible.

In 1998, Watanabe asked if every polygon can be untangled while keeping at least εn vertices fixed, for some $\varepsilon > 0$. Pach and Tardos [21] answered that question in the negative by providing an $\mathcal{O}((n \log n)^{2/3})$ upper bound on the number of fixed vertices. This has almost been matched by an $\Omega(n^{2/3})$ lower bound by Cibulka [8]. Several papers have studied the untangling problem [5, 6, 8, 15, 17, 21, 22]. Asymptotically tight bounds are known for paths [8], trees [15], outerplanar graphs [15], and planar graphs of treewidth two and three [9, 22].

For general planar graphs there is still a large gap. Namely, it is known that every planar graph can be untangled while keeping $\Omega(n^{0.25})$ vertices fixed [5] (this answered a question by Pach and Tardos [21]) and that there are planar graphs that cannot be untangled while keeping $\Omega(n^{0.4948})$ vertices fixed [6]. Theorem 1 can help close this gap, whenever a good bound on collinear sets is known. Namely, Bose et al. [5] (implicitly) and Ravsky and Verbitsky [22] (explicitly) proved that every straight-line drawing of a planar graph G can be untangled while keeping $\Omega(\sqrt{|S|})$ vertices fixed, where S is a free set of G . Together with Corollary 1 this implies that, for untangling, it is enough to find large collinear sets.

Theorem 3. *Let S be a collinear set of a planar graph G . Every straight-line drawing of G can be untangled while keeping $\Omega(\sqrt{|S|})$ vertices fixed.*

Da Lozzo, Dujmović, Frati, Mchedlidze, and Roselli [9] proved that every triconnected

cubic planar graph has a collinear set of size $\Omega(n)$. Then Theorem 3 implies the following new result, for which $\Omega(n^{0.25})$ was a previously best known untangling bound.

Corollary 3. *Every straight-line drawing of any n -vertex triconnected cubic planar graph can be untangled while keeping $\Omega(\sqrt{n})$ vertices fixed.*

Corollary 3 is almost tight due to the $\mathcal{O}(\sqrt{n \log^3 n})$ upper bound for triconnected cubic planar graphs of diameter $\mathcal{O}(\log n)$ [8]. Corollary 3 cannot be extended to all bounded-degree planar graphs, see [13, 20] for reasons why. Da Lozzo et al. [9] also proved that planar graphs of treewidth at least k have $\Omega(k^2)$ -size collinear sets. Together with Theorem 3, this implies that every straight-line drawing of an n -vertex planar graph of treewidth at least k can be untangled while keeping $\Omega(k)$ vertices fixed. This gives, for example, a tight $\Theta(\sqrt{n})$ untangling bound for planar graphs of treewidth $\Theta(\sqrt{n})$.

Universal Point Subsets. Closing the gap between $\Omega(n)$ and $\mathcal{O}(n^2)$ on the size of the smallest *universal point set* (a set of points on which every n -vertex planar graph can be drawn with straight edges by using n of these points as locations for the vertices) is a major, extensively studied, and difficult graph drawing problem, open since 1988 [2, 10, 18].

The interest universal point sets motivated the following notion introduced by Angelini et al. [1]. A *universal point subset* for a set \mathcal{G} of n -vertex planar graphs is a set P of $k \leq n$ points in the plane such that, for every $G \in \mathcal{G}$, there is a plane straight-line drawing of G in which k vertices of G are mapped to the k points in P . Every set of n points in general position is a universal point subset for n -vertex outerplanar graphs [4, 7, 16]; every set of $\lceil \frac{n-3}{8} \rceil$ points in the plane is a universal point subset for the n -vertex planar graphs of treewidth at most three [9]; and, every set of $\sqrt{\frac{n}{2}}$ points in the plane is a universal point subset for the n -vertex planar graphs [13].

Dujmović [13] proved that every set of $v(G)$ points in the plane is a universal point subset for a planar graph G . Together with Corollary 1 this implies that, in order to find large universal point subsets, it is enough to look for large collinear sets.

Theorem 4. *Let S be a collinear set for a graph G . Then every set of $|S|$ points in the plane is a universal point subset for G .*

As was the case with untangling, Theorem 4 implies new results for universal point subsets of triconnected cubic planar graphs and treewidth- k planar graphs. In particular, Theorem 4 and the fact that every triconnected cubic planar graph has a collinear set of size $\lceil \frac{n}{4} \rceil$ [9] imply the following asymptotically tight result. The previously best known bound was $\Omega(\sqrt{n})$ [13].

Corollary 4. *Every set of $\lceil \frac{n}{4} \rceil$ points in the plane is a universal point subset for every n -vertex triconnected cubic planar graph.*

Similarly, Theorem 4 and the fact that planar graphs of treewidth at least k have collinear sets of size ck^2 , for some constant c [9], imply that every set of ck^2 points in the plane is a universal point subset for such graphs. This gives, for example, an asymptotically tight $\Theta(n)$ result on the size of the largest universal point subset for planar graphs of treewidth $\Theta(\sqrt{n})$.

For similar applications of Theorem 1 and Corollary 1, such as *column planarity* [3, 9, 13] and *partial simultaneous geometric embeddings with and without mappings* [3, 12, 13] see a survey by Dujmović [13].

1.2 Proof Outline for Theorem 1

We assume w.l.o.g. that G is a plane straight-line drawing in which the collinear set $S \subseteq V(G)$ lies on the y -axis $Y = \{(0, y) : y \in \mathbb{R}\}$. Let $L = \{(x, y) \in \mathbb{R}^2 : x < 0\}$ and $R = \{(x, y) \in \mathbb{R}^2 : x > 0\}$ denote the open halfplanes to the left and right of Y . We consider the points on the y -axis Y as being ordered, with $(0, a)$ before $(0, b)$ if $a < b$. We assume, furthermore, that we are given $|S|$ distinct y -coordinates, and the goal is to find another plane straight-line drawing of G in which the vertices in S are positioned on Y with the given y -coordinates.

The difficulty comes from edges of G that cross Y . These edges must cross Y in prescribed intervals between the prescribed locations of vertices in S , and these intervals may be arbitrarily small. An extreme version of this subproblem is the one in which G is a drawing where every edge intersects Y in exactly one point (possibly an endpoint) and the location of each crossing point is prescribed. The most difficult instances occur when G is edge-maximal.

In Section 3 we describe these edge-maximal graphs, which we call A-graphs. A-graphs are a generalization of quadrangulations, in which every face is either a quadrangle whose every edge intersects Y or a triangle with one vertex in each of L , Y , and R . Theorem 6 shows that it is possible to find a plane straight-line drawing of any A-graph where the intersections of the drawing with Y occur at prescribed locations. For this purpose, we set up a system of linear equations and show that it has a unique solution. This proof involves linear algebra and continuity arguments.

In Section 4 we prove that every collinear set is free. The technical statement of this result, Theorem 7, shows a somewhat stronger result for triangulations that makes it possible not only to prescribe the locations of vertices on Y but also to nearly prescribe the points at which edges of the triangulation cross Y . This proof uses combinatorial reductions that are applied to a triangulation T that either reduce its size or increase the number of edges that cross Y . When none of these reductions is applicable to T , removing the edges of T that do not cross Y creates an A-graph, G , on which we can apply Theorem 6.

Section 2 begins our discussion with definitions that we use throughout.

2 Preliminaries

We recall some standard definitions.

A *curve* C is a continuous function from $[0, 1]$ to \mathbb{R}^2 . The points $C(0)$ and $C(1)$ are the *endpoints* of C . A curve C is *simple* if $C(s) \neq C(t)$ for $s \neq t$ except possibly for $s = 0$ and $t = 1$; it is *closed* if $C(0) = C(1)$. A *Jordan curve* $C : [0, 1] \rightarrow \mathbb{R}^2$ is a simple closed curve. We will often not distinguish between a curve C and its image $\{C(t) : 0 \leq t \leq 1\}$. The *open curve* is the set $\{C(t) : 0 < t < 1\}$. A point $x \in \mathbb{R}^2$ lies *on* C if $x \in C$. For any Jordan curve C , $\mathbb{R}^2 \setminus C$ has two connected components: One of these, C^- , is finite (the *interior* of C) and the other, C^+ , is infinite (the *exterior* of C).

All graphs considered in this paper are finite and simple. We use $V(G)$ and $E(G)$ to denote the vertex set and edge set of G , respectively. We use xy to denote the edge between the vertices x and y .

A *drawing* of a graph G consists of G together with a one-to-one mapping $\varphi: V(G) \rightarrow \mathbb{R}^2$ and a mapping ρ from $E(G)$ to curves in \mathbb{R}^2 such that, for each $xy \in E(G)$, $\rho(xy)$ has endpoints $\varphi(x)$ and $\varphi(y)$. We will not distinguish between a drawing G and the underlying graph G , and we will never explicitly refer to φ and ρ . In particular, we will sometimes have a drawing G and we will speak about constructing a different drawing of G , without danger of confusion. A drawing is *straight-line* if each edge is a straight-line segment. A drawing is *plane* if each edge is a simple curve, and no two edges intersect, except possibly at common endpoints. A *Fáry drawing* is a plane straight-line drawing. A *plane straight-line graph* is a planar graph G along with an associated Fáry drawing of G .

By default, an edge curve includes its endpoints, otherwise we refer to it as an *open edge*. The *faces* of a plane drawing G are the maximal connected subsets of $\mathbb{R}^2 \setminus \bigcup_{xy \in E(G)} xy$. One of these faces, the *outer face*, is unbounded; the other faces are called *inner* or *bounded faces*. A *boundary vertex* is incident to the outer face, other vertices are called *interior vertices*. If C is a cycle in a plane drawing, then there is a Jordan curve whose image is the union of edges in C . In this case, the interior and exterior of C refer to the interior and exterior of the corresponding Jordan curve.

A *triangulation* (a *quadrangulation*) is a plane drawing, not necessarily with straight edges, in which each face is bounded by a 3-cycle (respectively, a 4-cycle).

A *separating triangle* of a graph G is a cycle of length 3 whose removal disconnects G .

The *contraction* of an edge xy in a graph G identifies x and y into a new vertex v . Formally, we obtain a new graph G' with $V(G') = V(G) \cup \{v\} \setminus \{x, y\}$ and $E(G') = \{ab \in E(G) : \{a, b\} \cap \{x, y\} = \emptyset\} \cup \{va : xa \in E(G) \text{ or } ya \in E(G)\}$. If G is a triangulation and we contract the edge $xy \in E(G)$, then the resulting graph G' is also a triangulation provided that xy is not part of a separating triangle. Any plane drawing of G leads naturally to a plane drawing of G' .

2.1 Characterization of Collinear Sets

We will make use of the following strengthening of Theorem 2 which follows from the proof in [9]:

Theorem 5. *For any planar graph G , the following two statements are equivalent:*

1. *There is a plane drawing of G and a proper good curve $C: [0, 1] \rightarrow \mathbb{R}^2$ such that the sequence of edges and vertices intersected by C is r_1, \dots, r_k .*
2. *There is a Fáry drawing of G in which the sequence of edges and vertices intersected by the y -axis Y is r_1, \dots, r_k .*

3 A-Graphs

In this section, we study a special class of graphs that are closely related to quadrangulations in which every edge crosses Y . (See Figure 3 for an example.)

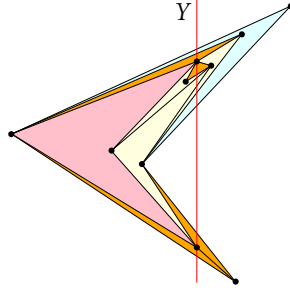


Figure 3: An A-graph with 2 vertices on Y .

Definition 2. An *A-graph*, G , is a plane straight-line graph with $n \geq 3$ vertices that has the following properties:

1. Every edge of G intersects Y in exactly one point, possibly an endpoint.
2. Every face of G , including the outer face, is a quadrilateral or a triangle (not containing any disconnected components inside).
3. Every quadrilateral face of G is non-convex.
4. Every triangular face contains one vertex on Y , one in L , and one in R .
5. Every vertex v on Y is incident to precisely two triangular faces, one “above v ”, which contains the line segment between v and $v + (0, \epsilon)$ for some $\epsilon > 0$, and one “below v ”, containing the line segment between v and $v - (0, \epsilon)$ for some $\epsilon > 0$.

In the special case where G has no vertices in Y , the graph G is a quadrangulation in which every edge crosses Y . Further, Property 5 applies even if v is on the outer face of G (in which case it implies that the outer face of G must be a triangle). Some additional properties of G follow from Properties 1–5.

6. G is connected.
7. Every vertex of G has degree at least 2.
8. If $n \geq 4$, then every vertex in Y has degree at least 3.

Property 6 follows directly from Property 2. Property 7 follows from the fact that every vertex is incident to at least one face and every face is a simple cycle. Property 8 follows from the fact that every vertex on Y is incident to at least two triangular faces, which involve at least 4 vertices, unless $n = 3$. (Property 3 and 4 are actually redundant—Property 3 follows from Properties 1 and 5; Property 4 follows from Property 1.)

In the following theorem, we will show that every A-graph G has a Fáry drawing with prescribed intersections with Y and a prescribed outer face.

Theorem 6.

- Let G be an A-graph.
- Let e_1, \dots, e_m be the sequence of edges in G , in the order they are intersected by Y . Ties between edges having a common endpoint on Y are broken arbitrarily, except that e_1 and e_m are always edges on the outer face.
- Let $y_1 \leq \dots \leq y_m$ be any sequence of numbers where, for each $i \in \{1, \dots, m-1\}$, $y_i = y_{i+1}$ if and only if e_i and e_{i+1} have a common endpoint in Y .

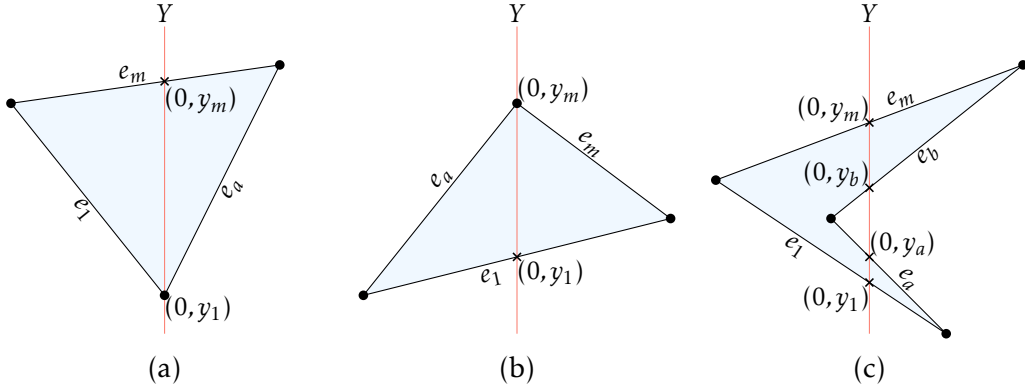


Figure 4: The three possibilities for the outer face in Theorem 6.

Then G has a Fáry drawing in which the intersection between e_i and Y is the single point $(0, y_i)$, for each $i \in \{1, \dots, m\}$.

Moreover, the shape Δ of the outer face can be prescribed, subject only to the constraint that Δ has to be consistent with the graph and the data y_1, \dots, y_m . Specifically, we have three possibilities, which are illustrated in Figure 4.

- a) If the outer face of G is a triangle containing the lowest vertex on Y , then Δ must be a triangle with a vertex at $(0, y_1)$, and the opposite edge crosses Y at $(0, y_m)$.
- b) Symmetrically, if the outer face of G is a triangle containing the highest vertex on Y , then Δ must be a triangle with a vertex at $(0, y_m)$, and the opposite edge crosses Y at $(0, y_1)$.
- c) Otherwise, the outer face of G is a quadrilateral. Let e_1 , e_a , e_b , and e_m be the edges of the outer face. Then Δ has to be a quadrilateral whose edges cross Y at $(0, y_1)$, $(0, y_a)$, $(0, y_b)$, and $(0, y_m)$.

It would have been more natural to represent the intersection of Y with G as a mixed sequence of vertices and edges. However, to simplify the statement of the theorem and its proof, we have chosen to specify the desired drawing by a number y_i for every edge, subject to equality constraints. The convention in Condition 2 about e_1 and e_m being boundary edges is introduced only for notational convenience.

The rest of this section is devoted to proving Theorem 6. We are going to prove Theorem 6 in its strongest form, in which the outer face Δ is prescribed. We begin by making some simplifying assumptions, all without loss of generality. First, we assume w.l.o.g. that Δ and all vertices of G are contained in the strip $[-1, 1] \times (-\infty, +\infty)$. This can be achieved by a uniform scaling. Second, if the outer face of G is a quadrilateral, we assume w.l.o.g. that the common vertex of e_1 and e_m in the given drawing of G is in L , as in Figures 3 and 4c, and the vertex of desired output shape Δ incident to e_1 and e_m is also in L ; this can be achieved by a reflection of G or Δ with respect to Y .

If $m = 3$ or $m = 4$, then G is a 3- or a 4-cycle, respectively, hence it suffices to draw it as Δ . Therefore we assume, from now on, that $m \geq 5$.

We will describe the desired Fáry drawing by assigning a slope s_i to each edge $e_i \in$

$E(G)$. Since there can be no vertical edges, each slope s_i is well-defined. We have $m = |E(G)|$ slope variables, s_1, \dots, s_m . We can see that these variables determine the drawing: Since every edge e_i contains the point $(0, y_i)$, the slope s_i fixes the line through e_i . Since every vertex v not on Y is incident to at least two edges that contain distinct points on Y , the location of v is fixed by any two of v 's incident edges. (The location of each vertex on Y is fixed by definition.) Our strategy is to construct a system of m linear equations in the m variables s_1, \dots, s_m , and to show that this system is feasible and that its solution gives the desired Fáry drawing of G .

A necessary condition for the slopes to determine a Fáry drawing of G is that the edges with a common vertex should be concurrent. Let v be a vertex not on Y , and let e_i, e_j, e_k be three edges incident to v . The fact that the supporting lines of e_i, e_j , and e_k meet at a common point (the location of v) is expressed by the following *concurrency constraint* in terms of the slopes s_i, s_j, s_k :

$$\begin{vmatrix} 1 & 1 & 1 \\ s_i & s_j & s_k \\ y_i & y_j & y_k \end{vmatrix} = (y_j - y_k)s_i + (y_k - y_i)s_j + (y_i - y_j)s_k = 0 \quad (1)$$

Since y_1, \dots, y_m are given, this is a linear equation in s_1, \dots, s_m . Writing this equation for all triplets of edges incident to a common vertex v will include many redundant equations. Indeed, if v has degree d_v , it suffices to take $d_v - 2$ equations: For each vertex $v \in V(G)$, we choose two fixed incident edges e_i and e_j and run e_k through the remaining $d_v - 2$ edges, specifying that e_k should go through the common vertex of e_i and e_j .

Whenever convenient, we will use edges of G as indices so that, if $e = e_i$ is an edge of G , then $s_e = s_i$ and $y_e = y_i$. Further, if e is a line segment that intersects Y in a point, we will use y_e to denote the y -coordinate of the intersection of e and Y and s_e to denote the slope of e .

We now introduce additional equations for the edges that emanate from a vertex on Y ; refer to Figure 5. Suppose that a vertex $v \in Y$ is incident to edges $a_1, \dots, a_k \in L \cup Y$ and $b_1, \dots, b_\ell \in Y \cup R$, ordered from bottom to top as in Figure 6.

From Property 4 of A-graphs we have $k, \ell \geq 1$ and in addition $k + \ell \geq 3$ by Property 8. Let us first look at the slopes on the right side. We want these slopes to be increasing: $s_{b_1} < s_{b_2} < \dots < s_{b_\ell}$. We stipulate a stronger condition: We require that the slopes $s_{b_2}, \dots, s_{b_{\ell-1}}$ partition the interval $[s_{b_1}, s_{b_\ell}]$ in fixed proportions. In other words:

$$s_{b_i} = s_{b_1} + \lambda_i(s_{b_\ell} - s_{b_1}), \quad (2)$$

for some fixed sequence $0 < \lambda_2 < \dots < \lambda_{\ell-1} < 1$.

For example, we might set $\lambda_i := (i - 1)/(\ell - 1)$. This gives $\ell - 2$ equations, for $\ell \geq 2$. Similarly, we get $k - 2$ equations for the slopes s_{a_1}, \dots, s_{a_k} of the edges on the left side, for $k \geq 2$. In addition, for $k \geq 2$ and $\ell \geq 2$, we require that the *range* of slopes on the two sides are in a fixed proportion:

$$s_{a_1} - s_{a_k} = \mu(s_{b_\ell} - s_{b_1}), \quad (3)$$

for some fixed value $\mu > 0$.

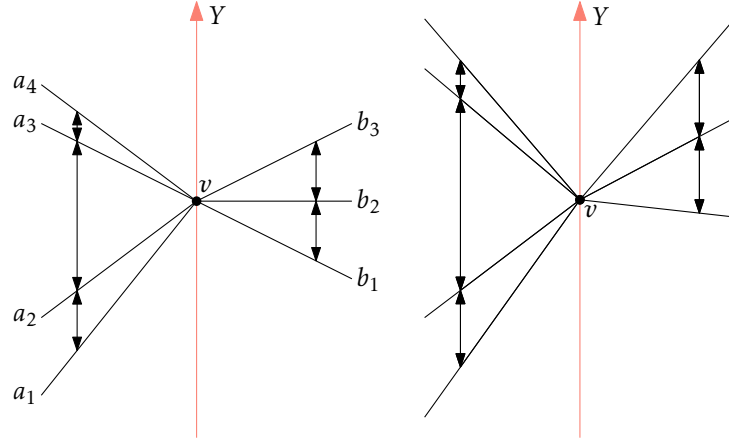


Figure 5: The proportionality constraints on slopes of edges incident to a vertex $v \in Y$.

We call the equations (2–3) the *proportionality constraints*. There are $(k + \ell) - 3$ such equations for the $k + \ell$ slopes, hence we have three degrees of freedom for the slopes incident to a vertex. Figure 5 illustrates these degrees of freedom: Namely, we can shear the edges on the right side vertically, adding the same constant to all slopes. We can independently shear all edges on the left side. In addition, we can vertically scale all lines jointly (both to the left and to the right), multiplying all slopes by the same constant factor. If this factor is negative, we would reverse the order of the slopes, simultaneously on the left and on the right. We will later see that this undesirable possibility is prevented in conjunction with other constraints that we are going to impose. We can already observe that any two slopes on one side determine all remaining slopes on that side. Moreover, the range of slopes on the other side ($s_{a_1} - s_{a_k}$ or $s_{b_\ell} - s_{b_1}$) is also determined. The notations λ_i and μ are here used in a local sense; for a different vertex v , we may choose different constants.

Lemma 1. *The total number of equations (1), (2), and (3) is $m - 4$.*

Proof. Let $n = |V|$ and let n_0 be the number of vertices on Y . Assume that G has f_3 triangular and f_4 quadrangular faces.

We have two triangles for every vertex on Y (Properties 4 and 5 of A-graphs):

$$f_3 = 2n_0 \quad (4)$$

Euler's formula gives

$$n + f_3 + f_4 = m + 2. \quad (5)$$

Double-counting of edge-face incidences leads to the relation

$$3f_3 + 4f_4 = 2m. \quad (6)$$

Denoting the degree of a vertex v by d_v , we have $d_v - 3$ equations for each of the n_0 vertices v on Y . For each of the $n - n_0$ vertices v not on Y , we have $d_v - 2$ equations. The total

number of equations is therefore

$$P = \sum_{v \in V \cap Y} (d_v - 3) + \sum_{v \in V \cap (L \cup R)} (d_v - 2) = \sum_{v \in V} (d_v - 2) - n_0 = 2m - 2n - n_0.$$

Using (4–6), this can be simplified to

$$\begin{aligned} P &= 2m - 2n - n_0 = 2m - 2n - (2f_3 + 2f_4) + (2f_3 + 2f_4) - n_0 \\ &= 2m - 2(n + f_3 + f_4) + \frac{1}{2}(4f_3 + 4f_4 - f_3) \\ &= 2m - 2(m + 2) + \frac{1}{2}(2m) = m - 4. \end{aligned} \quad \square$$

To achieve the desired number m of equations, we add four *boundary equations*. If the outer face is a quadrilateral, the desired slopes of the boundary edges already give us four equations: We set the slopes $s_1, s_a, s_b,$ and s_m of the boundary edges $e_1, e_a, e_b,$ and e_m to the fixed values of the slopes of the edges of Δ .

If the outer face is a triangle, the shape Δ gives us only three constraints for the slopes of the three edges e_1, e_a, e_m . If the triangle is $\alpha\beta\gamma$ with $\gamma \in Y$, we arbitrarily pick another (non-boundary) edge e_b incident to γ and set its slope s_b to an appropriate fixed value; this value has to be either larger or smaller than each of $s_1, s_a,$ and s_m depending on whether γ is the topmost or the bottommost point on Y and whether e_b lies in L or R . Together with the proportionality constraints, this effectively pins *all* slopes incident to γ to fixed values.

In both cases, we get 4 equations of the form

$$s_i = h_i, \tag{7}$$

where $i \in \{1, a, b, m\}$.

Altogether, we now have a system of m linear equations in the m unknowns $s = (s_1, \dots, s_m)$, which we can write compactly as $A \cdot s = b$, with a square matrix A whose entries come from (1–3) and (7). Only four entries of the right-hand side vector b are non-zero, due to the four boundary equations. We will show that $A \cdot s = b$ has a unique solution and that this solution gives a Fáry drawing of G .

3.1 Setting the Proportionality Constraints

Our plan is to construct the desired drawing by a continuous morph, starting from the given drawing of G . Since the proportionality constraints are not part of the output specification but were artificially added to achieve the right number of equations, we can make our life easy by just setting their coefficients so that they are satisfied by the initial drawing. Specifically, the statement of Theorem 6 assumes that G is a Fáry drawing. In this drawing, every edge e has a slope s'_e . We use these slopes to set the coefficients in the proportionality constraints. Consider a vertex $v \in Y$, incident to edges a_1, \dots, a_k and b_1, \dots, b_ℓ as described above. In the notation used in (2), we set

$$\lambda_i = (s'_{b_i} - s'_{b_1}) / (s'_{b_\ell} - s'_{b_1}).$$

The coefficients for the edges a_1, \dots, a_k on the left side are set similarly. If $k \geq 2$ and $l \geq 2$, we set

$$\mu = (s'_{a_1} - s'_{a_k}) / (s'_{b_\ell} - s'_{b_1})$$

in (3). This ensures that the initial slopes s'_1, \dots, s'_m satisfy the proportionality constraints.

3.2 Ordering constraints

We define a relation $<$ on the edges of G , where $e_1 < e_2$ if and only if

- $y_{e_1} < y_{e_2}$ and e_1 and e_2 have a common endpoint $v \in L$; or
- $y_{e_1} > y_{e_2}$ and e_1 and e_2 have a common endpoint $v \in R$.

We say that a vector $s = (s_1, \dots, s_m)$ satisfies the ordering constraints if $s_{e_1} < s_{e_2}$ for every pair $e_1, e_2 \in E(G)$ such that $e_1 < e_2$. This definition captures the condition that vertices of G in L (respectively, R) should be drawn so that they remain in L (respectively, R), as in the following.

Observation 1. *If a solution s to $A \cdot s = b$ satisfies the ordering constraints, then every vertex that is in L (in R) in G is also in L (respectively in R) in the drawing corresponding to s .*

Proof. Consider any vertex v that is in L in G and that is incident to (at least) two edges e_1 and e_2 with $y_{e_1} < y_{e_2}$, and hence $e_1 < e_2$. Since s satisfies the ordering constraints we have $s_{e_1} < s_{e_2}$, hence the lines with slopes s_{e_1} and s_{e_2} through $(0, y_{e_1})$ and $(0, y_{e_2})$, respectively, meet in L . The argument for the vertices in R is analogous. \square

By construction, the slopes $s'_{e_1}, \dots, s'_{e_m}$ of edges in G satisfy the ordering constraints, so the relation $<$ is acyclic.

Lemma 2. *Any solution s to $A \cdot s = b$ satisfying the ordering constraints yields a Fáry drawing of G with Δ as the outer face.*

Proof. If G is a plane drawing of a 2-connected graph, then another straight-line drawing G' of the same graph G is a Fáry drawing provided that two conditions are met: (i) For every vertex v , the clockwise order of the edges around v in G' is the same as in G ; and (ii) in the drawing G' , every face cycle of G is drawn without crossings (Devillers, Liotta, Preparata, and Tamassia [11, Lemma 16]).

In our case, G' is a straight-line drawing of G given by a solution to $A \cdot s = b$ that satisfies the ordering constraints.

First we show that G' satisfies condition (i). More specifically, we establish the following stronger property for every vertex v .

- (*) The edges going to the right from v are the same in G and G' , and their slopes have the same order in G and G' .

The same properties hold for the edges to the left.

We distinguish the following cases:

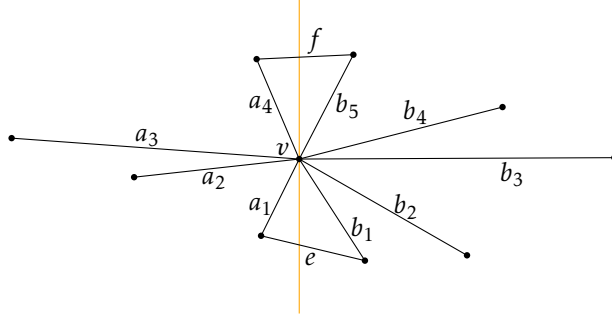


Figure 6: The ordering of the edges incident to a vertex v on Y .

1. $v \notin Y$. Since s satisfies the ordering constraints, by Observation 1 we know that v is on the same side (L or R) in G and in G' . All incident edges go to one side. This, together with the fact that the orders in which the edges incident to v intersect Y in G and G' agree implies that the slope orders of the edges around v in G and G' agree.
2. $v \in Y$, with incident edges $a_1, \dots, a_k \in L \cup Y$ and $b_1, \dots, b_\ell \in Y \cup R$ as in Figure 6. Again, Observation 1 ensures that these edges remain on the same side in G' .
 - (a) If v is a boundary vertex, then the boundary equations fix the slopes of the two incident boundary edges, a_1, b_1 or a_k, b_l , plus a third edge. As we already observed when the proportionality constraints were defined, these constraints then fix the slopes of all edges incident to v , so that their ordering agrees with that of G .
 - (b) If v is an interior vertex then, as discussed above, the proportionality constraints ensure that the slope order of v 's incident edges in G' either matches that of G on each side, or it is completely reversed on both sides. Let us assume for contradiction that the latter case happens:

$$s_{b_1} \geq s_{b_\ell} \text{ and } s_{a_k} \geq s_{a_1} \quad (8)$$

Let e be the third edge of the triangle with edges a_1 and b_1 , and let f be the third edge of the triangle with edges a_k and b_ℓ , see Figure 6. Then the ordering constraints for the endpoints of e imply $s_{b_1} < s_e < s_{a_1}$, and the ordering constraints for the endpoints of f imply $s_{a_k} < s_f < s_{b_\ell}$. Together with (8), this leads to a contradiction.

From the statement (*), it is now easy to derive that G' satisfies condition (ii). The graph G has triangle and quadrilateral faces. For a triangular face, (*) ensures that the triangle does not degenerate, and is therefore non-crossing, in G' . A quadrilateral face q must be non-convex in G by Property 3 of A-graphs, and for each vertex, the two incident edges of q go in the same direction (left or right). Thus, Property (*) ensures that q is non-crossing in G' .

Therefore, by the result of Devillers et al. cited above, G' is a Fáry drawing. That G' has Δ as the outer face follows from the inclusion of the boundary equations in $A \cdot s = b$. \square

Any solution s to $A \cdot s = b$ has the outer face drawn as Δ , by the boundary equations, and the intersection between e_i and Y is $(0, y_i)$ by construction. Hence, by Lemma 2, ensuring the existence of a solution s to $A \cdot s = b$ satisfying the ordering constraints is enough to prove Theorem 6.

3.3 Strong Ordering Constraints

For some $\epsilon > 0$, we say that $s = (s_1, \dots, s_m)$ satisfies the ϵ -strong ordering constraints if, for each $i, j \in \{1, \dots, m\}$ such that $e_i < e_j$, the inequality $s_j - s_i \geq \epsilon$ holds. Clearly, any s satisfying the ϵ -strong ordering constraints also satisfies the ordering constraints. The converse holds, for a suitably small ϵ (the inequalities being strict in the definition of ordering constraints). The following lemma tells us that this ϵ can be determined by Δ and by the sequence y_1, \dots, y_m .

Lemma 3. *If $\Delta \subset [-1, 1] \times (-\infty, +\infty)$, then any solution s to $A \cdot s = b$ that satisfies the ordering constraints also satisfies the ϵ -strong ordering constraints for all $\epsilon \leq \min\{|y_i - y_j| : e_i < e_j\}$.*

Proof. By Lemma 2 every vertex is contained in Δ . Hence, every x -coordinate is in the interval $[-1, 1]$. If $e_i < e_j$, then the common vertex of e_i and e_j has x -coordinate $(y_j - y_i)/(s_i - s_j)$. From $|(y_j - y_i)/(s_i - s_j)| \leq 1$ we derive $|s_i - s_j| \geq |y_j - y_i| \geq \epsilon$. \square

3.4 Uniqueness of Solutions Satisfying Ordering Constraints

Lemma 3 and the ϵ -strong ordering constraints play a crucial role in our proof because they allow us to appeal to continuity: If the slopes change continuously, it is impossible to violate the ordering constraints without first violating the ϵ -strong ordering constraints. But since the ordering constraints imply the ϵ -strong ordering constraints, it is impossible to violate the ordering constraints at all. An example of this argument will be seen in the following proof.

Lemma 4. *If s is a solution to $A \cdot s = b$ that satisfies the ordering constraints, then s is the unique solution to $A \cdot s = b$.*

Proof. Assume that ϵ is fixed so that $0 < \epsilon \leq \min\{|y_i - y_j| : e_i < e_j\}$.

Suppose, for contradiction, that there is a solution s to $A \cdot s = b$ that satisfies the ordering constraints, but is not unique. Since $A \cdot s = b$ is a linear system, it must then have a 1-parameter family of solutions $s + \lambda r$, $\lambda \in \mathbb{R}$, for some non-zero m -vector r .

Define the continuous (in fact, piecewise linear) function

$$f(\lambda) := \min\{(s_j + \lambda r_j) - (s_i + \lambda r_i) : e_i < e_j\}.$$

Let λ^* be the value with the smallest absolute value $|\lambda^*|$ such that $f(\lambda^*) \leq \epsilon/2$. In order to prove that λ^* exists, it suffices to prove that $f(\lambda) \leq 0$ can be achieved. The vector $r = (r_1, \dots, r_m)$ has at least four zero entries $r_1 = r_a = r_b = r_m = 0$ since the slopes s_1, s_a, s_b , and s_m are fixed. Since G is connected and $m \geq 5$, there is at least one vertex v with two incident edges e_k and e_ℓ such that $r_k = 0$ and $r_\ell \neq 0$. We can thus pick λ so that $(s_\ell + \lambda r_\ell) - (s_k + \lambda r_k) = s_\ell - s_k + \lambda r_\ell = 0$, and then $f(\lambda) \leq 0$. It follows that λ^* exists.

Now we know that, for any λ between 0 and λ^* and for any i and j such that $e_i < e_j$, the difference $(s_j + \lambda r_j) - (s_i + \lambda r_i)$ has the same sign as $s_j - s_i$. It follows that the slopes satisfy the ordering constraints throughout this interval, but then Lemma 3 implies that $f(\lambda^*) \geq \epsilon$, a contradiction. \square

3.5 A Parametric Family of Linear Systems

We now define a parametric family of linear systems $A^t \cdot s = b^t$, parameterized by $0 \leq t \leq 1$ by varying the intersection points $y = (y_1, \dots, y_m)$ and the boundary slopes $h = (h_1, h_a, h_b, h_m)$. Let us first see how the coefficients A and right-hand sides b of the system change when these data are changed. The coefficients of the concurrency constraints (1) depend linearly on y , whereas the proportionality constraints (2–3) remain unchanged, and the boundary constraints (7) have just the constant coefficient 1. In the right-hand sides b^t , the four nonzero entries are the four slopes $h = (h_1, h_a, h_b, h_m)$.

We derive the intermediate systems $A^t \cdot s = b^t$ by linear interpolation between the initial data and the target data: For the “starting system”, we use the intercepts $y^0 = (y_1^0, \dots, y_m^0)$ and the slopes $h^0 = (h_1^0, h_a^0, h_b^0, h_m^0)$ of the edges in the initial drawing G . In the “target system”, we use the specified target intercepts $y^1 = (y_1, \dots, y_m)$ and the slopes $h^1 = (h_1, h_a, h_b, h_m)$ from the target shape Δ . (If Δ is a triangle, this vector includes h_b as an arbitrarily chosen additional slope, as described earlier.)

We define the intermediate data y^t and h^t by linear interpolation:

$$y^t = (1-t)y^0 + ty^1, \quad h^t = (1-t)h^0 + th^1.$$

This defines the corresponding intermediate systems $A^t \cdot s = b^t$, whose coefficients and right-hand sides depend linearly on the parameter t .

It is important to note that the starting system $A^0 \cdot s = b^0$ has at least one solution, namely the slopes $s^0 = (s_1^0, \dots, s_m^0)$ of the edges in the initial drawing G . The proportionality constraints (2–3) were designed in this way, as described in Section 3.1. The concurrency constraints (1) are fulfilled because the initial drawing G is a straight-line drawing. The boundary constraints (7) are fulfilled by construction.

We will show that Lemma 3 can be applied to the system $A^t \cdot s = b^t$, for every $0 \leq t \leq 1$. We define an appropriate threshold value ϵ^* by

$$\begin{aligned} \epsilon^* &= \min_{e_i < e_j} \min_{0 \leq t \leq 1} |y_j^t - y_i^t| \\ &= \min_{e_i < e_j} \min\{|y_j^0 - y_i^0|, |y_j^1 - y_i^1|\} > 0 \end{aligned}$$

Lemma 5. *For every $0 \leq t \leq 1$, a solution s to $A^t \cdot s = b^t$ that satisfies the ordering constraints also satisfies the ϵ^* -strong ordering constraints.*

Proof. We denote by Δ^t the shape of the outer face as specified by y^t and h^t . It suffices to prove that this shape is contained in $[-1, 1] \times (-\infty, +\infty)$, at which point Lemma 3 applies.

We show that each vertex of Δ^t is in $[-1, 1] \times [-\infty, +\infty]$. If such a vertex v does not lie on Y and is incident to the two outer edges e_i and e_j , with $i, j \in \{1, a, b, m\}$ and $e_i < e_j$, it has

x -coordinate $(y_i^t - y_j^t)/(s_j^t - s_i^t)$. Consider the case that $v \in R$. So we want to show that

$$(y_i^t - y_j^t)/(s_j^t - s_i^t) \leq 1 . \quad (9)$$

By the ordering constraints, $s_j^t - s_i^t > 0$, so (9) is equivalent to

$$y_i^t - y_j^t \leq s_j^t - s_i^t.$$

This inequality holds for $t = 0$ and for $t = 1$. The left side is linear in t . Since e_i and e_j are boundary edges, the right side is also linear in t . So the inequality holds for every $t \in [0, 1]$. In the case $v \in L$, the proof that v 's x -coordinate is at least -1 is similar. \square

3.6 Existence (and uniqueness) of solutions to $A^t \cdot s = b^t$

We now prove the following lemma which, together with Lemma 2, completes the proof of Theorem 6.

Lemma 6. *For every $0 \leq t \leq 1$, the system $A^t \cdot s = b^t$ has a unique solution s^t , and this solution satisfies the ordering constraints.*

Proof. Since A^t is an $m \times m$ matrix, the system $A^t \cdot s = b^t$ has a unique solution s^t if and only if $\det A^t \neq 0$. When $\det A^t = 0$, the system may have no solutions or multiple solutions. When $\det A^t \neq 0$, Cramer's Rule states that the solution is $s^t = (s_1^t, \dots, s_m^t)$ where, for each $i \in \{1, \dots, m\}$,

$$s_i^t = \frac{\det A_i^t}{\det A^t}$$

and A_i^t denotes the matrix A^t with its i -th column replaced by b^t . The numerators $\det A_i^t$ and the common denominator $\det A^t$ are polynomials in t , and therefore continuous functions of t . The solution $s^t = (s_1^t, \dots, s_m^t)$ depends continuously on t as long as $\det A^t \neq 0$.

We have already established that $A^0 \cdot s = b^0$ has a solution s^0 that satisfies the ordering constraints. By Lemma 4, this solution is unique, so $\det A^0 \neq 0$.

Let t^* be the smallest $t > 0$ for which $\det A^t = 0$. If such a value does not exist we set $t^* = 2$.

First we argue that, for all $0 \leq t < \min\{1, t^*\}$, the unique solution s^t to $A^t \cdot s = b^t$ satisfies the ordering constraints. This argument is similar to the proof of Lemma 4. Suppose, for a contradiction, that there is a value $0 < t < \min\{1, t^*\}$ for which s^t does not satisfy the ordering constraints. As t increases its value from 0 to $\min\{1, t^*\}$, since s^t depends continuously on t , a value is reached in which s^t violates the ϵ^* -strong ordering constraints, while it does not violate the ordering constraints. However, this contradicts Lemma 3.

If $t^* > 1$ the same argument also extends to $t = 1$ and we are done. Let us therefore assume that $0 < t^* \leq 1$ and derive a contradiction. We look at the one-sided limit $s^* = \lim_{t \uparrow t^*} s^t$ as t approaches t^* from below. Each function s_i^t is a quotient of two polynomials. Thus, for $t \rightarrow t^*$ it can either converge to $s_i^{t^*}$, or diverge to $+\infty$ or $-\infty$. For $t < t^*$ all solutions s^t to the systems $A^t \cdot s = b^t$ satisfy the ϵ^* -strong ordering constraints. Hence, if the limit

exists, by continuity, it also satisfies $A^{t^*} \cdot s^* = b^{t^*}$ and the ϵ^* -strong ordering constraints. By Lemma 4, the solution s^* is the unique solution of $A^{t^*} \cdot s = b^{t^*}$, but this contradicts the assumption that $\det A^{t^*} = 0$.

It remains to rule out the possibility that $A^{t^*} \cdot s = b^{t^*}$ has no solution because $\lim_{t \uparrow t^*} s^t$ does not exist. Define the set $H = \{e_i \in \{e_1, \dots, e_m\} : \lim_{t \uparrow t^*} s_i^t \text{ exists}\}$. The set H corresponds to the edges of G with bounded slope; the remaining edges become vertical as $t \rightarrow t^*$. Lemma 7 below shows that H contains all edges of G . Hence $\lim_{t \uparrow t^*} s^t$ exists. This completes the proof of the lemma. \square

It remains to prove that the set H defined in the proof of Lemma 6 contains all edges in $E(G)$. We start by stating some properties of H .

Proposition 1. *The set H has the following properties:*

- (PR1) *H contains every edge incident to a vertex on the outer face of G .*
- (PR2) *If a vertex $v \notin Y$ has two incident edges in H , then all v 's incident edges belong to H .*
- (PR3) *If a vertex $v \in Y$ has two incident edges $vx, vy \in H$ with $x, y \in L$ or $x, y \in R$, then all v 's incident edges belong to H .*
- (PR4) *If $e_i < e_j < e_k$ and $e_i, e_k \in H$, then $e_j \in H$.*

Proof. (PR1) If v is a boundary vertex with $v \notin Y$, then the location of v is fixed and the y -intercepts and therefore slopes of v 's incident edges are fixed. If v is a boundary vertex with $v \in Y$ then Δ is a triangle and v has three incident edges with slopes fixed by the boundary equations. Two of these edges are boundary edges, so two of these edges lie on the same side, say L , and the third edge lies on the other side, say R . By the proportionality constraints (2) all edges in L are bounded, and thus belong to H . By the proportionality constraints (3) the range of slopes used by the edges in R is bounded, and as one of them is fixed all of them have bounded slopes, and thus belong to H .

(PR2) If v does not lie on Y and two incident edges have bounded slope, then the location of v is fixed in the limit. By the concurrency constraints, the slopes of the remaining incident edges are also bounded.

(PR3) The case where v lies on the outer face is subsumed by (PR1). Assume therefore that v lies on Y and is an interior vertex of G . Define the edges a_1, \dots, a_k and b_1, \dots, b_ℓ incident to v as in Figure 6. Let e be the third edge of the triangle with edges a_1 and b_1 , and let f be the third edge of the triangle with edges a_k and b_ℓ . Assume without loss of generality that two of the edges a_i belong to H . Then, by the proportionality constraints, all edges a_i belong to H , and moreover the range $s_{b_\ell}^t - s_{b_1}^t$ converges to a bounded limit as $t \rightarrow t^*$. It follows that either all slopes of the edges b_j are bounded, or they all diverge to $+\infty$, or they all diverge to $-\infty$. The ordering constraints for the endpoints of e imply $s_{b_1} < s_e < s_{a_1}$. This is inconsistent with $\lim_{t \uparrow t^*} s_{b_1}^t = +\infty$. The ordering constraints for the endpoints of f imply $s_{a_k} < s_f < s_{b_\ell}$. This is inconsistent with $\lim_{t \uparrow t^*} s_{b_\ell}^t = -\infty$. Thus, the only possibility is that all slopes of the edges incident to v are bounded.

(PR4) This follows from the ordering constraints, since, for all $0 \leq t < t^*$, $s_i^t < s_j^t < s_k^t$ and both $\lim_{t \uparrow t^*} s_i^t$ and $\lim_{t \uparrow t^*} s_k^t$ are defined. \square

We now present Lemma 7, which completes the proof of Lemma 6 and Theorem 6. The lemma is proved by induction on something that starts as an A -graph but is then dismantled into something more general. A *near- A -graph* is a graph that satisfies all conditions of an A -graph except that its outer face can be arbitrarily complex, even disconnected. More specifically, each edge of a near- A -graph intersects Y in exactly one point; each inner face is a triangle or a quadrilateral, without any disconnected components inside; each triangular face contains one vertex in each of Y , L , and R ; and for every vertex v on Y each of the faces directly above and below v is either a triangular face or the outer face.

Lemma 7. *Let G be a near- A -graph and let $H \subseteq E(G)$ be a set of edges satisfying Properties (PR1)–(PR4) of Proposition 1. Then $H = E(G)$.*

Proof. The proof is by induction primarily on the number of inner faces of G and secondarily on the number of vertices of G . We dismantle G from outside while maintaining Properties (PR1)–(PR4). In particular:

- If G is not 2-connected but has more than one edge, we will cut it into pieces with fewer edges.
- If G is 2-connected, we will modify it and reduce it to a graph with fewer interior faces, keeping the number of edges fixed.

Eventually, we will reduce to a graph with a single edge, and here the claim is trivial because the edge belongs to the boundary.

We will now go into the details of the proof. We refer to the edges of H simply as *H -edges*.

If G is not connected then we can independently apply induction on each component of G .

If G has a cut vertex v whose removal splits G into components A_1, \dots, A_r then, for each $i \in \{1, \dots, r\}$, we can independently apply induction on the subgraph G_i of G induced by $V(A_i) \cup \{v\}$. Every edge of G_i inherits its classification as an H -edge from its corresponding edge in G . Then it is easy to see that Properties (PR1)–(PR4) are satisfied by G_i . Properties (PR2)–(PR4) are obviously preserved under taking subgraphs. Property (PR1) follows from the fact that every boundary vertex of G_i is also a boundary vertex of G ; this is because each inner face of G is a quadrangle or a triangle, hence G_i cannot be nested inside a different subgraph G_j of G .

We are left with the case in which G is a 2-connected near- A -graph whose outer face is delimited by a simple cycle F . We distinguish two cases.

Case 1. The cycle F contains a vertex v on Y that is incident to an inner triangular face vab . In this case we *open up vab* , merging it into the outer face. Figure 7 illustrates the procedure for the case that ab lies below v , with $a \in L$ and $b \in R$. Let u and w be the predecessor and successor of v on the counterclockwise cycle F , and assume w.l.o.g. that $u \in R$. We construct a new graph G' by splitting v into two vertices x and y that both lie on Y , with y above x . We make x adjacent to u and to every neighbor of v between b and u .

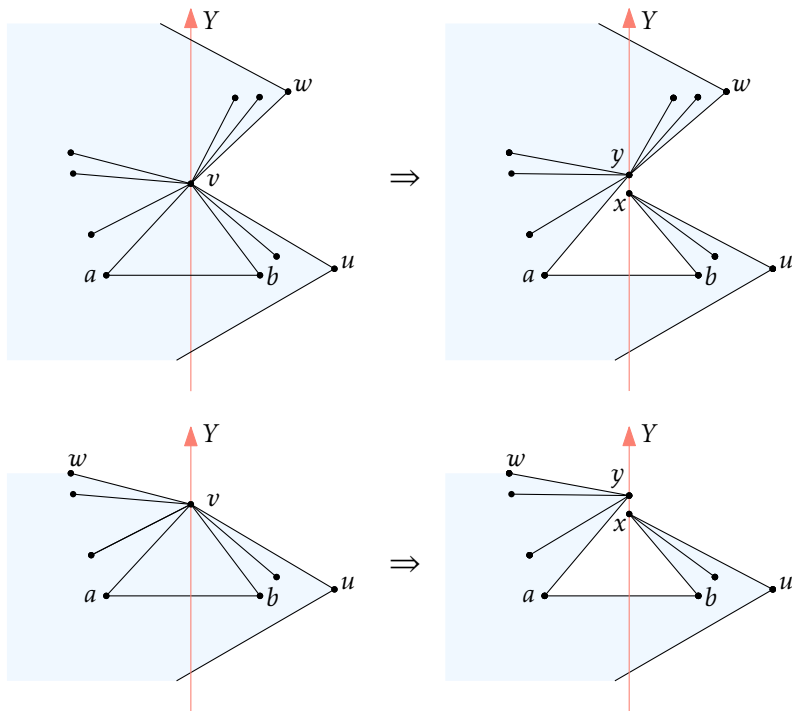


Figure 7: Proof of Lemma 7 with a vertex v on Y . Integrating a triangle into the outer face.

We make y adjacent to the remaining neighbors of v . Figure 7 shows that this procedure works both for $w \in R$ and for $w \in L$. Note that G' has one inner face less than G , hence induction applies.

Case 2. If Case 1 does not hold, every vertex of F on Y has all its neighbours in L or has all its neighbours in R . Since the same is already true for every vertex not on Y , a traversal of F has to zigzag/alternate between edges that move to the left and edges that move to the right. Thus, F must contain some reflex vertex v . The vertex v cannot lie on Y , because otherwise we would be in Case 1. Let uv and vw be the two consecutive edges of F incident on v and let p and q be the intersections of uv and vw with Y , with q above p . (Note that u and/or w may be contained in Y .) This implies that v is a reflex vertex of some inner face $q = vabc$ of G . Indeed, vc is the first edge incident to v intersected by Y and va is the last edge incident to v intersected by Y . (Note the possibility that $a = w$ and/or $c = u$.) We construct a new graph G' by splitting v into two vertices x and y . We make the vertex x adjacent to u and every neighbour z of v such that Y intersects vz before vu . We make y adjacent to all of v 's neighbors that are not adjacent to x . Figure 8 illustrates this procedure; the lower half illustrates the case where $w \in Y$. In G' , q is part of the outer face, so G' has one less inner face than G , hence induction applies.

This finishes the description of how we modify G into G' . Every edge of G' inherits its classification as an H -edge from its corresponding edge in G . We have to show that G' satisfies Properties (PR1)–(PR4). Actually Property (PR1) is the only property that needs to be discussed, as the other properties follow trivially from the fact that G satisfies them.

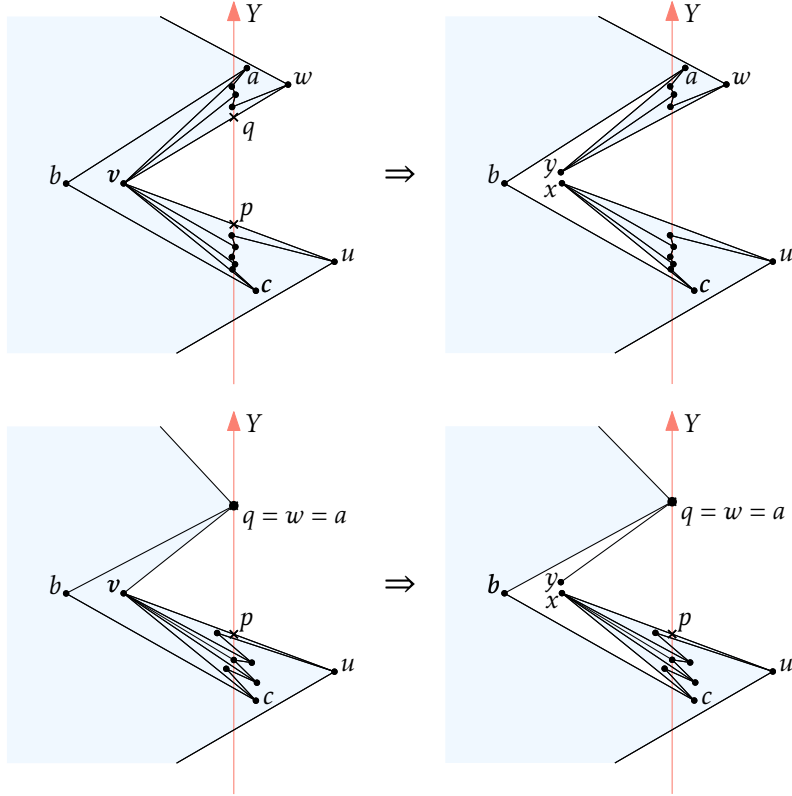


Figure 8: Proof of Lemma 7 for a reflex vertex v . Integrating a quadrilateral into the outer face.

First, note that all edges incident to the new vertices x or y were incident to v before, and thus they are H -edges. Second, all edges incident to any boundary vertex of G' that is also a boundary vertex of G are H -edges, since the edges of G' inherit their classification as H -edges from their corresponding edges in G . It remains to deal with the boundary vertices of G' that are inner vertices of G .

In Case 1 we have two boundary vertices of G' that might be inner vertices of G , namely a and b . These vertices do not lie on Y . By Property (PR1) for G , both va and vb are H -edges, since they are incident to the boundary vertex v . From the ordering constraints around a and b we get $va < ab < vb$ or $vb < ab < va$, and thus, by Property (PR4) for G , we have $ab \in H$. Now we have two H -edges va and ab incident to a , and by Property (PR2) for G all edges incident to a belong to H . It follows that all edges incident to a in G' are H -edges, and similarly for b .

In Case 2 we have three boundary vertices of G' that might be inner vertices of G , namely a , b , and c . By Property (PR1) for G , both va and vc are H -edges, since they are incident to the boundary vertex v . Consider the quadrilateral $q = vabc$ of G . By the ordering constraints, we get $vc < bc < ba < va$ or $va < ba < bc < vc$, depending on whether $v \in L$ or $v \in R$. Thus, by Property (PR4) for G , the edges bc and ba are also H -edges. The vertex b does not lie on Y . The vertex a might lie on Y or not, but if it does, then the

two incident edges va and ab lie in the same half-plane. The same holds for c . Thus by Properties (PR2) or (PR3) for G all edges incident to a , b and c in G belong to H . It follows that all edges incident to a , b , and c in G' are H -edges.

Since G' satisfies Properties (PR1)–(PR4) induction applies and all edges of G' (and thus all edges of G) are H -edges. This completes the proof. \square

4 Triangulations

So far, we have shown that every collinear set in an A-graph is free, and we can even specify for edges that cross the line Y the place where this crossing occurs. We will now apply this to prove for arbitrary planar graphs G that every collinear set is free. We might as well assume that G is a maximal planar graph, i.e., a triangulation.

Theorem 7.

- Let T be a triangulation, i.e., a (not necessarily straight-line) plane drawing of an edge-maximal planar graph.
- Let C be a good proper curve for T . (This means that $C(0) = C(1)$ is in the outer face and the intersection between C and each edge e of T is either empty, a single point, or the entire edge e .)
- Let r_1, \dots, r_k be the mixed sequence of vertices and open edges of T that are intersected by C , in the order in which they are intersected by C —edges of T that lie entirely on C are omitted from this sequence. (They are implicitly represented by their endvertices, which are two consecutive elements r_i and r_{i+1} .)
- Let $y_1 < \dots < y_k$ be a sequence of numbers.
- Let $\epsilon > 0$ be a tolerance parameter.

Then T has a Fáry drawing such that, for each $i \in \{1, \dots, k\}$:

- if r_i is a vertex, then it is drawn at $(0, y_i)$; and
- if r_i is an edge, then the intersection of r_i with Y has y -coordinate in the interval $[y_i - \epsilon, y_i + \epsilon]$.

Moreover, we can specify the shape Δ of the outer triangle, subject to obvious compatibility constraint that it intersects Y in the specified points.

The last condition can be formulated more explicitly: The triangle $\Delta = \alpha\beta\gamma$ is compatible with the given data r_1, \dots, r_k and y_1, \dots, y_k if the following conditions hold:

- If r_1 is a vertex, then $\beta = (0, y_1)$, otherwise $(0, y_1)$ is in the interior of the edge $r_1 = \beta\gamma$; and
- If r_m is a vertex, then $\alpha = (0, y_m)$, otherwise $(0, y_m)$ is in the interior of the edge $r_m = \alpha\gamma$.

If the tolerance ϵ is large, the statement of the theorem allows the order in which the edges cross Y to change. This is not intended, and it can be excluded if we choose $\epsilon < \min\{(y_{i+1} - y_i)/2 : i \in \{1, \dots, k-1\}\}$. In the proof, we will make this assumption.

Proof. We start by classifying the edges of T . An edge that has one endpoint in C^- and the other endpoint in C^+ is a *crossing edge*, otherwise it is a *non-crossing edge*. An edge is

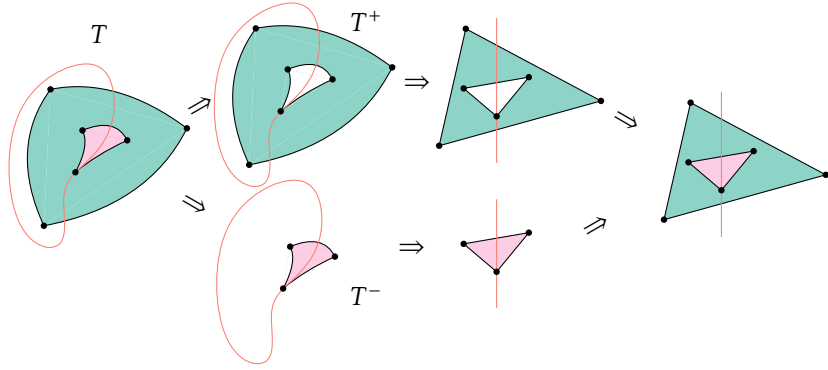


Figure 9: Recursing on separating triangles in the proof of Theorem 7

marked if it intersects C , otherwise it is *unmarked*. The unmarked edges are completely disjoint from C . The marked edges include all crossing edges, but also the edges with one endpoint on C and the edges that lie completely on C .

The proof is a double induction on the number of vertices of T , primarily, and on the number of non-crossing edges of T , secondarily. We begin by describing reductions that allow us to apply the inductive hypothesis. When none of these reductions applies, we arrive at our base case. To handle the base case, we remove every unmarked edge of T , and we will show that we obtain an A-graph, to which we apply Theorem 6.

Separating Triangles. (See Figure 9.) If T contains a separating triangle xyz , then denote by T^+ (respectively, T^-) the triangulation obtained from T by removing the vertices in the interior (respectively, exterior) of xyz . The triangle that xyz delimits an inner face of T^+ and the outer face of T^- . Both $|V(T^+)| < |V(T)|$ and $|V(T^-)| < |V(T)|$, so we can apply induction if necessary.

The case that the interior of xyz does not intersect C is easy. We draw T^+ by induction. In this drawing, we take the triangle representing the cycle xyz , and we draw T^- so that its outer face coincides with this triangle, for example by Tutte's Convex Drawing Theorem [23].

Consider now the case that C intersects the interior of xyz . Then C intersects the boundary in two points: either it passes through a vertex of xyz and the opposite open edge, or it intersects two open edges of xyz . In both cases, the vertices and edges of T intersected by C that are not in T^+ form a nonempty contiguous subsequence r_i, \dots, r_j of r_1, \dots, r_k . Each of r_{i-1} and r_{j+1} is either an edge or a vertex of xyz .

Apply induction on T^+ with the value $\epsilon' := \epsilon/2$ and the sequences $r_1, \dots, r_{i-1}, r_{j+1}, \dots, r_k$ and $y_1, \dots, y_{i-1}, y_{j+1}, \dots, y_k$. In the obtained Fáry drawing of T^+ let Δ' be the triangle representing xyz and let y'_{i-1} and y'_{j+1} be the respective y -coordinates of the intersections of r_{i-1} and r_{j+1} with Y . By the choice of ϵ' we have $y'_{i-1} < y_i < \dots < y_j < y'_{j+1}$. Observe that Δ' is compatible with r_{i-1}, \dots, r_{j+1} and $y'_{i-1}, y_i, \dots, y_j, y'_{j+1}$. We apply induction on T^- with value ϵ using the triangle Δ' and the sequences r_{i-1}, \dots, r_{j+1} and $y'_{i-1}, y_i, \dots, y_j, y'_{j+1}$. Combining the Fáry drawings of T^+ and T^- yields the desired Fáry drawing of T . Thus, from now on we

assume that T has no separating triangles.

Contractible Edges. (See Figure 10.) A face of T is a *crossing face* if it is incident to two crossing edges. We declare an unmarked edge of T to be *contractible* if it is not contained in the boundary of any crossing face.

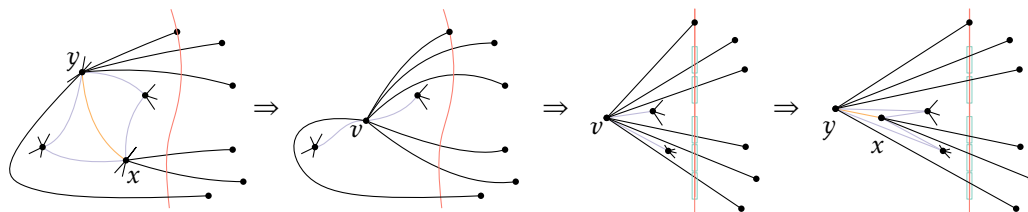


Figure 10: Contracting and uncontracting edges in the proof of Theorem 7

If T contains a contractible edge xy then we contract xy to obtain a new vertex v in a smaller triangulation T' . We then apply induction on T' with the value $\epsilon' = \epsilon/2$ to obtain a Fáry drawing of T' such that each crossing edge e_i crosses Y in the interval $[y_i - \epsilon/2, y_i + \epsilon/2]$.

To obtain a Fáry drawing of T we uncontract v by placing x and y within a ball of radius $\epsilon/2$ centered at v . (Such a placement is always possible, by a standard argument, see, e.g., [14, 24].) Since the distance between y and v and the distance between x and v are each at most $\epsilon/2$, each crossing edge r_i incident to x or y crosses Y in the interval $[y_i - \epsilon, y_i + \epsilon]$, as required. Thus, in the following we assume that T has no separating triangles or contractible edges.

Flippable edges. (See Figure 11.) We declare an unmarked edge xy of T to be *flippable* if there exist distinct vertices z, a, b , and c such that:

- (1) xyz is a non-crossing face of T ;
- (2) xyb, zyc, xza are crossing faces of T ; and either
- (3a) C intersects za, xa, xb, yb, yc , and zc in this order, or
- (3b) C intersects xa, xb, yb, yc, zc , and za in this order. (This case can only occur when xza is the outer face, otherwise xza would be a separating triangle.)

If T contains the flippable edge xy then we remove xy and replace it with zb to obtain a new triangulation T' . Note that, since T has no separating triangles, the edge zb is not already present in T . Further, T' has the same number of vertices of T and one less non-crossing edge. After choosing a crossing coordinate y_{zb} for zb between those y_{xb} and y_{yb} of xb and yb , we can inductively draw T' with tolerance ϵ and sequences $r_1, \dots, xb, zb, yb, \dots, r_k$ and $y_1, \dots, y_{xb}, y_{zb}, y_{yb}, \dots, y_k$.

We claim that in the resulting Fáry drawing of T' , we can replace zb by xy without creating a crossing, thus producing the desired Fáry drawing of T . We show this by establishing that both b and z are convex vertices in $xbyz$. The vertex b is not a reflex vertex in $xbyz$, since bx and by are crossing edges. In Case (3a), the existence of the edges za and zc ensures that, in the Fáry drawing of T' , $xbyz$ is convex. In Case (3b), the triangle xza is convex and $xbyz$ is contained in this triangle, therefore z is a convex vertex in $xbyz$.

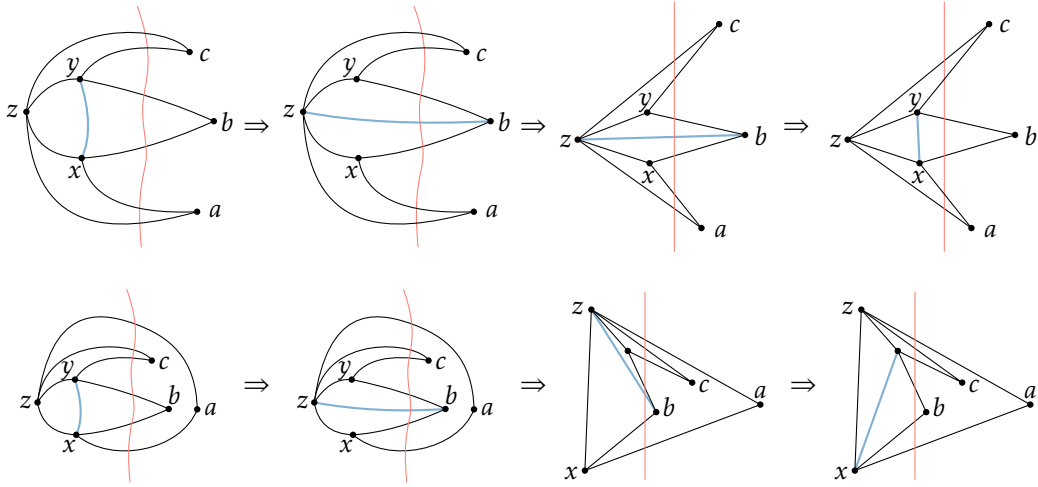


Figure 11: Flipping edges in the proof of Theorem 7. Top row: Case 3a. Bottom row: Case 3b

Edges on C . If T contains an edge xy that lies on C , then we treat it as we treated flip-pable edges. In this case, xy is incident to two triangles xyz and yxb with $z \in C^+$ and $b \in C^-$. We replace xy with an edge zb to obtain a new triangulation T' with the same number of vertices of T and one less non-crossing edge. We apply induction and get a Fáry drawing of T' , in which z and b are on opposite sides of Y and x and y are on Y , hence neither z nor b is a reflex vertex of the quadrilateral $xzyb$. Thus, removing zb and adding xy gives a Fáry drawing of T .

The Base Case. We are left with the case in which T is a triangulation with no separating triangles, no contractible edges, no flippable edges, and no edge contained in C . If T is the complete graph on three or four vertices, then the proof is trivial, so we may assume that T has at least 5 vertices.

We will simply omit the marked edges. The result will be an A-graph, to which we can apply Theorem 6. In the resulting drawing, we will see that we can reinsert the omitted edges without producing crossings.

Claim 1. Any unmarked edge xy in C^+ is on the boundary of two faces xyz and yxb where $z, b \in C \cup C^-$, see Figure 12a–b.

Proof. Since xy is not contractible, at least one of xyz and yxb is a crossing triangle, so at least one of z and b , say b , is in C^- . Suppose then, for the sake of contradiction, that $z \in C^+$. Since neither zx nor yz is contractible, they must be incident to crossing faces xza and zyc , respectively, see Figure 12c. If $a = b = c$, then T is the complete graph on four vertices, which we have already ruled out. Therefore, assume without loss of generality that $b \neq c$. We have $a \neq c$, because otherwise xya would be a separating triangle that separates z from b . Similarly, $a \neq b$, otherwise byz would separate x from a . This leaves us in the situation in which we have distinct vertices x, y, z, a, b , and c such that $xyz \in C^+$, such that $xyb, zyc,$

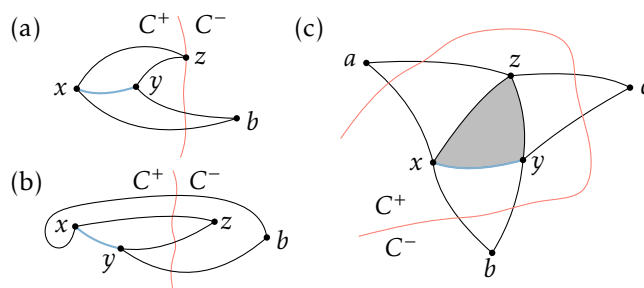


Figure 12: (a–b) shows how the two triangles incident to an unmarked edge xy could look. In case (b), yxz is the outer face. (c) The triangles adjacent to xyz in the proof of Claim 1.

and xza are crossing faces of T , and such that xyz is a non-crossing face of T . Then at least one of xy , yz , or zx is a flippable edge. This contradiction proves the claim. \square

Symmetrically, every unmarked edge xy in C^- is incident to two faces xyz and yxz with $z, b \in C \cup C^+$. This implies that no face of T contains more than one unmarked edge.

Thus, every unmarked edge of T is incident to two faces that intersect C . The union of these two faces is a quadrilateral whose boundary consists of four edges that intersect C . Let \tilde{G} denote the plane drawing obtained by removing all unmarked edges from T . By Theorem 5, we know that \tilde{G} has a Fáry drawing G whose edges and vertices intersect Y in the same order as they intersect C in \tilde{G} . We have the following.

Claim 2. G is an A -graph.

Proof. In order to prove the claim, we check each of the properties of an A -graph (Definition 2).

1. The removal of unmarked edges and the fact that T has no edge entirely on C ensure that every edge of G intersects Y in exactly one point.
2. Because no face of T is incident to more than one unmarked edge, each face of G is a quadrilateral or a triangle.
3. A quadrilateral face $q = abcd$ appears in G when we remove the unmarked edge ac from T . This, and the fact that every edge of q intersects Y , ensures that a or c is a reflex vertex of q .
4. The only triangular faces of G are those consisting of three marked edges, which necessarily have one vertex in each of Y , L , and R .
5. Since T has no edge on C , every vertex of T on C is incident to two triangular faces (one above and one below) each having three marked edges. These faces are still present in G .

This concludes the proof of the claim. \square

We would now like to apply Theorem 6 to obtain a Fáry drawing of G in which, for each $i \in \{1, \dots, k\}$, the intersection of r_i with Y is at $(0, y_i)$ and the appropriate vertices on the outer face of T map to the vertices of the triangle Δ . Before doing so, we must first

prescribe an outer face Δ' for the Fáry drawing of G . If the outer face of G is a 3-cycle, then we use $\Delta' = \Delta$. Otherwise, suppose the outer face of G is a 4-cycle $\alpha\beta\gamma$ and $\alpha\beta$ is an unmarked edge of T . In this case, the locations of α , β , and γ are given by the three vertices of Δ (with α and β both on the same side of Y). If x lies on Y , then $x = r_i$ for some i , and the position of x is determined by y_i . Otherwise, it is determined by the positions of α and β and the values y_i and y_j , where $r_i = \alpha x$ and $r_j = \beta x$.

In this way, we can apply Theorem 6 to obtain a Fáry drawing of G in which the intersection of r_i with Y is at $(0, y_i)$. Each internal edge ac of T not in G corresponds to a quadrangular face $q = abcd$ of G in which a or c is a reflex vertex. Therefore, the edge ac can be added to the drawing without introducing crossings. A single external edge $\alpha\beta$ on the outer face of T might not appear in G . In this case the outer face of G is a quadrilateral $q' = \alpha x \beta \gamma$ in which x is a reflex vertex, so the segment $\alpha\beta$ lies outside of q' , and the edge $\alpha\beta$ can therefore be added to the drawing of G without introducing crossings. Therefore reinserting each edge of T not in G gives the desired Fáry drawing of T (and the choice of Δ' ensures that the outer face of this drawing is Δ). This concludes the proof of Theorem 7. \square

We are finally ready to prove Theorem 1. Given a plane drawing of a graph G , a collinear set S in G , and any $y'_1 < \dots < y'_{|S|}$, we need to prove that G has a Fáry drawing in which the vertices in S are drawn at $(0, y'_1), \dots, (0, y'_{|S|})$. Let C be the proper good curve that contains S and let $v_1, \dots, v_{|S|}$ denote the vertices of S in the order they are encountered when traversing C clockwise starting at the outer face.

If G is not a triangulation then we add edges to triangulate it in such a way that each edge we add has a proper intersection with C . To do this, we first add each edge $v_i v_{i+1}$ where v_i and v_{i+1} are in a common face of G to obtain an augmented graph G' . If each edge $v_i v_{i+1}$ added this way is drawn so that it coincides with the subcurve of C joining v_i and v_{i+1} , then C will be a proper good curve for G' . The interior of each face of G' is either entirely contained in the interior of C or entirely contained in the exterior of C . At this point we can greedily add edges to G' until it becomes a triangulation. By Theorem 2, the property that S is collinear set is preserved.

Theorem 5 implies that there exists a Jordan curve C that is admissible for G and that contains the vertices of S in some order, say $v_1, \dots, v_{|S|}$. The curve C intersects a subset of the edges and vertices of G in some order r_1, \dots, r_k . We extend the sequence y'_i by inserting additional elements, resulting in a sequence $y_1 < \dots < y_k$ so that, whenever $r_i = v_j$ for some $i \in \{1, \dots, k\}$ and $j \in \{1, \dots, |S|\}$, then $y_i = y'_j$. We select any triangle Δ that is compatible with r_1, \dots, r_k and y_1, \dots, y_k and choose $\epsilon = (1/3) \min\{y_{i+1} - y_i : 1 \leq i \leq k-1\}$. Theorem 7 then gives us a Fáry drawing of G in which the vertices in S are at $(0, y'_1), \dots, (0, y'_{|S|})$, as required by Theorem 1. Finally, edges that were inserted to create a triangulation are simply removed, and we obtain the desired Fáry drawing of the initial graph. \square

5 Open Problems

In this paper we proved that every collinear set is a free set. Several problems concerning collinear and free sets remain open. Here we mention our favorite two.

Let $f(n)$ be the minimum, over all n -vertex planar graphs G , of the size of the largest collinear set in G . What is the growth rate of $f(n)$? The best known bounds are $f(n) \in \Omega(\sqrt{n})$ and $f(n) \in \mathcal{O}(n^\sigma)$, for $\sigma < 0.986$ [5, 22]. Our results prove that $f(n)$ is also the minimum size of the largest free set over all n -vertex planar graphs; this makes determining the growth rate of $f(n)$ even more relevant. For example, any improvement in the lower bound would immediately give an improved result for untangling planar graphs.

We find it interesting to understand whether our main theorem, Theorem 1, can be generalized so that the y -coordinates are arbitrarily prescribed not only for the vertices on Y , but also for the crossing points of the edges with Y . Note that Theorem 7 *almost* gives this generalization, as every edge crossing Y is at most ϵ away from its prescribed crossing point, for any arbitrarily small ϵ .

Acknowledgement

Part of this research was conducted during the 5th and the 6th Workshops on Geometry and Graphs, held at the Bellairs Research Institute, March 5–10, 2017 and March 11–16, 2018. We are grateful to the organizers and participants for providing a stimulating research environment.

References

- [1] Patrizio Angelini, Carla Binucci, William S. Evans, Ferran Hurtado, Giuseppe Liotta, Tamara Mchedlidze, Henk Meijer, and Yoshio Okamoto. Universal point subsets for planar graphs. In Kun-Mao Chao, Tsan-sheng Hsu, and Der-Tsai Lee, editors, *23rd International Symposium on Algorithms and Computation (ISAAC 2012)*, volume 7676 of *LNCS*, pages 423–432. Springer, 2012.
- [2] Michael J. Bannister, Zhanpeng Cheng, William E. Devanny, and David Eppstein. Superpatterns and universal point sets. *J. Graph Algorithms Appl.*, 18(2):177–209, 2014.
- [3] Luis Barba, William Evans, Michael Hoffmann, Vincent Kusters, Maria Saumell, and Bettina Speckmann. Column planarity and partially-simultaneous geometric embedding. *J. Graph Algorithms Appl.*, 21(6):983–1002, 2017.
- [4] Prosenjit Bose. On embedding an outer-planar graph in a point set. *Comput. Geom.*, 23(3):303–312, 2002.
- [5] Prosenjit Bose, Vida Dujmović, Ferran Hurtado, Stefan Langerman, Pat Morin, and David R. Wood. A polynomial bound for untangling geometric planar graphs. *Discrete & Computational Geometry*, 42(4):570–585, 2009.
- [6] Javier Cano, Csaba D. Tóth, and Jorge Urrutia. Upper bound constructions for untangling planar geometric graphs. *SIAM J. Discrete Math.*, 28(4):1935–1943, 2014.
- [7] Netzahualcoyotl Castañeda and Jorge Urrutia. Straight line embeddings of planar graphs on point sets. In *Proc. the 8th Canadian Conference on Computational Geometry (CCCG)*, pages 312–318, 1996.

-
- [8] Josef Cibulka. Untangling polygons and graphs. *Discrete & Computational Geometry*, 43(2):402–411, 2010.
- [9] Giordano Da Lozzo, Vida Dujmović, Fabrizio Frati, Tamara Mchedlidze, and Vincenzo Roselli. Drawing planar graphs with many collinear vertices. *Journal of Computational Geometry*, 9(1):94–130, 2018.
- [10] H. de Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41–51, 1990.
- [11] Olivier Devillers, Giuseppe Liotta, Franco P. Preparata, and Roberto Tamassia. Checking the convexity of polytopes and the planarity of subdivisions. *Comput. Geom.*, 11(3-4):187–208, 1998.
- [12] Emilio Di Giacomo, Walter Didimo, Giuseppe Liotta, Henk Meijer, and Stephen K. Wismath. Planar and quasi-planar simultaneous geometric embedding. *The Computer Journal*, 58(11):3126–3140, 2015.
- [13] Vida Dujmović. The utility of untangling. *J. Graph Algorithms Appl.*, 21(1):121–134, 2017.
- [14] István Fáry. On straight line representations of planar graphs. *Acta Sci. Math. (Szeged)*, 11:229–233, 1948.
- [15] Xavier Goaoc, Jan Kratochvíl, Yoshio Okamoto, Chan-Su Shin, Andreas Spillner, and Alexander Wolff. Untangling a planar graph. *Discrete & Computational Geometry*, 42(4):542–569, 2009.
- [16] P. Gritzmann, B. Mohar, J. Pach, and R. Pollack. Embedding a planar triangulation with vertices at specified points (solution to problem e3341). *Amer. Math. Monthly*, 98:165–166, 1991.
- [17] Mihyun Kang, Oleg Pikhurko, Alexander Ravsky, Mathias Schacht, and Oleg Verbitsky. Untangling planar graphs from a specified vertex position - Hard cases. *Discrete Applied Mathematics*, 159(8):789–799, 2011.
- [18] Maciej Kurowski. A 1.235 lower bound on the number of points needed to draw all n -vertex planar graphs. *Information Processing Letters*, 92(2):95–98, 2004.
- [19] Tamara Mchedlidze, Marcel Radermacher, and Ignaz Rutter. Aligned drawings of planar graphs. In Fabrizio Frati and Kwan-Liu Ma, editors, *25th International Symposium on Graph Drawing and Network Visualization (GD 2017)*, volume 10692 of LNCS, pages 3–16. Springer, 2017.
- [20] P. J. Owens. Non-Hamiltonian simple 3-polytopes whose faces are all 5-gons or 7-gons. *Discrete Mathematics*, 36(2):227–230, 1981.
- [21] János Pach and Gábor Tardos. Untangling a polygon. *Discrete & Computational Geometry*, 28(4):585–592, 2002.

-
- [22] Alexander Ravsky and Oleg Verbitsky. On collinear sets in straight-line drawings. In Petr Kolman and Jan Kratochvíl, editors, *37th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2011)*, volume 6986 of *LNCS*, pages 295–306. Springer, 2011. Preliminary version on arXiv:[0806.0253](https://arxiv.org/abs/0806.0253).
- [23] William T. Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, 13:743–768, 1963.
- [24] David R. Wood. A simple proof of the Fáry-Wagner theorem. *CoRR*, abs/cs/0505047, 2005.