On Graphs Supported by Line Sets *

Vida Dujmović[1], William Evans[2], Stephen Kobourov[3], Giuseppe Liotta[4], Christophe Weibel[5], and Stephen Wismath[6]

[1] School of Computer Science Carleton University `cgm.cs.mcgill.ca/~vida`
[2] Dept. of Computer Science, Univ. of British Columbia `www.cs.ubc.ca/~will`
[3] Dept. of Computer Science, Univ. of Arizona `www.cs.arizona.edu/~kobourov`
[4] Dept. of Computer Science, Univ. of Perugia `www.diei.unipg.it/~liotta`
[5] Dept. of Mathematics, McGill University `www.math.mcgill.ca/~weibel`
[6] Dept. of Computer Science, Univ. of Lethbridge `www.cs.uleth.ca/~wismath`

Abstract. For a set S of n lines labeled from 1 to n, we say that S supports an n-vertex planar graph G if for every labeling from 1 to n of its vertices, G has a straight-line crossing-free drawing with each vertex drawn as a point on its associated line. It is known from previous work [4] that no set of n parallel lines supports all n-vertex planar graphs. We show that intersecting lines, even if they intersect at a common point, are more "powerful" than a set of parallel lines. In particular, we prove that every such set of lines supports outerpaths, lobsters, and squids, none of which are supported by any set of parallel lines. On the negative side, we prove that no set of n lines that intersect in a common point supports all n-vertex planar graphs. Finally, we show that there exists a set of n lines in general position that does not support all n-vertex planar graphs.

1 Introduction

We consider the effect of restricting the placement of vertices in a planar, straight-line, crossing-free embedding of a planar graph. Every vertex has an associate region of the plane where it can be placed. If each region is the whole plane then the regions support all planar graphs. If the regions are points then they fail to support even such a simple class of graphs as paths. Our interest is in what classes of planar graphs are supported by particular families of vertex regions. Specifically, in this paper we focus on vertex regions that are lines.

A set of segments is *crossing-free* if no two segments intersect in their interiors. A *vertex labeling* of a graph $G = (V, E)$ is a bijection $\pi : V \to [n]$. A set R of n regions (subsets of $\mathbb{R}^2$) labeled from 1 to n *supports* a graph G *with vertex labeling* π if there exists a set of distinct points $p_1, p_2, \ldots, p_n$ such that p_i lies in region i for all i and the segments $\overline{p_{\pi(u)}p_{\pi(v)}}$ for $(u, v) \in E$ are crossing-free. The set R of n labeled regions *supports* a graph G if R supports G with vertex labeling π for every vertex labeling π. As an example of the use of this terminology, we show that every n-pinwheel (set of n labeled lines that share a common point) supports any n-squid (see definition below).

* Research supported in part by: NSERC, MIUR under project AlgoDEEP prot. 2008TFBWL4. The research in this paper started during the *McGill/INRIA Workshop at Bellairs*. The authors thank the organizers and the other participants for useful discussions.

While we focus on embeddings that prescribe a specific region for each vertex, the problem is also interesting if each vertex may be placed in any one of the regions. In this variant, a set of regions R *supports* a graph $G = (V, E)$ *without mapping* if there *exists* a bijection π from V to R such that R supports G with vertex labeling π. Rosenstiehl and Tarjan [8] posed the question of whether there exists a point set of size n that supports without mapping all n-vertex planar graphs: a *universal* point set for all planar graphs. De Fraysseix *et al.* [3] resolved the question in the negative by presenting a set of n-vertex planar graphs that requires a point set of size $\Omega(n + \sqrt{n})$. For some classes of n-vertex planar graphs, universal point sets of size n have been found. In particular, Gritzmann *et al.* [6] showed that any set of n points in general position forms a universal point set for trees and indeed for all outerplanar graphs, for which Bose [2] gave an efficient drawing algorithm.

Embedding with mapping is an even more restricted version of the problem. For example, any set of n points supports without mapping all n-vertex paths. Whereas, for large enough n, no set of n points supports all n-vertex paths. For $n \geq 5$, every set of n points contains a subset of three collinear points or four points in convex position.[7] In both cases, it is easy to devise a vertex mapping of three (respectively four) consecutive vertices of any n-vertex path to that subset of points that forces an edge crossing.

If the straight-line edge condition is relaxed in this mapped setting, Pach and Wenger [7] showed that any set of n points supports all n-vertex planar graphs, however $\Omega(n)$ bends per edge may be necessary in any crossing-free drawing. Even if the mapping constraint is relaxed to just two colors: the red vertices must be mapped to any red point and the blue vertices to any blue point, Badent *et al.* [1] proved that $\Omega(n)$ bends per edge are sometimes necessary.

Estrella-Balderrama *et al.* [4] show that any set of n parallel lines supports exactly the class of *unlabeled level planar (ULP) graphs*. This class of graphs contains several sub-classes of trees (namely caterpillars, radius-2 stars and degree-3 spiders) [4] and a restricted set of graphs with cycles (such as generalized caterpillars) [5]. The simplest class of trees not supported by parallel lines is the class of lobsters.

We show that any set of n lines that intersect at a common point supports a larger sub-class of n-vertex trees than the ULP graphs. We further show that no set of n lines that intersect at a common point supports all n-vertex planar graphs. Whether such a set of lines supports all trees is a natural open question. We also show that there exists a set of n lines in general position that does not support all n-vertex planar graphs. Here, a set of lines is considered in general position if no two lines are parallel and no three lines intersect in a common point. The main open question remaining is whether there exists a set of lines in general position that supports all planar graphs.

2 Pinwheels

Definition 1. A *pinwheel* is an arrangement of n distinct lines that intersect the origin and are labeled from 1 to n in clockwise order. Each line in the pinwheel is called a *track*.

[7] Eszter Klein's Happy Ending problem.

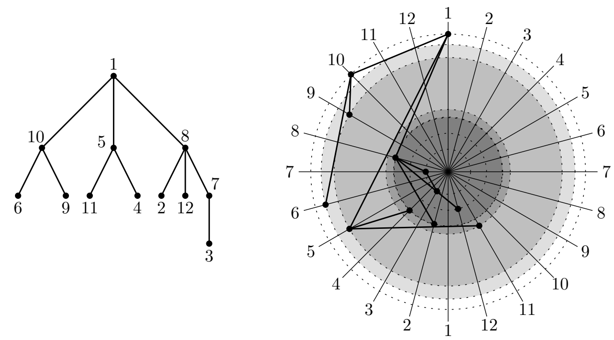

Fig. 1. (a) A labeled lobster with spine vertices $v_1 = 1$, $v_2 = 8$. (b) An embedding using the algorithm. The dotted circles indicate the empty discs at each step.

As the next lemma shows, pinwheels are an interesting family of line sets to consider when investigating whether more general families support planar graphs.

Lemma 1. *Any class of graphs supported by every n-line pinwheel is also supported by every arrangement of n lines, no two of which are parallel.*

Proof: Determine a circle that contains all line intersections. By scaling the arrangement down we can make the radius of this circle arbitrarily small, rendering it effectively into a pinwheel. ■

3 Graphs Supported by Arrangements of Lines

In this section we describe non-ULP families of planar graphs that are supported by every arrangement of lines, no two of which are parallel. We use pinwheels as supporting sets for the graphs in these families since by Lemma 1 the results will then apply to the more general arrangements. We begin by studying lobsters, then extend the result to squids and finally consider outerpaths.

- A *caterpillar* is a graph in which the removal of all degree one vertices and their incident edges results in a path. This path is called the *spine* of the graph.
- A *lobster* is a graph in which the removal of all degree one vertices and their incident edges results in a caterpillar.
- A *squid* is a subdivision of a lobster.
- An *outerpath* is an outerplanar graph whose weak dual is a path (where the weak dual is obtained from the dual by removing the vertex corresponding to the outerface and its adjacent edges).

Lemma 2. *Every n-line pinwheel supports any n-vertex lobster.*

Proof: Let L be a lobster with n vertices and spine vertices $v_1, v_2, \ldots, v_k$. We compute a straight-line embedding of L on any labeled n-pinwheel such that no two edges cross for any vertex labeling of L. We place the vertices in order of a preorder traversal of L, where L is considered as a tree rooted at v_1 and such that each spine vertex is the last descendant of its parent. At any step, there is a set of vertices not all of whose children have been drawn – call these the *active* vertices. We maintain the invariant that all active vertices can "see" the origin (*i.e.*, the segment from the embedded active vertex to the origin does not intersect a segment of the drawing). As a consequence, all active vertices can see an empty disk (*i.e.*, a disk that does not intersect a segment of the current drawing) of nonzero radius centered at the origin and intersecting every track twice; see Fig. 1. At each step, the current vertex is placed at one of the two intersection points of its track and the boundary of the largest empty disk centered at the origin that is seen by all active vertices. The intersection point that is chosen is the one that is encountered first in counter-clockwise radial order from the track of its already-placed parent. (The first vertex, v_1, is initially placed on its corresponding track at an arbitrary point that is not the origin.)

The correctness of the drawing algorithm is proved by induction on the length of the spine. While a spine vertex is active, only vertices at distance at most two from it are drawn. Since the radial distance between a vertex and its parent is less than 180 degrees, we maintain the invariant that each active vertex sees the origin. ■

Lemma 3. *Every n-line pinwheel supports any n-vertex squid.*

Proof: We extend the algorithm of Lemma 2 to the drawing of squids. A squid G' can be obtained from a lobster G by subdividing edges of G. For each vertex v created by subdividing an edge (u, w) of G, we define v's *lobster parent* as the closer of u or w to the root v_1. We draw the vertices in order of a preorder traversal of the graph. But at each step, the position chosen for a vertex is on the track that is encountered first in counter-clockwise radial order from the track of its lobster parent instead of its parent. As a result, the whole path obtained by subdividing an edge is drawn at the radial distance of at most 180 degrees from its lobster parent. As in the proof of Lemma 2, every active vertex can see the origin. ■

With similar techniques as those of Lemmas 2 and 3, the following can be proved.

Lemma 4. *Every n-line pinwheel supports any n-vertex outerpath.*

Lemmas 3, 4, and 1 imply the following.

Theorem 1. *Every arrangement of n lines, no two of which are parallel, supports any n-vertex squid and any n-vertex outerpath.*

4 Non-supporting Line Sets

In this section, we show that there is a labeled planar graph that is not supported by any pinwheel. We also show that there exists a family of n-line sets, where each set is in general position, that does not support all n-vertex planar graphs. Note that this doesn't

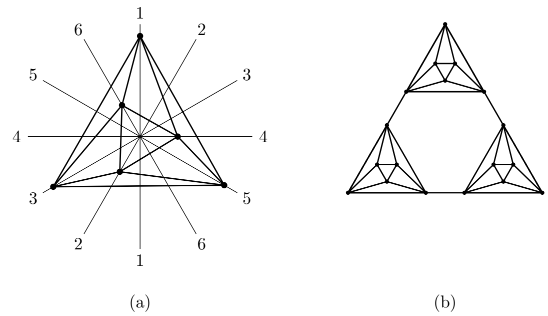

Fig. 2. (a) Graph G_6 can be realized only if the origin of the pinwheel is contained in an internal face. (b) Graph G_6^3 consists of three copies of G_6 connected by three edges.

rule out the possibility that some family of n-line sets in general position could support all n-vertex planar graphs.

Both arguments rely on graphs that use as a building block the graph G_6 in Fig. 2(a). It is not difficult to show that any straight-line, crossing-free embedding of G_6 with the given labeling requires that the origin of the pinwheel (with tracks labeled in clockwise order) is in an internal face. We prove a slightly stronger statement since we will need it in the proof of Theorem 3. For a set, S, of lines, no two of which are parallel, define the *core*, $C(S)$, of S to be the union of the intersections, finite edges, and bounded cells of the arrangement of S.

Lemma 5. *Let S be any set of lines, no two of which are parallel, labeled so that they intersect some line at infinity in the order 1,2,3,4,5,6. In any straight-line, crossing-free embedding of G_6 (labeled as in Fig. 2(a)) on S, the core $C(S)$ intersects some internal face of the embedding.*

Proof: Suppose for the sake of contradiction that the core $C(S)$ lies in the external face of some embedding of G_6. Thus each vertex of G_6 lies on a half-line (of the arrangement of S) that does not intersect $C(S)$ and these six half-lines intersect a line at infinity in some order. Let us assume initially that this order is $1, 2, 3, 4, 5, 6$.

Consider edge $(1, 5)$ in the embedding of G_6, and the line ℓ that contains edge $(1, 5)$. For any pair of vertices $a, b \in \{2, 3, 4\}$ of G_6, if a and b are in the distinct half-planes bounded by ℓ then a does not see b, that is, the segment between a and b intersects edge $(1, 5)$. Since both 3 and 4 are adjacent to 2 in G_6, all three of these vertices are in the same half-plane determined by ℓ. Moreover, 2 is contained inside of a triangle, T, determined by either $1, 3, 5$ or $1, 4, 5$. Since 6 is adjacent to 2 in G_6, 6 is inside of T as well. However, line 6 does not intersect T, which provides a desired contradiction.

A similar argument holds for the other possible half-line orders. ∎

To construct the graph that is not supported by any pinwheel, we make three copies of G_6 and label them so that the vertex labeled k in the original graph is labeled $6(i-1)+k$ in the ith copy. Finally, the graph G_6^3 is created by connecting the three labeled copies of G_6 with the help of three additional edges, as shown in Fig. 2(b).

Theorem 2. *Planar graph G_6^3 is not supported by any pinwheel.*

Proof: Assume for the sake of contradiction that there is a straight-line, crossing-free drawing of the labeled graph G_6^3 on the pinwheel. By Lemma 5, each of the copies of G_6 can be realized crossing-free and with straight line edges only if the origin of the pinwheel is contained in an internal face. Without loss of generality, that implies that the first copy of G_6 is inside an internal face of the second copy of G_6, and both are inside an internal face of the third copy of G_6. That provides a desired contradiction, since the edge connecting the first copy with the third copy must cross some edge of the second copy. ■

We now turn our attention to lines in general position. One might hope that lines in general position (*i.e.*, no two lines are parallel and no three lines intersect in a common point) provide enough freedom in the placement of vertices to support any planar graph. We show that the general position assumption alone is not sufficient. Specifically, we prove that there exists a family of n-line sets such that each set is in general position and not all planar graphs are supported by a line set in the family.

A *parabolic grid* of n lines consists of lines T_j for $j = 1, \ldots, n$ through points $(0, (j-1)/(n-1))$ and $((n-j)/(n-1), 0)$. By again using a graph containing multiple copies of G_6 and by exploiting a geometric property of triangulated graphs drawn on parabolic grids, we show (see the appendix for details):

Theorem 3. *Parabolic grids do not support all planar graphs.*

5 Conclusion and Open Problems

Whether there exists some set of n lines that does support all n-vertex planar graphs is a natural question that is still open. It is also not known whether pinwheels support all trees.

References

1. M. Badent, E. D. Giacomo, and G. Liotta. Drawing colored graphs on colored points. *Theor. Comput. Sci.*, 408(2-3):129–142, 2008.
2. P. Bose. On embedding an outer-planar graph in a point set. *CGTA: Computational Geometry: Theory and Applications*, 23(3):303–312, 2002.
3. H. de Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41–51, 1990.
4. A. Estrella-Balderrama, J. J. Fowler, and S. G. Kobourov. Characterization of unlabeled level planar trees. *Computational Geometry: Theory and Applications*, 42(7):704–721, 2009.
5. J. J. Fowler and S. G. Kobourov. Characterization of unlabeled level planar graphs. In *15th Symposium on Graph Drawing (GD)*, pages 37–49, 2007.
6. P. Gritzmann, B. Mohar, J. Pach, and R. Pollack. Embedding a planar triangulation with vertices at specified points. *American Math. Monthly*, 98:165–166, 1991.
7. J. Pach and R. Wenger. Embedding planar graphs at fixed vertex locations. *Graphs and Combinatorics*, 17:717–728, 2001.
8. P. Rosenstiehl and R. E. Tarjan. Rectilinear planar layouts and bipolar orientations of planar graphs. *Discrete Computational Geometry*, 1(4):343–353, 1986.

A Proof of Lemma 4

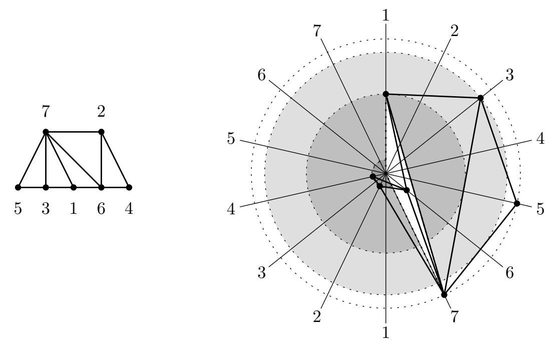

Fig. 3. (a) A labeled outerpath on 7 vertices. (b) An embedding using the algorithm. The vertices are processed in the order $v_1 = 5$, $v_2 = 7$, $v_3 = 3$, $v_4 = 1$, $v_5 = 6$, $v_6 = 2$, and $v_7 = 4$. The dotted circles represent the empty discs at each step. Vertex 6 is placed inside the empty triangle, $\varDelta$, shown in the figure.

Proof: Let $G = (V, E)$ be an n-vertex outerpath and let v_1 and v_n be the degree-2 vertices in G. Let $e_1, e_2, \ldots, e_{n-1}$ be the edges of G that are crossed by the Hamiltonian cycle in the dual of G in order from the outerface. Let (v_1, v_2) be the edge e_1 (with v_1 of degree 2) and in general let v_{i+1} be the new vertex introduced by edge e_i.

Place vertices v_1 and v_2 on their corresponding tracks (not at the origin). Note that all points of the embedded edge (v_1, v_2) see an empty disk of positive radius centered at the origin.

For $i = 2$ to $n-1$, place vertex v_{i+1} on its corresponding track so that the embedded edge e_i sees the origin.

For the correctness of the algorithm it suffices to show that we can maintain the following invariant: the embedded edge e_i sees an empty disk of positive radius centered at the origin. The invariant holds for e_1. In general, when placing v_{i+1} we know that the embedded edge e_i sees the origin and therefore the triangle $\varDelta$ defined by the origin O and embedded edge e_i is empty. If the track for v_{i+1} intersects $\varDelta$ then we pick any point of that track that is inside $\varDelta$ (other than the origin) and maintain the invariant. Otherwise the track for v_{i+1} intersects the empty disk at two points. At least one of these two points forms a triangle with O and the already embedded endpoint of e_i that is contained in the union of $\varDelta$ and the empty disc (see Fig. 3). Placing v_{i+1} at this point preserves the invariant. ■

B Parabolic grids do not support all planar graphs

Definition 2. A *parabolic grid* of n lines consists of lines T_j for $j = 1, \ldots, n$ through points $(0, (j-1)/(n-1))$ and $((n-j)/(n-1), 0)$.

Definition 3. The *closed corner* is the core of the lines $T_1, T_2, \dots, T_n$, i.e., the union of the intersections, finite edges, and bounded cells in the arrangement of $T_1, T_2, \dots, T_n$.

Note that with the coordinates chosen, any intersection of two lines in the parabolic grid is contained in the square $[0, 1] \times [0, 1]$.

First, we show that if parabolic grids support all planar graphs then any planar graph can be embedded using an area of the closed corner as small as desired.

Lemma 6. *Suppose parabolic grids support all planar graphs. Then, for any $\varepsilon > 0$, any graph can be drawn inside a triangle, such that the area of the intersection of the triangle with the closed corner is less than ε.*

Proof: We can assume that the graph is triangulated, and so its external face is a triangle for any planar embedding.

Let us define $T(x)$ as a track going through the points $(0, x)$ and $(1 - x, 0)$ for any $x \in [0, 1]$. In particular, the track T_j of a parabolic grid of n lines is the track $T((j-1)/(n-1))$. Suppose there is a certain graph G, and values $x_1, \dots, x_k$ such that any embedding of G mapped to the tracks $T(x_1), \dots, T(x_k)$ must have an intersection of area more than ε with the closed corner. We show that there is a contradiction with the assumption that parabolic grids support all planar graphs.

Let $f(\delta)$ be the minimum area of the intersection with the closed corner of any embedding of G on the tracks $T(x_1 + \delta), \dots, T(x_k + \delta)$. In particular, $f(0) = \varepsilon$. The function $f(\delta)$ being continuous, there is a certain interval $[-\delta', \delta']$ such that $f(\delta) \geq \varepsilon/2$ in this interval.

We set the number of tracks, n, to be greater than $1/(\varepsilon\delta')$. As a result, there are $2/\varepsilon$ values of δ within the interval $[-\delta', \delta']$ such that $T(x_i + \delta)$ is a track of the parabolic grid. For each of these values, we create a corresponding copy of the graph G to be mapped to the tracks $T(x_1 + \delta), \dots, T(x_k + \delta)$. By assumption, each copy cannot be embedded on a triangle that has an intersection of area less than $\varepsilon/2$ with the closed corner. If we connect all of the copies by additional edges forming a cycle, it becomes impossible to embed all the copies simultaneously such that they are pairwise contained in an internal face of another. Therefore, the total area of the intersection of the triangles containing the copies with the closed corner is more than one, which is impossible since the closed corner itself has area less than one. ■

Note that for any fixed n, we can choose ε as small as desired.

We prove now that parabolic grids do not support all planar graphs by showing that a particular graph cannot be embedded on the grid without containing a fixed positive area inside the closed corner.

Theorem 3. *Parabolic grids do not support all planar graphs.*

Proof: Let G be a planar triangulation on 24 vertices that has four vertex disjoint copies of G_6 as its subgraphs. The four copies of G_6 are mapped to four sets of consecutive tracks S_1, S_2, S_3 and S_4 on the parabolic grid. The cores $C(S_i)$ of the closed corner are represented on Figure 4. Note that the sets $C(S_i)$ are disjoint and separated by a positive distance.

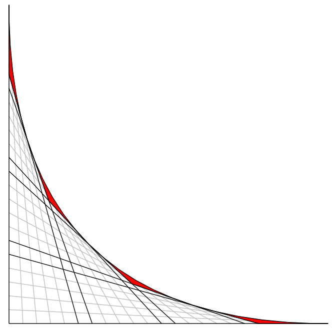

Fig. 4. It is impossible for a triangle to intersect all four of the shaded regions without intersecting the closed corner in at least some fixed positive area.

Lemma 5 implies that the ith copy of G_6 in G cannot be embedded on its subset S_i of the parabolic grid without intersecting $C(S_i)$. Thus, any embedding of G must have as external face a triangle that intersects each $C(S_i)$. However, it is clearly impossible to have any triangle intersecting all of them without intersecting the closed corner in at least some fixed positive area. Lemma 6 thus implies that the parabolic grid does not support all planar graphs. ■