2012

Geometric Graph Theory and Wireless Sensor Networks

Deniz Sarioz
Graduate Center, City University of New York

GEOMETRIC GRAPH THEORY AND WIRELESS SENSOR NETWORKS

by

DENIZ SARIOZ

A dissertation submitted to the Graduate Faculty in Computer Science in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2012

This manuscript has been read and accepted for the Graduate Faculty in Computer Science in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

János Pach

Date — Chair of Examining Committee

Theodore Brown

Date — Executive Officer

Boris Aronov

Amotz Bar-Noy

Bilal Khan

Supervision Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

GEOMETRIC GRAPH THEORY AND WIRELESS SENSOR NETWORKS

by

Deniz Sarioz

Advisor: Distinguished Professor János Pach

In this work, we apply geometric and combinatorial methods to explore a variety of problems motivated by wireless sensor networks. Imagine sensors capable of communicating along straight lines except through obstacles like buildings or barriers, such that the communication network topology of the sensors is their visibility graph. Using a standard distributed algorithm, the sensors can build common knowledge of their network topology.

We first study the following inverse visibility problem: What positions of sensors and obstacles define the computed visibility graph, with fewest obstacles? This is the problem of finding a minimum obstacle representation of a graph. This minimum number is the obstacle number of the graph. Using tools from extremal graph theory and discrete geometry, we obtain for every constant h that the number of n-vertex graphs that admit representations with h obstacles is $2^{o(n^2)}$. We improve this bound to show that graphs requiring $\Omega(n/\log^2 n)$ obstacles exist.

We also study restrictions to convex obstacles, and to obstacles that are line segments. For example, we show that every outerplanar graph admits a representation with five convex obstacles, and that allowing obstacles to intersect sometimes decreases their required number.

Finally, we study the corresponding problem for sensors equipped with GPS. Positional information allows sensors to establish common knowledge of their communication network geometry, hence we wish to compute a minimum obstacle representation of a given straight-line graph drawing. We prove that this problem is NP-complete, and provide a $O(\log OPT)$-factor approximation algorithm by showing that the corresponding hypergraph family has bounded Vapnik-Chervonenkis dimension.

Acknowledgements

First and foremost, I would like to thank my advisor János Pach. His deep love for his subject matter and perseverance in problem solving are infectious. He taught me in and out of the classroom, introduced me to lots of interesting research directions, and spent a lot of time reading and commenting on my manuscripts. Before I acknowledge them for other reasons, I would like to thank each member of my committee for being on my committee. I would like to thank Ted Brown, Amotz Bar-Noy, Gábor T. Herman, and T. Yung Kong—I learned a lot while working on projects under their supervision. János Pach, Ted Brown, Gábor T. Herman, Peter Brass, and the CUNY Doctoral Students' Council have provided me with financial support and travel grants. I would like to thank Boris Aronov for serving as the extramural member of my committee and for his detailed feedback on drafts of this work. I would like to thank Radoslav Fulek, Matthew P. Johnson, Filip Morić, Padmini Mukkamala, and Noushin Saeedi for many useful discussions and collaboration. I would like to thank Alon Efrat, Bilal Khan, Stephen Kobourov, and Peter Brass for discussing problems and offering guidance. I would like to acknowledge Robert M. Haralick and Stathis Zachos for helping me to become a more well-rounded student of computer science through their courses and seminars. I would like to thank Ted Brown, Joe Driscoll, and Lina Garcia for their support and help related to matters of CUNY and the Ph. D. Program in Computer Science. Last but not least, I am grateful to Rose Weixel for keeping me company while writing this dissertation, and for the better part of the last five years. All errors, omissions, and idiosyncrasies to follow are mine.

Table of Contents

List of Figures

Chapter 1

Introduction and Motivation

The purpose of this work is to apply geometric and combinatorial methods to explore a family of problems motivated by wireless sensor networks. Imagine sensors capable of communicating along straight lines except through obstacles like buildings or barriers, and close enough to each other to communicate subject to this visibility constraint. That is, the communication network topology of the sensors corresponds to their visibility graph in the presence of obstacles. Using a standard distributed algorithm, the sensors can attain common knowledge of their network topology.

For purposes of environmental modeling in the absence of additional information, it may be desirable for the sensors to automatically conjecture the obstacle shapes and locations, as well as the positions of the sensors, purely based on the network topology. More succinctly, we would like to infer a plausible obstacle representation of the network topology. An *obstacle representation* [4] of an abstract graph G is a point set in the plane together with a set of obstacles, such that the visibility graph on the point set is isomorphic to G. Without loss of generality, every obstacle is a simple polygon, and the graph vertices together with the obstacles' vertices are in general position, that is, no three are collinear.

Since there are many obstacle representations for any given abstract graph, simpler explanations ought to be preferred over more complicated ones. Variations on what ‘simple’ means as well as additional assumptions beget a family of inverse problems.

We first study the following inverse visibility problem: What positions of sensors and obstacles define the computed visibility graph, with fewest obstacles? From the point of view of application, this question corresponds to the “simplest means fewest” approach, giving equal weight to a triangular obstacle and an obstacle with a million sides. This is the problem of finding a minimum cardinality obstacle representation of a graph. This graph parameter is the *obstacle number* of the graph. Using tools from extremal graph theory and discrete geometry, we obtain the following result.

Theorem 2.1.2. *For every constant h, the number of n-vertex graphs with obstacle number at most h is* $2^{o(n^2)}$.

As we will show, this is sufficient to conclude that the obstacle numbers for many families, such as bipartite graphs, is unbounded. This result has been published in Graphs and Combinatorics [36]. Using facts regarding order types and line segment arrangements, we improve this bound to prove the following:

Theorem 2.4.2. *There are graphs on n vertices with obstacle number at least* $\Omega\left(n/\log^2 n\right)$.

Prior information on the nature of the obstacles can be very useful. Many environments are likely to have artificial structures that are convex. Hence, we also study restriction of the above problem to convex obstacles. An obstacle representation of a graph G with only convex obstacles is called a *convex obstacle representation* of G, and the fewest number of obstacles over all convex obstacle representations of G is called the *convex obstacle number* of G.

Theorem 3.1.1. *There are graphs on n vertices with convex obstacle number at least* $\Omega(n/\log n)$.

Another meaningful measure of complexity for an obstacle representation with polygonal obstacles is the total number of vertices of the obstacles. Hence, a representation that minimizes this may be preferred. Notice that this minimum is of the same order of magnitude as the minimum number of obstacles each of which is a straight line segment. An obstacle representation of a graph G in which every obstacle is a straight line segment is called a *segment obstacle representation* of G, and the fewest number of obstacles over all segment obstacle representations of G is called the *segment obstacle number* of G. For this we have:

Theorem 3.1.2. *There are graphs on n vertices with segment obstacle number at least* $\Omega(n^2/\log n)$.

These results were presented at the 36th International Workshop on Graph Theoretic Concepts in Computer Science, whose proceedings were published in the Lecture Notes in Computer Science series of Springer [32].

Alpert, Koch, and Laison [4] showed that every outerplanar graph has obstacle number at most one. They queried whether the convex obstacle number is a bounded parameter for outerplanar graphs. We give a construction to show:

Theorem 3.2.1. *Every outerplanar graph has convex obstacle number at most* five.

Using a similar technique, we prove an upper bound of *four* for the convex obstacle numbers of bipartite permutation graphs. The results to be presented in this dissertation regarding outerplanar graphs and bipartite permutation graphs have been accepted to appear in a volume on Geometric Graph Theory [19].

We also show that allowing segment obstacles to intersect sometimes decreases their required number, as we formulate exact dependencies on the number of vertices for certain families of graphs.

Finally, we study the corresponding problem for sensors equipped with GPS. Positional information allows sensors to establish common knowledge of the geometry of their communication network in addition to its topology. In mathematical terms, we wish to compute a minimum obstacle representation of a given straight-line graph drawing.

We show that this problem can be formulated as a hypergraph transversal problem of size polynomial in the number of graph vertices, establishing that it is in NP. We prove that this problem is NP-hard by giving a reduction from planar vertex cover. We provide a $O(\log OPT)$-factor approximation algorithm by bounding the Vapnik-Chervonenkis dimension for the corresponding hypergraph family. This positive result was presented at the 23rd Canadian Conference on Computational Geometry [42]. An interesting special case is that of the family of plane graphs—drawings of planar graphs in the plane without edge crossings. Even though the problem is NP-hard already for this special case, we show the following positive algorithmic results by a reduction to maximum degree 3 planar vertex cover.

Corollary 4.4.2. *There is a polynomial-time approximation scheme (PTAS) for computing the obstacle number of a plane graph.*

Corollary 4.4.3. *Computing the obstacle number of a plane graph is fixed parameter tractable (FPT).*

These reductions and their complexity-theoretic consequences were first given in the e-print [28].

Recurrent notations and terminology

All of our graphs are finite and simple. A graph G consists of a finite set of elements called the vertices of G and denoted by $V(G)$, together with a specified set of vertex pairs called the edges of G and denoted by $E(G)$. A graph *on* V has vertex set V. For vertices u and v of a graph G, we refer to the pair $\{u,v\}$ as uv. We denote by $[n]$ the set $\{1,2,3,\ldots,n\}$. A *labeled graph* G has vertex set $[n]$, and its edges are specified pairs in $[n]$. An *unlabeled graph* on n vertices is an equivalence class of labeled graphs on n vertices closed under graph isomorphism. A *hypergraph*, also known as a set system, is the following generalization of a graph. A hypergraph $\mathscr{H}$ consists of a set of elements called the vertices of $\mathscr{H}$, together with a collection of specified subsets of its vertex set called the *hyperedges* of $\mathscr{H}$.

While some concepts that we study may have reasonable projective analogues, we use the word *plane* to always indicate the Euclidean plane.

The *arrangement* $\mathscr{A}$ of a set S of line segments in the plane is the incidence structure among the vertices, edges, and cells of the arrangement, defined in the following way. A segment endpoint or an intersection point of two segments is a *vertex* of $\mathscr{A}$. An open interval on a segment defining the arrangement between two vertices of the arrangement that contains no vertex of the arrangement is an *edge* of $\mathscr{A}$. A connected component of the complement of the union of the segments in S is a *cell* of $\mathscr{A}$. Arrangements of line segments are readily extended to arrangements of lines. The *graph of* $\mathscr{A}$ is the graph on the vertices of $\mathscr{A}$ with pairs defined by the edges of $\mathscr{A}$.

These notations and terms are in keeping with community standards, e.g., [30].

Chapter 2

Obstacle numbers of graphs: Bounds and ad-hoc computations

2.1 Obstacle representations of graphs and the obstacle number of a graph

Consider a set P of points in the plane and a set of closed polygonal obstacles whose vertices together with the points in P are in *general position*, that is, no *three* of them are on a line. The corresponding *visibility graph* has P as its vertex set, two points $p, q \in P$ being connected by an edge if and only if the segment pq does not meet any of the obstacles. Visibility graphs are extensively studied and used in computational geometry, robot motion planning, and sensor networks; see [10], [22], [33], [34], [45].

Alpert, Koch, and Laison [4] introduced an interesting new parameter of graphs, closely related to visibility graphs. Given a graph G, we say that a set of points and a set of polygonal obstacles as above constitute an *obstacle representation* of G if the corresponding visibility graph is isomorphic to G. A representation with h obstacles is called an *h-obstacle representation*. The smallest number of obstacles over all obstacle representations of G is called the *obstacle number* of G, which we denote by $obs(G)$.

Given any placement (embedding) of the vertices of G in general position in the plane, a straight-line drawing of G consists of the image of the embedding and the set of open straight line segments connecting all pairs of points that correspond to the edges of G. We refer to a pair of vertices in G that does not define an edge in G as a *non-edge* of G. If there is no danger of confusion, we make no notational difference between the vertices of G and the corresponding points, or between vertex pairs and corresponding open segments. The complement of the set of all points that correspond to a vertex or belong to at least one edge of G falls into connected components. Each of these components is called a *face* of the drawing. In every drawing of a graph, there is a unique unbounded face, referred to as the *outside face*. A 1-obstacle representation in which the obstacle lies on the outside face is called an *outside obstacle representation*, and such an obstacle is called an *outside obstacle*. Notice that if G has an obstacle representation with a particular placement of its vertex set, then

1. each obstacle must lie entirely in one face of the corresponding drawing of G, and

2. each non-edge of G must be blocked by at least one of the obstacles.

Therefore, the problem of finding the minimum number of obstacles required for a given graph drawing can be reformulated as: What is the smallest number of faces that together block all non-edges?

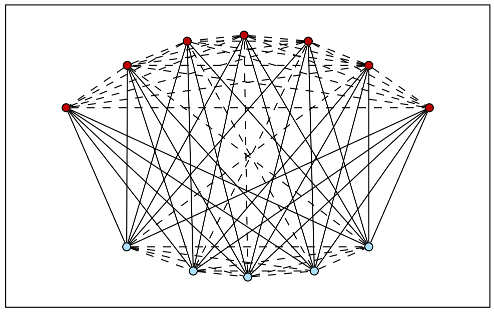

Figure 2.1 – Drawing of G_1 that can be completed to a 2-obstacle representation

Alpert, Koch, and Laison [4] showed that any representation of the bipartite graph G_1 which can be obtained by removing a maximum matching from a complete bipartite graph $K_{5,7}$, requires at least *two* obstacles. See Fig. 2.1. They also constructed a *split graph* G_2, i.e., a graph that splits into a complete subgraph and an independent set, with a number of edges running between them, which has obstacle number at least *two*. See Fig. 2.2. We

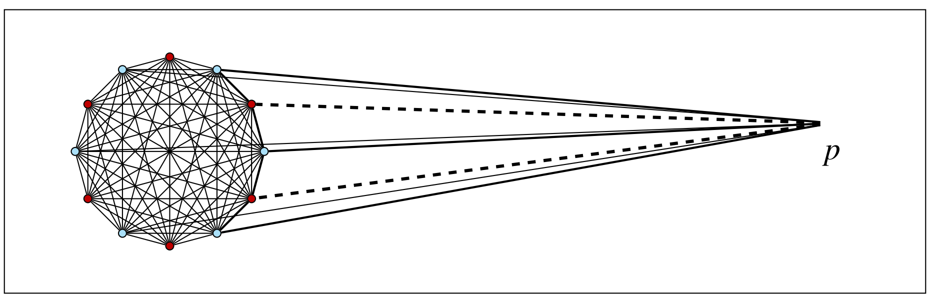

Figure 2.2 – $V(G_2)$ is the union of a clique A of 92379 vertices, and an independent set I of $\binom{92379}{6}$ vertices of degree 6 with distinct neighborhoods. Out of every 92379 points in general position, at least 12 are in convex position. For some drawing of G_2, we show the drawing induced on such 12 vertices comprising A' and a vertex $p \in I$ with edges to 6 vertices in A' that alternate around conv(A'). In every drawing of G_2, every such choice of A' and p implies the presence of at least two interior-disjoint solid quadrilaterals with non-edges inside each.

augment the above examples with a third one: a graph G_3 with obstacle number at least *two*, whose complement is a bipartite graph.

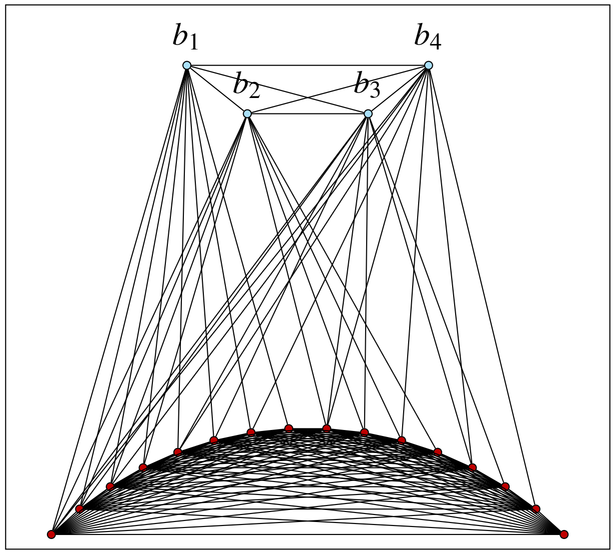

Figure 2.3 – Drawing of G_3, the unique 20-vertex graph whose vertex set consists of a clique of *four* (light blue) vertices and a clique of *sixteen* (dark red) vertices such that every red vertex is adjacent to a distinct set of blue vertices

Lemma 2.1.1. *The graph G_3 specified in Fig. 2.3, which consists of two complete subgraphs with many edges between them, has obstacle number at least* two*.*

We will not prove this lemma here. Instead, in Section 2.3, we shall introduce smaller graphs G'_1, G'_2, and G'_3 which are respectively bipartite, split, and bipartite complement graphs, and show that each of them has obstacle number at least *two*.

Alpert, Koch, and Laison applied the Erdős-Szekeres convex n-gon theorem [16] to generalize their construction of G_2 to produce a sequence of graphs with arbitrarily large obstacle numbers. Here we demonstrate that the existence of such graphs is a simple consequence of the fact that no graph of obstacle number *one* contains a subgraph isomorphic to G_1, G_2, or G_3. In Section 2.2, we will show that this set of forbidden graphs allows us to utilize some extremal-graph-theoretic tools. They yield that the number of graphs with n vertices and bounded obstacle number is very small, compared to the total number of labeled graphs, which is $2^{\binom{n}{2}}$. More precisely, we obtain

Theorem 2.1.2. *For any fixed positive integer h, the number of graphs on n (labeled) vertices with obstacle number at most h is at most* $2^{o(n^2)}$.

One of the unsolved questions left open in [4] was whether there exist bipartite graphs with arbitrarily large obstacle number. Since total number of labeled bipartite graphs with n vertices is at least $2^{n^2/4}$, Theorem 2.1.2 immediately implies that the answer to the above question is in the affirmative. We also give a construction from scratch.

Corollary 2.1.3. *For any positive integer h, there exists a bipartite graph with obstacle number larger than h.*

2.2 Hereditary properties—Proof of Theorem 2.1.2

In 1985, Erdős, Kleitman, and Rothschild [15] proved that, as n tends to infinity, the number of all K_ℓ-free graphs on n vertices is asymptotically equal to the number of $(\ell-1)$-partite graphs with n vertices with as equal vertex classes as possible. This result was soon generalized to graphs that do not contain some fixed (not necessarily induced) subgraph H [14].

Analogous questions based on the *induced* subgraph relation were investigated in [38], [40], and [39]. If a graph G does not contain an induced subgraph isomorphic to a fixed graph H, then the same is true for every induced subgraph of G. Therefore, this property is *hereditary*. As we explain later, a hereditary graph property can be characterized by its set of forbidden induced subgraphs. In order to formulate an Erdős-Kleitman-Rothschild type theorem valid for any hereditary graph property, we need some definitions and notations.

We identify a graph property $\mathscr{P}$ with the set of all graphs that satisfy this property. In the same spirit, we denote by $\mathscr{P}^n$ the set of all graphs on n labeled vertices which satisfy property $\mathscr{P}$. From now on, we consider only properties that hold for infinitely many graphs.

A graph is (r,s)-colorable if its vertex set can be partitioned into r blocks, out of which s are cliques and every remaining block is an independent set. Let $\mathscr{C}(r,s)$ denote the set of all (r,s)-colorable graphs. A graph property which holds for all graphs is called *trivial*. Given any nontrivial hereditary graph property $\mathscr{P}$, its *coloring number* is defined as

$$r(\mathscr{P}) = \max\{r \mid \exists s : \mathscr{C}(r,s) \subseteq \mathscr{P}\}.$$

We now argue that the parameter $r(\mathscr{P})$ exists and is at least *one*. By Ramsey's Theorem, every graph on $n > R(k,\ell)$ vertices has either a clique on k vertices or an independent set on ℓ vertices. Hence, if a property were to exclude both a complete graph and a clique, it would hold for finitely many graphs, contrary to hypothesis. In other words, it must be the case that $\mathscr{C}(1,0) \subseteq \mathscr{P}$ or $\mathscr{C}(1,1) \subseteq \mathscr{P}$, hence $r(\mathscr{P}) \geq 1$. Since $r(\mathscr{P})$ is strictly less than the number of vertices of any graph that does not belong to $\mathscr{P}$, it is also bounded from above.

Theorem 2.2.1 (Bollobás, Thomason [7])**.** *For any nontrivial hereditary graph property $\mathscr{P}$, the number of (labeled) graphs on n vertices with property $\mathscr{P}$ is*

$$|\mathscr{P}^n| = 2^{\left(1-\frac{1}{r(\mathscr{P})}+o(1)\right)\binom{n}{2}}.$$

Here, it does not matter whether we count labeled or unlabeled graphs, because the corresponding quantities differ only by a factor of at most $n! = 2^{O(n\log n)}$. If for some value r there is no s such that $\mathscr{C}(r,s) \subseteq \mathscr{P}$, then for every $r' > r$ there is no s for which $\mathscr{C}(r',s) \subseteq \mathscr{P}$. If we can find $(2,0)$-colorable, $(2,1)$-colorable, and $(2,2)$-colorable graphs, *none* of which has property $\mathscr{P}$, then, by the preceding observations, $r(\mathscr{P}) = 1$. Thus, by

Theorem 2.2.1, we can conclude that the number of graphs on n vertices with property $\mathscr{P}$ is $2^{o(n^2)}$.

The familiar term for a $(2,0)$-colorable graph is bipartite. A $(2,1)$-colorable graph consists of a clique and an independent set, possibly with edges running between them; such a graph is often called a *split graph* [18], [44]. A $(2,2)$-colorable graph consists of two cliques, possibly with edges running between them—its complement is bipartite.

Apply Theorem 2.2.1 to the hereditary property that a graph admits a 1-obstacle representation. The graphs G_1, G_2, and G_3 introduced in Section 2.1 are $(2,0)$-, $(2,1)$- and $(2,2)$-colorable. Thus, in view of the fact that, according to Alpert et al. and Lemma 2.1.1, none of them admits a 1-obstacle representation, we can conclude that the number of all graphs on n (labeled) vertices with obstacle number at most 1 is $2^{o(n^2)}$. In other words, Theorem 2.1.2 holds for $h=1$.

Denote the set of the first n positive integers by $[n]$. Given $h>1$, consider a graph G on the vertex set $[n]$ with obstacle number at most h, and fix an obstacle representation R for it with h obstacles $O_1,O_2,\ldots,O_h$. As usual, we do not distinguish between $V(G)$ and the point set corresponding to it in R. For each $i\in[h]$, let G_i be the visibility graph on $V(G)$ determined only by the obstacle O_i. It is easy to see that G is a subgraph of G_i, since O_i by itself blocks no more visibilities among $V(G)$ than do all h obstacles combined. In other words, $E(G)\subseteq\cap_{i\in[h]}E(G_i)$. In fact, we have that $E(G)=\cap_{i\in[h]}E(G_i)$, since for every edge $uv\in E(G)$, the segment uv avoids all obstacles specified in R. Let us denote by $\mathscr{G}_h^n$ the set of labeled graphs on $[n]$ with obstacle numbers at most h. Since every $G\in\mathscr{G}_h^n$ is uniquely determined by the above graphs $G_1,G_2,\ldots,G_h\in\mathscr{G}_1^n$, we have $\left|\mathscr{G}_h^n\right|\leq\left|\mathscr{G}_1^n\right|^h$. Using the fact that $\left|\mathscr{G}_1^n\right|=2^{o(n^2)}$, we can conclude that $\left|\mathscr{G}_h^n\right|=2^{o(n^2)}$ for any fixed h.

This completes the proof of Theorem 2.1.2.

2.3 Small $(2,s)$-colorable graphs without 1-obstacle representations

We show that a particular 10-vertex $(2,0)$-colorable (i.e., bipartite) graph G_1' has obstacle number greater than *one*. This improves upon the 12-vertex bipartite graph G_1 in [4], and settles a conjecture therein in the affirmative. We also show that a particular 70-vertex $(2,1)$-colorable (i.e., split) graph G_2' has obstacle number greater than *one*, improving on the $\left(92379+\binom{92379}{6}\right)$-vertex graph implied by a construction in [4]. We finally show that a $(2,2)$-colorable 10-vertex graph G_3' has obstacle number greater than *one*, improving on our own 20-vertex graph G_3. The graphs G_1, G_2, and G_3 were depicted in Section 2.1.

2.3.1 A 10-vertex bipartite graph

In [4], $K^*_{m,n}$s has been defined as the graph obtained from the complete bipartite graph $K_{m,n}$ by removing a maximum matching. There, it was shown that every $K^*_{m,n}$ graph admits a 2-obstacle representation: The two independent sets are placed within disjoint half-planes, such that the non-edges in the removed matching meet at a single point so that a single non-outside obstacle is sufficient to meet them, while the non-edges within the independent sets meet the outside face so that an outside obstacle is sufficient to meet them. The authors also gave a strong hint for obtaining an outside obstacle representation of $K^*_{4,n}$ for every n by providing an easily generalizable outside obstacle representation for $K^*_{4,5}$. Furthermore, they proved that $G_1 := K^*_{5,7}$ does not admit a 1-obstacle representation. We dedicate the rest of this section to proving the following conjecture of theirs.

Theorem 2.3.1. *$G_1' := K^*_{5,5}$, the graph obtained from $K_{5,5}$ by removing a perfect matching, has obstacle number 2.*

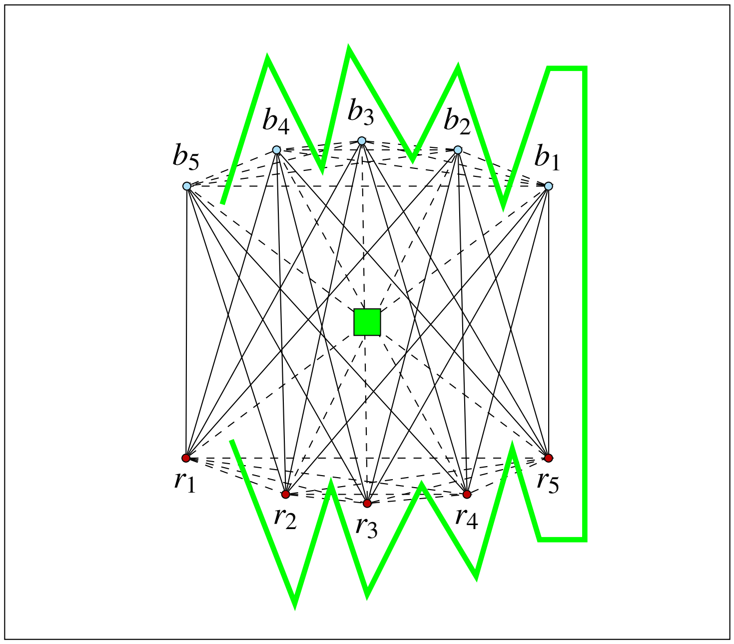

Figure 2.4 – A 2-obstacle representation of the bipartite graph G_1', i.e., $K^*_{5,5}$

Proof. To able to refer to individual vertices of $K^*_{5,5}$, let $V(K^*_{5,5}) = B \uplus R$ such that $B = \{b_1, b_2, b_3, b_4, b_5\}$ (the set of light blue vertices) and $R = \{r_1, r_2, r_3, r_4, r_5\}$ (the set of dark red vertices) are independent sets and there is an edge from a blue vertex b_i to a red vertex r_j if and only if $i \neq j$. See Fig. 2.4 for a 2-obstacle representation of G_1'.

Before we proceed, we borrow some definitions and two facts from [4]. Given points a, b, c in the plane we say *a sees b to the left of c* (equivalently, *sees c to the right of b*) if the points a, b, and c appear in clockwise order. If a point a is outside the convex hull of some set S of points, the relation "a sees to the left of" is transitive on S, hence is a total ordering of S, called the *a-sight ordering* of S.

We paraphrase Lemma 3 of [4] in the following way.

Lemma 2.3.2. *If a graph having $K_{2,3}$ as an induced subgraph has a 1-obstacle representation, then in such a representation the two parts (independent sets) of the induced $K_{2,3}$ are* *linearly separable. Moreover, for each part S, every vertex in the other part induces the same sight ordering of S.*

Lastly, we paraphrase a fact used in the original proof of Lemma 2.3.2. We denote by $\text{conv}(P)$ for the convex hull of a point set P.

Lemma 2.3.3. *In a 1-obstacle representation of $K^*_{5,5}$, every vertex subset S consisting of 2 red vertices and 2 blue vertices with 4 distinct subscripts (the necessary and sufficient condition for a $K_{2,2}$ to be induced) is in convex position, with both color classes appearing contiguously around* $\text{conv}(S)$. *Hence the drawing induced on S (i.e., the drawing of every induced $K_{2,2}$ in $K^*_{5,5}$) is self-intersecting, a* bowtie.

We now state and prove a new lemma, one of many to help prune the space of vertex arrangements potentially amenable to 1-obstacle representations of $K^*_{5,5}$.

For any three points p,q,r, we denote by $\angle pqr$ the union of the rays $\overrightarrow{qp}$ and $\overrightarrow{qr}$.

Lemma 2.3.4. *Every 1-obstacle representation of $K^*_{5,5}$ is an outside obstacle representation.*

Proof. Assume that we are given a 1-obstacle representation of $K^*_{5,5}$ that is not an outside obstacle representation. At least three vertices are on the convex hull boundary of the vertices by the general position assumption. Every pair of vertices appearing consecutively around the convex hull boundary must constitute an edge, otherwise an outside obstacle would be required to block it. Then without loss of generality b_1, r_2, b_3 appear consecutively on the bounding polygon. All other vertices including r_4 are inside $\text{conv}(\angle b_1 r_2 b_3)$. Hence the drawing of the $K_{2,2}$ induced on $\{b_1, r_2, b_3, r_4\}$ is not a bowtie, which contradicts Lemma 2.3.3. □

Lemma 2.3.5. *In every 1-obstacle representation of $K^*_{5,5}$, every vertex v is linearly separable from the set S of its neighbors, defining a v-sight ordering on S.*

Proof. Assume for contradiction that we are given a 1-obstacle representation of $K_{5,5}$ in which some vertex, without loss of generality, b_1, is not linearly separable from the set of its neighbors. Then b_1 is in the convex hull of $\{r_2, r_3, r_4, r_5\}$. By the general position assumption, a triangulation of $\{r_2, r_3, r_4, r_5\}$ will reveal that b_1 is inside some triangle with red vertices, Without loss of generality, $\Delta r_3 r_4 r_5$. Then by the general position assumption, the ray $\overrightarrow{b_1 b_2}$ meets an interior point of some edge of this triangle, Without loss of generality, $\overline{r_4 r_5}$. This implies that the drawing of $K_{2,2}$ induced on $\{b_1, r_4, b_2, r_5\}$ is not a bowtie, which contradicts Lemma 2.3.3. □

In a graph drawing or obstacle representation, we say that a polygon (by which we mean simple closed polygonal curve) is *solid* if it is a subset of the drawing: if every point on it is a vertex or on an edge.

Lemma 2.3.6. *In every 1-obstacle representation of $K^*_{5,5}$, every vertex in R (respectively, B) is linearly separable from B (respectively, R).*

Proof. We will show that in every 1-obstacle representation of $K^*_{5,5}$, each blue vertex is linearly separable from R. The analogous statement about each red vertex and B can be proved symmetrically.

Assume for contradiction that we are given a 1-obstacle representation of $K^*_{5,5}$ in which some blue vertex is in $\mathrm{conv}(R)$. Without loss of generality, $b_1 \in \mathrm{conv}(R)$. By Lemma 2.3.5, b_1 is linearly separable from $\{r_2, r_3, r_4, r_5\}$. Without loss of generality, $\overleftrightarrow{r_2 r_5}$ is a horizontal line with r_2 to the left of r_5 such that b_1 is above $\overleftrightarrow{r_2 r_5}$ and $r_3 r_4$ is inside $\mathrm{conv}(\angle r_2 b_1 r_5)$. Call the four open regions delineated by the lines $\overleftrightarrow{r_2 b_1}$ and $\overleftrightarrow{r_5 b_1}$ the left, right, upper, and lower

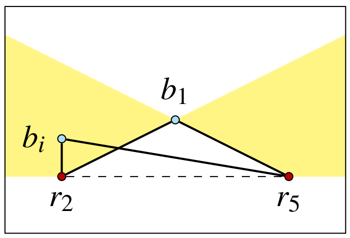

(a) For every $i \in \{3,4\}$, b_i must be in the shaded region.

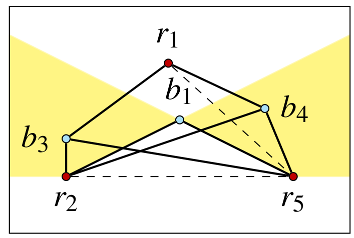

(b) If b_3 and b_4 are on opposite quadrants of b_1 as shown, the non-edge r_1r_5 cannot be blocked by the outside face.

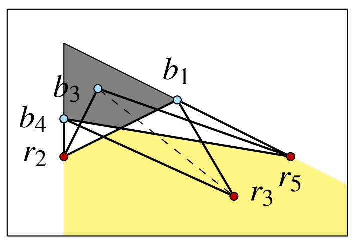

(c) The "possibility regions" of b_3 and r_3 are shown in different hues. Even if both regions are unbounded, the non-edge r_3b_3 cannot be blocked by the outside face.

Figure 2.5 – Subfigures (a), (b), and (c) respectively accompany the second, third, and last paragraphs in the proof of Lemma 2.3.6. Some edges and non-edges are omitted for clarity, as they often will be in subsequent figures.

b_1-quadrants. For each $i \in \{3,4\}$, since the drawing of $K_{2,2}$ induced on $\{r_2, b_i, r_5, b_1\}$ must be a bowtie by Lemma 2.3.3, b_i is above $\overleftrightarrow{r_2 r_5}$ and in either the left or the right b_1-quadrant. (See Fig. 2.5(a).)

Without loss of generality, b_3 is in the left b_1-quadrant. Assume for contradiction that b_4 is in the right b_1-quadrant. Since $b_1 \in \text{conv}(R)$, r_1 is in the upper b_1-quadrant. Then b_3, r_1, b_4 are respectively in the left, upper, and right b_1-quadrants, and r_5 is on the boundary of the right and lower b_1-quadrants. This implies that the drawing of $K_{2,2}$ induced on $\{b_3, r_1, b_4, r_5\}$ is non-self-intersecting, not a bowtie. By Lemma 2.3.3, this means b_4 is in the left b_1-quadrant along with b_3. (See Fig. 2.5(b).)

Notice that $K_{2,3}$ is induced on $\{r_2, r_5, b_1, b_3, b_4\}$. Then by Lemma 2.3.2, b_3 and b_4 are above $\overleftrightarrow{r_2 r_5}$ along with b_1, and the r_2- and r_5-sight orderings of $\{b_1, b_3, b_4\}$ are the same, with b_1 appearing rightmost. Without loss of generality, r_2 and r_5 see b_4 to the left of b_3. Hence, b_3 is inside $\text{conv}(\angle b_4 r_2 b_1)$ in addition to being inside $\text{conv}(\angle b_4 r_5 b_1)$. By the same token, since b_1 sees r_3 to be between r_2 and r_5, so does b_4. Hence, r_3 is inside $\text{conv}(\angle r_2 b_4 r_5)$, in addition to being inside $\text{conv}(\angle r_2 b_1 r_5)$. These conditions ensure that $\text{conv}(\angle r_2 b_3 r_5)$ and $\text{conv}(\angle b_4 r_3 b_1)$ meet to give a convex quadrilateral region with solid boundary that has $\overline{b_3 r_3}$ as a diagonal. This implies that the non-edge $b_3 r_3$ is not blocked by the outside face, in contradiction to Lemma 2.3.4. (See Fig. 2.5(c).) □

Denote by $K_{3,3}^-$ the graph obtained by removing an edge from $K_{3,3}$. Note that our proof of Lemma 2.3.6 relies on showing that the assumptions lead to a drawing of $K_{2,2}$ forbidden by Lemma 2.3.3, or to a forbidden drawing of $K_{3,3}^-$ like the one shown in Fig. 2.5(c).

Lemma 2.3.7. *In every 1-obstacle representation of $K_{5,5}^*$, the convex hulls of R and B are disjoint, hence, there is a line separating R from B.*

Proof. Assume for contradiction that we are given a 1-obstacle representation of $K_{5,5}^*$ in which $\text{conv}(R) \cap \text{conv}(B) \neq \varnothing$. Let X denote $\text{conv}(R) \cap \text{conv}(B)$. But by Lemma 2.3.6, $(R \cup B) \cap X = \varnothing$. This means that X is a $2k$-gonal shape ($2 \leq k \leq 5$) separating $\text{conv}(B)$ and $\text{conv}(R)$ into k pieces each, alternating around it.

If $k \geq 3$, without loss of generality, $r_1 b_{i_2} r_2 b_{i_1} r_3 b_{i_3}$ is a counterclockwise enumeration of some convex hexagon H. Take $\overleftrightarrow{r_2 r_3}$ as horizontal. Without loss of generality, b_4 is below $\overleftrightarrow{r_2 r_3}$. By Lemma 2.3.2, b_1 and b_5 are also below $\overleftrightarrow{r_2 r_3}$. This means that $\{b_{i_2}, b_{i_3}\} = \{b_2, b_3\}$. If $i_2 = 3$, then H is solid and has the non-edge $b_2 b_3$ as an internal diagonal, which therefore requires an internal obstacle, contradicting Lemma 2.3.4. Otherwise, $i_2 = 2$ and $\overline{r_2 b_3}$ meets $\overline{b_2 r_3}$ at some point q, so the solid convex quadrilateral $r_1 b_2 q b_3$ has $b_2 b_3$ as an internal diagonal, which once again requires an internal obstacle, contradicting Lemma 2.3.4.

Therefore, X separates $\text{conv}(R)$ and $\text{conv}(B)$ into 2 pieces each. Denote by R_1 and R_2 the subsets of R induced by this partition, and define B_1 and B_2 similarly. Without loss of generality, $|R_1| \in \{1, 2\}$ and $|B_1| \in \{1, 2\}$. Now we will show that $|R_1| = |B_1| = 1$.

Assume otherwise for contradiction. Without loss of generality, $R_1 = \{r_1, r_2\}$ and $R_2 = \{r_3, r_4, r_5\}$. By Lemma 2.3.2, $\overline{r_1 r_3}$ is linearly separable from $\Delta b_2 b_4 b_5$. The line $\overleftrightarrow{r_1 r_3}$ separates B_1 from B_2. This implies that $\{b_2, b_4, b_5\} \subseteq B_2$. Similarly, $\overline{r_2 r_4}$ is linearly separable from $\Delta b_1 b_3 b_5$, which implies $\{b_1, b_3, b_5\} \subseteq B_2$. But then we have $|B_2| = 5$, a contradiction.

Without loss of generality, let $R_1 = \{r_1\}$. To see that this forces $B_1 = \{b_1\}$, assume for contradiction that (without loss of generality) $B_1 = \{b_2\}$. Then $\overline{r_1 r_3}$ meets $\overline{b_2 b_4}$, contradicting Lemma 2.3.3.

Without loss of generality, the sets R_1, B_1, R_2, B_2 appear clockwise around X, in this order. Notice that every red vertex in R_2 sees b_1 rightmost in B. Without loss of generality, let the b_1-sight ordering of R be r_5, r_4, r_3, r_2, r_1. To highlight the resemblance to the proof of Lemma 2.3.6, take the line $\overleftrightarrow{r_2r_5}$ to be horizontal with r_2 to the left of r_5.

Since $K_{2,3}$ is induced on $\{b_1,b_3,b_4,r_2,r_5\}$, by Lemma 2.3.2 the r_2- and r_5-sight orderings of $\{b_4,b_3,b_1\}$ are the same. Since r_2 and r_5 are in R_2, they see b_1 as the rightmost blue vertex and without loss of generality they see b_4 to the left of b_3. Thus we have exactly the same conditions as those used in the last paragraph of the proof of Lemma 2.3.6 to conclude that $\overline{b_3r_3}$ is an interior diagonal of a solid quadrilateral, hence an outside obstacle is insufficient in this case too.

Therefore, in a 1-obstacle representation of $K^*_{5,5}$, conv(B) and conv(R) are disjoint. □

Armed with the knowledge that every 1-obstacle representation of $K^*_{5,5}$ is an outside obstacle representation and requires R and B to be linearly separable, assume for contradiction that we are given a drawing of $K^*_{5,5}$ that admits a 1-obstacle representation. We will argue that such a drawing necessarily contains a drawing of $K_{2,2}$ requiring more than one obstacle or a drawing of $K^-_{3,3}$ requiring more than one obstacle. We justify the existence of such a forbidden configuration by using an algorithm that removes vertices from the drawing until inspecting the convex hull boundary of the vertices must reveal the existence of such a configuration.

Now we give some terminology needed to describe the algorithm. By Lemma 2.3.7 the convex hulls of B and R are disjoint. Without loss of generality, the x-axis separates B from R, with B above and R below. Let U be a subset of the vertex set that has at least three blue and at least three red vertices. Consider the clockwise cyclic order of the vertices in the bounding polygon of U. If u appears immediately before v in this order, we say that the ordered pair (u,v) is clockwise in U. The general position assumption and the linear separation of R from B imply that a unique ordered pair of the form (r_i,b_j) is clockwise in U, and a unique ordered pair of the form (b_k,r_ℓ) is clockwise in U. Refer to $\{r_i,b_j\}$ as the *left*

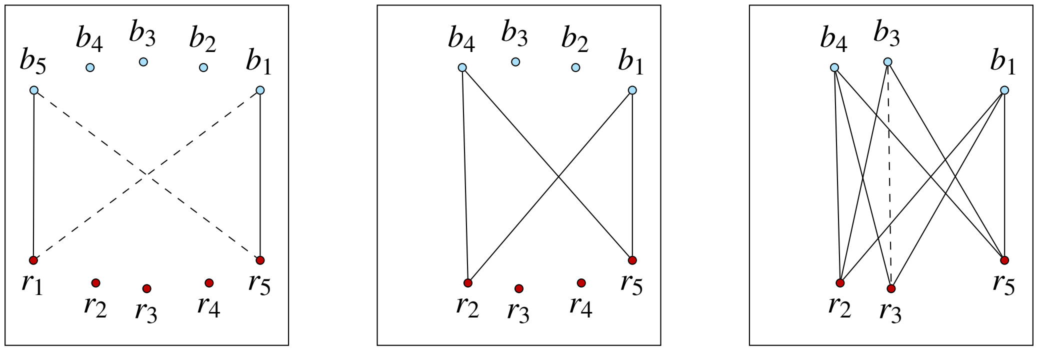

(a) No subscript in W is unique, so no internal vertex is removed. Both diagonals are then non-edges, so the left wall is removed.

(b) Halting condition $|W| = |I|$ attained with $|W| = 4$ and $\{r_3, b_3\} \subseteq U \setminus W$, which we can tell by inspecting W since $3 \in [5] \setminus I$ and I grows monotonically.

(c) The forbidden drawing of $K^-_{3,3}$ implied in (b) shown. The obstacle is outside, yet the non-edge $b_3 r_3$ shown here is not blocked by the outside face.

Figure 2.6 – A run of the algorithm in the proof of Theorem 2.3.1. Notice that the initial state features a placement of $V(K^*_{5,5})$ in which R is linearly separable from B and below B as required. In all but the last subfigure, only the dichromatic pairs induced on W are shown.

wall of U and denote it by $w_{\text{left}} = w_{\text{left}}(U)$. Similarly, refer to $\{b_k, r_\ell\}$ as the *right wall of U* and denote it by $w_{\text{right}} = w_{\text{right}}(U)$. Let $W = W(U) = \{r_i, b_j, b_k, r_\ell\}$, the *wall vertices of U*. The assumptions on U imply $3 \leq |W| \leq 4$. Denote by $I = I(U)$ the set of actual subscripts of the vertices in $W(U)$. Then $2 \leq |I| \leq 4$. Refer to a pair of vertices as *dichromatic* if they belong to different color classes and *monochromatic* otherwise. Observe that $|W| - |I|$ is the number of dichromatic non-edges of $K^*_{5,5}$ induced on W.

Here is the algorithm sketch. Initialize $U := V(K^*_{5,5})$. The halting condition is $|W| = |I|$, i.e., that every vertex in W has a distinct subscript. Repeat the following until the halting condition arises. For every $i \in [5]$, we say that r_i and b_i are twins. For every vertex with a unique subscript in W, remove its twin from U (unless it has already been removed). Remove at least one vertex in W with a twin also in W, the specifics to be described later.

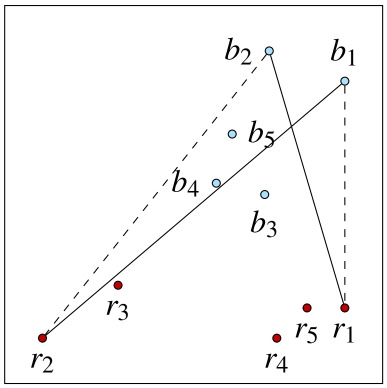

(a) $|W| = 4$ and w_{right} is a non-edge: Removing b_1 will not evict r_1 from the right wall, so b_1 is removed.

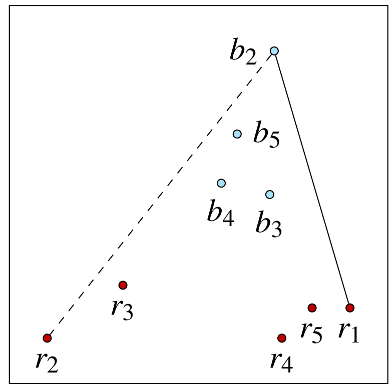

(b) $|W| = 3$ and w_{left} is the unique dichromatic non-edge induced on W. The vertex b_2 is kept since it is in both walls, while its twin r_2 is removed.

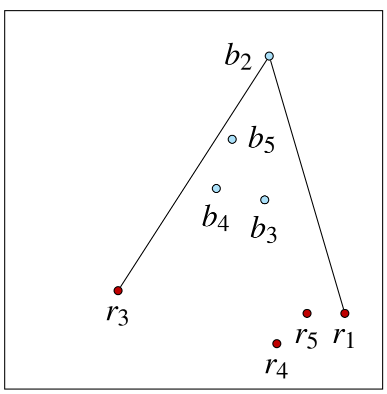

(c) Halting condition $|W| = |I|$ attained with $|W| = 3$ and $b_4 \in U \setminus W$, which we can tell by inspecting W since $4 \in [5] \setminus I$ and I grows monotonically.

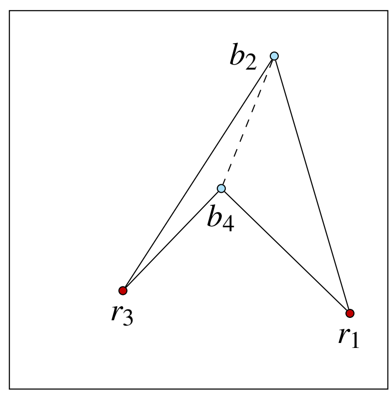

(d) The forbidden drawing of $K_{2,2}$ implied in (c) shown. The obstacle is outside, yet the non-edge $b_2 b_4$ is not blocked by the outside face.

Figure 2.7 – Another run of the algorithm in the proof of Lemma 2.3.1, illustrating configurations distinct from those shown in Fig. 2.6

A vertex v is removed from U only if its twin $\overline{v}$ is in W, and removing v will cause $\overline{v}$ to be locked in W for the rest of the algorithm execution, due to the careful way in which we remove a wall vertex. Assuming that this claim holds, I grows monotonically. This means $\{r_j, b_j\} \subseteq U \setminus W$ for every $j \in [5] \setminus I$. Let us call the vertices in $U \setminus W$ the interior vertices of U, and a pair $\{r_j, b_j\} \subseteq U \setminus W$ an interior non-edge of U. Since $|I| \leq |W| \leq 4$ and I grows monotonically, U always has some interior non-edge. Furthermore, at most two vertices from each color class are ever removed, ensuring the propagation of the precondition $|U \cap R| \geq 3$ and $|U \cap B| \geq 3$ and proper termination. We now show why the halting condition implies a forbidden configuration. The halting condition $|W| = |I|$ arises in two cases:

1. $|W| = 3$. Without loss of generality, $W = \{b_1, r_2, b_3\}$. Then r_4 is an interior vertex of U by the monotonicity of I. We will show that the copy of $K_{2,2}$ induced on $\{b_1, r_2, b_3, r_4\}$ gives a contradiction. r_4 is inside conv($\angle b_1 r_2 b_3$), hence the drawing of the $K_{2,2}$ induced on $\{b_1, r_2, b_3, r_4\}$ is not a bowtie, which by Lemma 2.3.3 yields a contradiction.

2. $|W| = 4$. Without loss of generality, $w_{\text{left}} = \{r_1, b_2\}$ and $w_{\text{right}} = \{b_3, r_4\}$. Then $\{b_5, r_5\}$ is an interior non-edge of U by the monotonicity of I. We will show that the copy of $K_{3,3}^-$ induced on $\{r_1, r_4, r_5, b_2, b_3, b_5\}$ gives a contradiction. Notice that $K_{2,3}$ is induced on $\{r_1, r_4, b_2, b_3, b_5\}$ and on $\{r_1, r_4, r_5, b_2, b_3\}$. The vertex b_2 is leftmost in the r_1-sight ordering of $\{b_2, b_3, b_5\}$, r_1 is rightmost in the b_2-sight ordering of $\{r_1, r_4, r_5\}$, b_3 is rightmost in the r_4-sight ordering of $\{b_2, b_3, b_5\}$, and r_4 is leftmost in the b_3-sight ordering of $\{r_1, r_4, r_5\}$. By applying Lemma 2.3.2 to the aforementioned two vertex sets on which $K_{2,3}$ is induced, we obtain that b_2 and b_3 both see r_5 between r_1 and r_4, and that r_1 and r_4 both see b_5 between b_2 and b_3. These conditions are

sufficient to ensure that b_5r_5 is an internal diagonal of a solid quadrilateral and hence cannot be blocked by the outside face, contradicting Lemma 2.3.4.

Now we describe how to remove wall vertices in a way that guarantees the "locking" described above, and hence the monotonicity of I. Note that removing a vertex does not affect a wall that it is not in.

If $|W| = 4$, $w_{\text{left}} = \{r_i, b_j\}$, and $w_{\text{right}} = \{b_k, r_\ell\}$, then we call $\{r_i, b_k\}$ and $\{b_j, r_\ell\}$ the diagonals of U. If both diagonals of U are non-edges, remove from U both vertices in w_{left}. If a single diagonal of U is a non-edge, then without loss of generality, $w_{\text{left}} = \{r_1, b_2\}$ and $w_{\text{right}} = \{b_1, r_3\}$. In this case, proceed to the next iteration by removing b_1 from U. Now we argue why this ensures that r_3 gets locked in the right wall. For every $i \in \{4,5\}$, $K_{2,2}$ is induced on $\{b_1, r_3, b_2, r_i\}$, hence by Lemma 2.3.3, $r_i \notin \text{int}\angle b_2r_3b_1$. Recalling that r_2 has already been removed, the next counterclockwise vertex after r_3 on the resulting convex hull boundary after removing b_1 will still be blue. Therefore, r_3 remains in the right wall.

If some wall is a non-edge, then without loss of generality, $w_{\text{right}} = \{b_1, r_1\}$. If $|W| = 3$, Without loss of generality, $w_{\text{left}} = \{r_2, b_1\}$. Remove r_1 from U, so that r_2 and b_1 will be locked in w_{left}. If $|W| = 4$, pick the vertex to remove from w_{right} in the following way. If $r_1 \in w_{\text{right}}(U \setminus \{b_1\})$ then remove b_1, otherwise remove r_1. To show why this simple action guarantees that the twin of the removed vertex 'stays' in the right wall, we need to justify that if $r_1 \notin w_{\text{right}}(U \setminus \{b_1\})$ then $b_1 \in w_{\text{right}}(U \setminus \{r_1\})$.

By hypothesis, $w_{\text{right}}(U \setminus \{b_1\}) = \{b', r'\}$ where $r' \in R \setminus \{r_1\}$. First we must explain why b_1 is the unique blue vertex in U to the right of the line $\overleftrightarrow{r'b'}$. By the definition of right wall, no vertex in $U \setminus \{b_1\}$ is to the right of the line $\overleftrightarrow{r'b'}$. But if b_1 were also to the left of the line $\overleftrightarrow{r'b'}$, then r' together with b' would constitute the right wall of U, contradicting $\{b_1, r_1\} = w_{\text{right}}(U)$. Therefore, b_1 is the unique blue vertex to the right of $\overleftrightarrow{r'b'}$. Initialize a dynamic line L to $\overleftrightarrow{r'b'}$. Rotate L clockwise around $\text{conv}(U \setminus \{b_1, r_1\})$ until it becomes horizontal, allowing it to sweep the entire portion of the half-plane above the x-axis to the right of $\overleftrightarrow{r'b'}$. The vertex b_1 is the unique blue vertex of U swept by L. Denote by $\hat{r}$ the other vertex of $U \setminus \{r_1\}$ on L at the precise moment when b_1 is swept by L, which is unique by the general position assumption. It is possible that $\hat{r} = r'$. No vertex of $U \setminus \{r_1\}$ is to the right of the line $\overleftrightarrow{\hat{r}b_1}$. Therefore, $\{\hat{r}, b_1\} = w_{\text{right}}(U \setminus \{r_1\})$.

This completes an informal and yet complete specification of the algorithm that shows that every 1-obstacle representation of $K^*_{5,5}$ has a forbidden configuration of vertices resulting in a contradiction. Therefore, the obstacle number of G'_1, i.e., $K^*_{5,5}$, is greater than *one*. This implies that the obstacle number of G'_1 is *two*, per its obstacle representation in Fig. 2.4.

□

2.3.2 A 70-vertex split graph

Theorem 2.3.8. *The split graph $G'_2 := CE(6)$, consisting of a clique of 6 blue vertices and an independent set of 64 red vertices each of which has a distinct set of neighbors, has obstacle number greater than* one.

Proof. While the graph $CE(6)$ is defined unambiguously by the theorem statement, we give the following definition of the graph family $CE(k)$ in order to assign unique names to the vertices of $CE(6)$, and to be able to refer to its induced subgraphs. Denote by $[k]$ the set of integers $\{1, 2, \dots, k\}$. For $k \in \mathbb{Z}^+$, let $B(k) = \{b_1, b_2, \dots, b_k\}$ be a set of k light blue vertices, and let $R(k) = \{r_A \mid A \subseteq [k]\}$ be a set of 2^k dark red vertices. Let $CE(k)$ be the graph on $B(k) \uplus R(k)$ in which $B(k)$ is a clique, $R(k)$ is an independent set, and there is an edge between $b_i \in B(k)$ and $r_A \in R(k)$ if and only if $i \in A$.

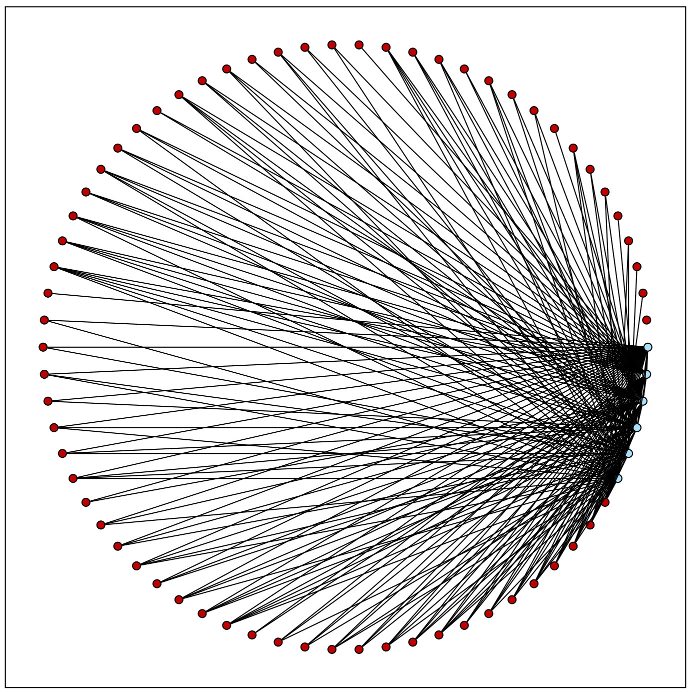

Figure 2.8 – Drawing of G_2', i.e., $CE(6)$, whose vertex set consists of a clique (light blue) of six vertices and an independent set (dark red) of 64 vertices with distinct neighborhoods

First we present lemmas regarding 1-obstacle representations of $CE(4)$ that will prove instrumental in showing that $CE(6)$ does not have a 1-obstacle representation. We do this by exploiting the hereditary nature of the CE family, that is, whenever $k' < k$, copies of $CE(k')$ can be found as an induced subgraph of $CE(k)$ in a color-preserving fashion.

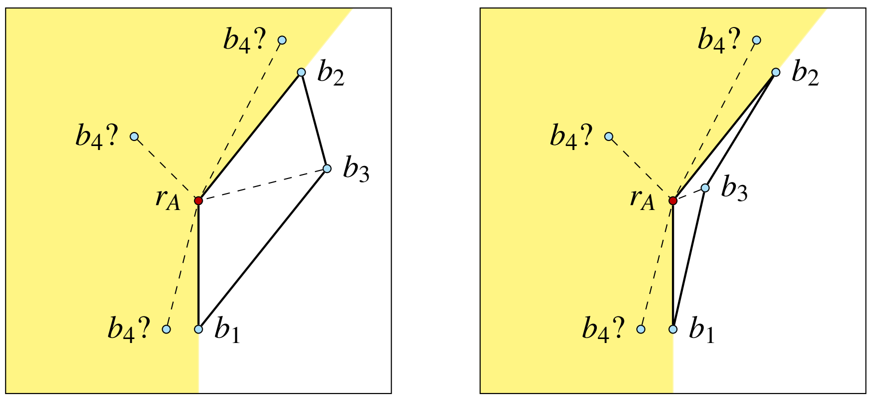

(a) Case of convex $b_1 r_A b_2 b_3$ (b) Case of concave $b_1 r_A b_2 b_3$

Figure 2.9 – For the proof of Lemma 2.3.9. The red vertex r_A is *A-fragmented* with $1,2 \in A$ and $3,4 \in \overline{A}$. Without loss of generality, b_3 is in $\operatorname{conv}(\angle b_1 r_A b_2)$ (unshaded) while b_4 is in the complement of $\operatorname{conv}(\angle b_1 r_A b_2)$ (shaded).

When considering obstacle representations for $CE(k)$, for a fixed index set $A \subseteq [k]$ we denote $[k] \setminus A$ by $\overline{A}$. For a fixed placement of the blue vertices $B(k)$ in general position, we say a point p in general position with respect to $B(k)$ is *A-fragmented* if there are distinct $i_1, i_2 \in A$ and distinct $i_3, i_4 \in \overline{A}$ such that $\angle b_{i_1} p b_{i_2}$ separates b_{i_3} from b_{i_4}.

Lemma 2.3.9. *For every integer $k \geq 4$, every obstacle representation of $CE(k)$ in which some $r_A \in R(k)$ is* A-fragmented *involves at least two obstacles.*

Proof. For an arbitrary $k \geq 4$, consider an obstacle representation of $CE(k)$ in which for a certain $A \subseteq [k]$, r_A is *A-fragmented*. Without loss of generality $1,2 \in A$ and $3,4 \in \overline{A}$ with $b_3 \in \operatorname{conv}(\angle b_1 r_A b_2)$ and $b_4 \notin \operatorname{conv}(\angle b_1 r_A b_2)$. (See Fig. 2.9.)

Then the quadrilateral $Q = b_1 r_A b_2 b_3$ is non-self-intersecting and has $r_A b_3$ as an internal diagonal. Hence, an obstacle is needed inside Q, which is interior-disjoint from the complement of $\mathrm{conv}(\angle b_1 r_A b_2)$, in order to block $r_A b_3$. Since $r_4 \notin \mathrm{conv}(\angle b_1 r_A b_2)$, so is $r_A b_4$, therefore a different obstacle must block $r_A b_4$. $\square$

To simplify the notation for red vertices, from now on we will write the subscript i instead of $\{i\}$, and $\bar{i}$ instead of $[k] \setminus \{i\}$ whenever convenient. We will also write B instead of $B(k)$ where the value of k is clear from context.

Lemma 2.3.10. *For every integer $k \geq 4$, in every 1-obstacle representation of $CE(k)$, B is in convex position.*

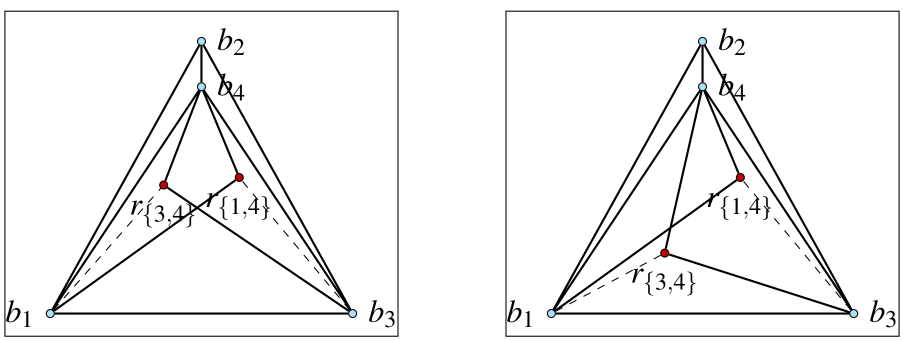

(a) Subcase of $r_{\{3,4\}}$ above $\overline{b_1 r_{\{1,4\}}}$ (b) Subcase of $r_{\{3,4\}}$ below $\overline{b_1 r_{\{1,4\}}}$

Figure 2.10 – For the proof of Lemma 2.3.10 Case 1

Proof. By Carathéodory's Theorem, it is sufficient to prove the result for $k = 4$.

Assume for contradiction that we are given a 1-obstacle representation of $CE(4)$ in which B is not in convex position. Without loss of generality, b_4 is inside the triangle $\Delta b_1 b_2 b_3$. There are two cases to consider.

Case 1: The obstacle is in $\mathrm{conv}(B)$. Without loss of generality, the obstacle is inside $\Delta b_1 b_4 b_3$. The vertex $r_{\{1,4\}}$ has non-edges to b_2 and b_3, so if it were outside of $\Delta b_1 b_2 b_3$ then at least one of these two non-edges would be outside of $\Delta b_1b_2b_3$, requiring a second obstacle. Nor can $r_{\{1,4\}}$ be inside $\Delta b_1b_2b_4$ or $\Delta b_2b_3b_4$, since that would cause its non-edge with b_2 to be in inside that triangle, again requiring a second obstacle. A symmetric argument applies to $r_{\{3,4\}}$. Notice that $r_{\{1,4\}} \in \text{conv}(\angle b_4b_2b_3)$, lest it be $\{1,4\}$-*fragmented*. Likewise, $r_{\{3,4\}} \in \text{conv}(\angle b_1b_2b_4)$, lest it be $\{3,4\}$-*fragmented*. Then without loss of generality, $r_{\{1,4\}}$ is inside $\Delta r_{\{3,4\}}b_4b_3$ which $b_1r_{\{3,4\}}$ is outside of. (See Fig. 2.10.) Hence, distinct obstacles are required to block $b_1r_{\{3,4\}}$ and $b_3r_{\{1,4\}}$, a contradiction.

Case 2: The obstacle is outside of $\text{conv}(B)$. Then $r_{\overline{4}} \notin \text{conv}(B)$ and without loss of generality, $r_{\overline{4}} \in \text{conv}(\angle b_1b_4b_3)$. Hence the obstacle is inside $\Delta b_1b_3r_{\overline{4}}$. Since the quadrilateral $Q = b_1b_4b_3r_{\overline{4}}$ is convex, every point outside of Q has a segment joining it to b_4 or $r_{\overline{4}}$ without crossing Q. Therefore, every remaining red vertex without an edge to b_4, in particular, $r_{\{1,3\}}$, is inside Q. The introduction of $r_{\{1,3\}}$ into the drawing results in interior-disjoint quadrilaterals $Q' = b_1r_{\{1,3\}}b_3r_{\overline{4}}$ and $Q'' = b_1r_{\{1,3\}}b_3b_4$. (See Fig. 2.11.) Since $r_{\{1,3\}}r_{\overline{4}}$ is inside Q' and $r_{\{1,3\}}b_4$ is inside Q'', distinct obstacles are required to block these non-edges, a contradiction. □

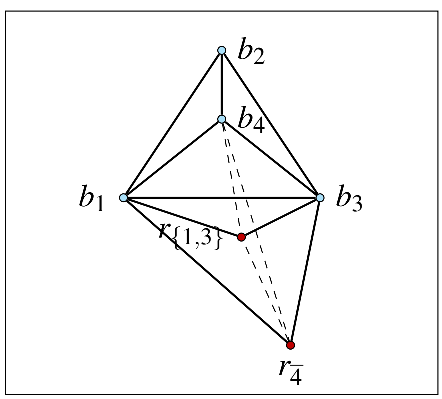

Figure 2.11 – For the proof of Lemma 2.3.10 Case 2. The assumptions lead without loss of generality to the configuration shown here, with $r_{\{1,3\}}b_4$ and $r_{\{1,3\}}r_{\overline{4}}$ requiring distinct obstacles.

Now that we have some restrictions on the relative positions of blue vertices in all 1-obstacle representations of $CE(k)$ for all $k \geq 4$, we pursue the question of where the red vertices can be positioned with respect to the blue vertices.

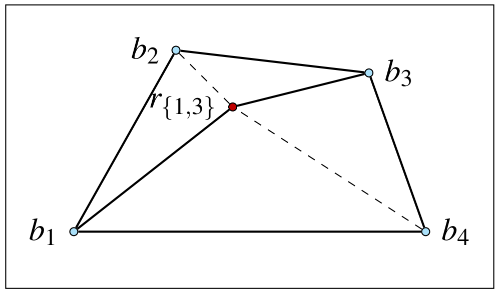

Figure 2.12 – For the proof of Lemma 2.3.11 Case 2. The vertex $r_{\{1,3\}}$ is $\{1,3\}$-*fragmented*.

Lemma 2.3.11. *For every integer $k \geq 4$, in every 1-obstacle representation of $CE(k)$ the obstacle is outside of* $\text{conv}(B)$*, and hence $R \setminus \{r_B\}$ is outside of* $\text{conv}(B)$.

Proof. By Carathéodory's Theorem, it is sufficient to establish that the obstacle is outside of $\text{conv}(B)$ in the special case $k = 4$.

Assume for contradiction that we are given a 1-obstacle representation of $CE(4)$ such that the obstacle is in $\text{conv}(B)$. By Lemma 2.3.10, B is in convex position. Without loss of generality, $b_1b_2b_3b_4$ is a clockwise enumeration of B.

Case 1: $r_{\{1,3\}} \notin \text{conv}(B)$. Imagine the polygon $b_1b_2b_3b_4$ bounding $\text{conv}(B)$ as opaque: Since it is convex, $r_{\{1,3\}}$ sees some side of $\text{conv}(B)$ in its entirety, hence $r_{\{1,3\}}$ sees b_i for a certain even i. This means $r_{\{1,3\}}b_i$ is outside $\text{conv}(B)$, hence it will require a separate obstacle, a contradiction.

Case 2: $r_{\{1,3\}} \in \text{conv}(B)$. (See Fig. 2.12.) By the convexity of B, $r_{\{1,3\}}$ is $\{1,3\}$-*fragmented*, a contradiction.

We have established that the obstacle is outside of conv(B). Assume for contradiction that some vertex $r \in R \setminus \{r_B\}$ is inside conv(B). Since r has a non-edge to some vertex in B, this non-edge must be inside conv(B). Since the boundary of conv(B) is a solid polygon, this requires an obstacle inside conv(B), contrary to what we just proved. □

We introduce some further terminology to use in the context of $CE(k)$ (for any integer $k > 0$) for a fixed arrangement of B. The following definitions are meant only for points outside of conv(B) and in general position with respect to B.

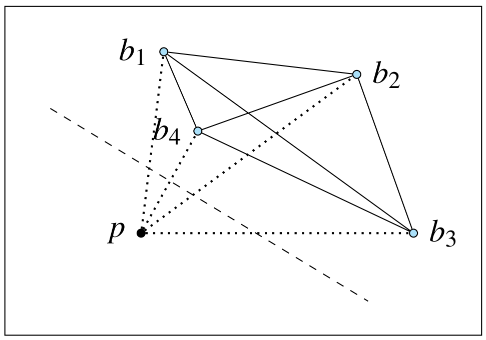

Figure 2.13 – Illustration of the concepts *A-straight*, *A-convex*, and *A-reflex*. For $CE(4)$ consider the given placement of B. The point p is $\{2,3\}$-*straight*, $\{4\}$-*convex*, and $\{1,4,3\}$-*reflex*.

For a given $A \subseteq [k]$, let B_A denote $\{b_i \mid i \in A\}$. We say that a point p is *A-straight* if some line through p separates B_A and $B_{\overline{A}}$ (vacuously true if $A \in \{\varnothing, [k]\}$). We say that a point p is *A-convex* if it sees $B_A \neq \varnothing$ between two non-empty parts of $B_{\overline{A}}$ that comprise $B_{\overline{A}}$. If a point is $\overline{A}$-*convex*, we say it is *A-reflex*. Observe that if r_A is *A-reflex*, then an obstacle is required in a bounded face, but not necessarily in the unbounded face. Note that for every $A \subseteq [k]$, every point $p \notin$ conv(B) in general position with respect to B is either *A-straight*, *A-convex*, *A-reflex*, or *A-fragmented*. (See Fig. 2.13.)

Now we can finish proving with relative ease that $CE(6)$ does not admit a 1-obstacle representation. Assume for contradiction that we are given a 1-obstacle representation of $CE(6)$. By Lemma 2.3.10, B is in convex position. Without loss of generality, $b_1b_2b_3b_4b_5b_6$ is a clockwise enumeration of B. By Lemma 2.3.11, $R \setminus \{r_B\}$ is outside of conv(B). In particular, $r_{\{1,3,5\}}$ is outside of conv(B). We will show that every point outside of conv(B) and in general position with respect to B is $\{1,3,5\}$-*fragmented* by showing that it is neither $\{1,3,5\}$-*straight* nor $\{1,3,5\}$-*convex* nor $\{1,3,5\}$-*reflex*.

Since $\{b_1,b_3,b_5\}$ is not linearly separable from $\{b_2,b_4,b_6\}$, no $\{1,3,5\}$-*straight* point exists.

Assume for contradiction that some point p is $\{1,3,5\}$-*convex*. Hence p sees odd-subscripted blue vertices together between two sets of even-subscripted blue vertices. Then p is $\{i,j\}$-*straight* for some $\{i,j\} \subseteq \{2,4,6\}$. But b_ib_j is a diagonal of the bounding hexagon of B, which contradicts that it is linearly separable from $B \setminus \{b_i,b_j\}$. By a symmetric argument, no point is $\{2,4,6\}$-*convex* (i.e., $\{1,3,5\}$-*reflex*) either. Therefore, $r_{\{1,3,5\}}$

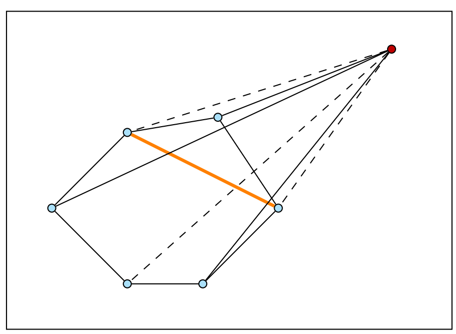

Figure 2.14 – For the proof of Theorem 2.3.8. A red vertex r_A adjacent exactly to blue vertices non-adjacent in the bounding polygon of B is *A-fragmented* no matter what, as in this example.

is $\{1,3,5\}$-*fragmented*, requiring two obstacles, a contradiction.

Therefore, G_2', i.e., $CE(6)$, has obstacle number greater than *one*. $\square$

2.3.3 A 10-vertex graph with bipartite complement

We showed in [36] that the $(2,2)$-colorable 20-vertex graph G_3 has obstacle number greater than 1. One can obtain G_3 from $CE(4)$ by adding all possible edges among the vertices in the independent set of 16 red vertices. Here, we show that a 10-vertex induced subgraph of it, G'_3, also has obstacle number greater than 1.

Let G'_3 be the graph consisting of a clique of light blue vertices $B = \{b_i \mid i \in [4]\}$, a clique of dark red vertices $R = \{r_A \mid A \in \binom{[4]}{2}\}$, and additional edges between every b_i and every r_A with $i \in A$. (See Fig. 2.15.)

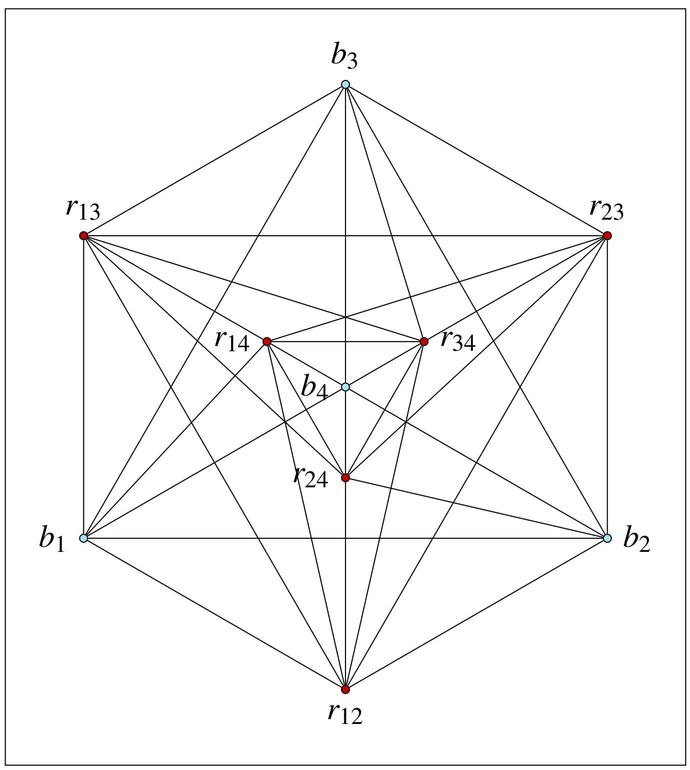

Figure 2.15 – Drawing of the bipartite complement graph G'_3 with rotational symmetry

Theorem 2.3.12. *G'_3, a $(2,2)$-colorable graph on* ten *vertices, has obstacle number greater than* one.

Proof. Assume for contradiction that we are given a 1-obstacle representation of G'_3. Following the terminology in the preceding section, we shall say that a red vertex r_A is *fragmented* if it is *A-fragmented*. That is, a vertex r_A is not fragmented if and only if there are points p and q such that $\angle p r_A q$ strictly separates $\{b_i \mid i \in A\}$ from the remaining blue vertices. If some red vertex r_A is fragmented, then *two* obstacles will be required due to $\{r_A\} \cup B$, a contradiction.

Case 1: B is not in convex position. Without loss of generality, b_4 is inside the triangle $\Delta b_1 b_2 b_3$.

Subcase 1a: The obstacle is in $\mathrm{conv}(B)$. Case 1 of the proof of Lemma 2.3.10 is based only on the vertices b_1, b_2, b_3, b_4, $r_{\{1,4\}}$ and $r_{\{2,4\}}$, under the same conditions, hence that argument applies verbatim to yield a contradiction here.

Subcase 1b: The obstacle is outside of $\mathrm{conv}(B)$. Let $C_{\{1,2\}} = \mathrm{conv}(\angle b_2 b_4 b_1)$, $C_{\{1,3\}} = \mathrm{conv}(\angle b_1 b_4 b_3)$, and $C_{\{2,3\}} = \mathrm{conv}(\angle b_3 b_4 b_2)$. Every red vertex is in precisely one of these regions and outside of $\mathrm{conv}(B)$. Let $f : \binom{[3]}{2} \to \binom{[3]}{2}$ be the map such that $r_A \in C_{f(A)}$ whenever $A \in \binom{[3]}{2}$. We will show that every possible assumption about f leads to a contradiction.

Assume for contradiction that f has a fixed point. Without loss of generality, $r_{\{1,2\}} \in C_{\{1,2\}}$. This means that $Q = b_2 b_4 b_1 r_{\{1,2\}}$ is a solid convex quadrilateral, hence to block $b_4 r_{\{1,2\}}$, the obstacle is inside Q. Then, $r_{\{3,4\}}$ must be inside Q in order for the obstacle to block both $b_1 r_{\{3,4\}}$ and $b_2 r_{\{3,4\}}$. But then, $\angle b_4 r_{\{3,4\}} r_{\{1,2\}}$ partitions Q into disjoint quadrilateral regions with solid boundaries that respectively contain $b_1 r_{\{3,4\}}$ and $b_2 r_{\{3,4\}}$. Hence, two obstacles are required, a contradiction. Therefore, f has no fixed point.

Assume for contradiction that f is not a permutation. Without loss of generality, $r_{\{1,3\}}$ and $r_{\{2,3\}}$ are both in $C_{\{1,2\}}$. In order for both of these red vertices to not be fragmented,

$\overleftrightarrow{b_3b_4}$ must separate $b_1r_{\{1,3\}}$ and $b_2r_{\{2,3\}}$. Hence, $Q = b_2b_1r_{\{1,3\}}r_{\{2,3\}}$ is a solid, non-self-intersecting quadrilateral. If Q is concave, we get an immediate contradiction due to Q separating its diagonals, both of which are non-edges in G_3'. If Q is convex, the obstacle is inside Q in order to block its diagonals. But since $r_{\{1,2\}}$ is outside of $C_{\{1,2\}}$, it does not meet $\mathrm{conv}(Q)$, requiring another obstacle, a contradiction. Therefore, f is a permutation.

Since f is a permutation of three elements with no fixed point, it is cyclic. Without loss of generality, $r_{\{1,2\}} \in C_{\{2,3\}}$ and $r_{\{1,3\}} \in C_{\{1,2\}}$. In order to not be fragmented, $r_{\{1,2\}}$ is on the same side of $\overleftrightarrow{b_1b_4}$ as b_2, and $r_{\{1,3\}}$ is on the same side of $\overleftrightarrow{b_3b_4}$ as b_1. These conditions ensure that b_1b_4 does not meet $r_{\{1,2\}}r_{\{1,3\}}$. Indeed, if b_2b_4 and $r_{\{1,2\}}r_{\{1,3\}}$ did meet at some point p, then the convex solid quadrilateral $b_1r_{\{1,3\}}pb_4$ would have $b_2r_{\{1,3\}}$ inside and $b_4r_{\{1,2\}}$ outside, requiring two obstacles, a contradiction. If not, then $b_1r_{\{1,3\}}r_{\{1,2\}}b_2b_4$ is a non-self-intersecting solid pentagon with $b_2r_{\{1,3\}}$ inside and $b_4r_{\{1,2\}}$ outside, requiring two obstacles, a contradiction.

Having exhausted all possibilities, we have shown that the assumptions of Subcase 1b lead to a contradiction.

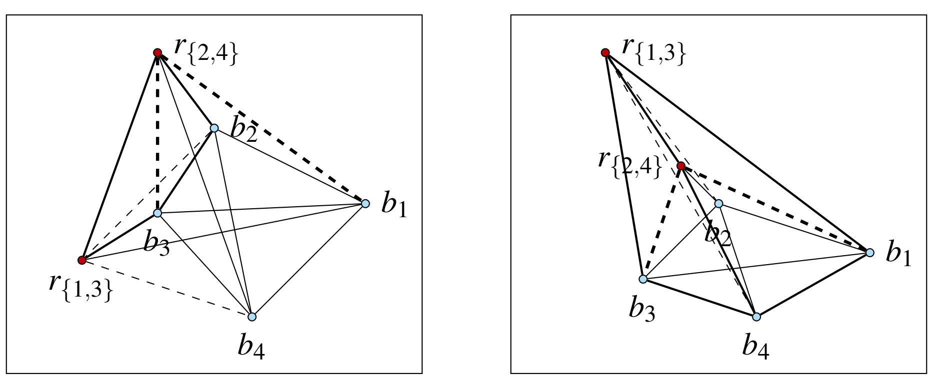

(a) Subcase 2a (b) Subcase 2b

Figure 2.16 – For the proof of Theorem 2.3.12 Case 2. The thick dashed non-edges require distinct obstacles.

Case 2: B is in convex position. Without loss of generality, the bounding polygon of B is $b_1b_2b_3b_4$. To not be fragmented,

(i) $r_{\{1,3\}}$ and $r_{\{2,4\}}$ must lie outside of $\mathrm{conv}(B)$;

(ii) for $r_{\{1,3\}}$, either $b_1, b_3 \in \mathrm{conv}(\angle b_2 r_{\{1,3\}} b_4)$ or $b_2, b_4 \in \mathrm{conv}(\angle b_1 r_{\{1,3\}} b_3)$; and

(iii) for $r_{\{2,4\}}$, either $b_1, b_3 \in \mathrm{conv}(\angle b_2 r_{\{2,4\}} b_4)$ or $b_2, b_4 \in \mathrm{conv}(\angle b_1 r_{\{2,4\}} b_3)$.

Subcase 2a: $b_1, b_3 \in \mathrm{conv}(\angle b_2 r_{\{1,3\}} b_4)$ and $b_2, b_4 \in \mathrm{conv}(\angle b_1 r_{\{2,4\}} b_3)$. Without loss of generality, the quadrilateral $b_4b_1b_2r_{\{1,3\}}$ is convex and has b_3 inside, and without loss of generality, the quadrilateral $b_3b_4b_1r_{\{2,4\}}$ is convex and has b_2 inside. Hence, $b_2b_3r_{\{1,3\}}r_{\{2,4\}}$ is a solid convex quadrilateral with $b_1r_{\{2,4\}}$ outside and $b_3r_{\{2,4\}}$ inside. Therefore, two obstacles are required, a contradiction.

Subcase 2b: $b_2, b_4 \in \mathrm{conv}(\angle b_1 r_{\{1,3\}} b_3)$ or $b_1, b_3 \in \mathrm{conv}(\angle b_2 r_{\{2,4\}} b_4)$. Due to symmetry, we proceed assuming the former. Without loss of generality, $Q = b_3b_4b_1r_{\{1,3\}}$ is a convex quadrilateral. The obstacle is inside Q due to $r_{\{1,3\}}b_4$. In order for $b_1r_{\{2,4\}}$ and $b_3r_{\{2,4\}}$ to be blocked, $r_{\{2,4\}}$ is inside Q. Hence, $\angle r_{\{1,3\}} r_{\{2,4\}} b_4$ partitions $\mathrm{conv}(Q)$ into two regions with solid boundaries that respectively contain $b_1r_{\{2,4\}}$ and $r_{\{2,4\}}b_3$. Therefore, two obstacles are required, a contradiction.

Therefore, G'_3 has obstacle number greater than *one*. □

2.4 Improved bounds via order types

In this section, we show that there exist graphs on n vertices with obstacle number at least $\Omega(n/\log^2 n)$. Better results are proved if we are allowed to use only *convex* obstacles or polygonal obstacles with a small number of sides.

We improve upon Theorem 2.1.2 by using some estimates on the number of different *order types* of n points in the Euclidean plane, discovered by Goodman and Pollack [24],

[25] (see also Alon [2]). We establish the following results.

Theorem 2.4.1. *The number of graphs on n (labeled) vertices with obstacle number at most h is at most*

$$2^{O(hn\log^2 n)}.$$

Theorem 2.4.2. *There exist graphs G on n vertices with obstacle numbers*

$$\mathrm{o}bs(G) \geq \Omega\left(n/\log^2 n\right).$$

Note that Theorem 2.4.2 directly follows from Theorem 2.4.1. Indeed, since the total number of (labeled) graphs with n vertices is at least $2^{\Omega(n^2)}$, the existence of an n-vertex graph with obstacle number greater than $h(n)$ is guaranteed for sufficiently large n for any function h such that $h(n) \cdot n\log^2 n = o(n^2)$, that is, provided that $h(n) = o(n/\log^2 n)$.

The aim of this section is to prove Theorems 2.4.1and 2.4.2. The idea is to find a short encoding of the obstacle representations of graphs, and to use this to give an upper bound on the number of graphs with low obstacle number.

We need to review some simple facts from combinatorial geometry. Two sets of points, P_1 and P_2, in general position in the plane are said to have the same *order type* if there is a one to one correspondence between them with the property that the orientation of any triple in P_1 is the same as the orientation of the corresponding triple in P_2. Counting the number of different order types is a classical task.

Theorem 2.4.3 (Goodman, Pollack [24])**.** *The number of different order types of n points in general position in the plane is* $2^{O(n\log n)}$.

Remark 2.4.1. *The upper bound* $2^{O(n\log n)}$ *holds also for the number of order types of n* labeled *points, since the number of permutations of n points is* $n! = 2^{O(n\log n)}$.

In a graph drawing, the *complexity* of a face is the number of edges of the face, when viewed as a polygonal set. The following result was proved by Arkin, Halperin, Kedem, Mitchell, and Naor (see Matoušek, Valtr [31] for its sharpness).

Theorem 2.4.4 (Arkin et al. [5])**.** *The complexity of a single face in a drawing of a graph with n vertices is at most* $O(n \log n)$.

Note that this bound does not depend of the number of edges of the graph.

Proof of Theorem 2.4.1. For any given graph G with n vertices that admits an h-obstacle representation, fix such a representation and consider the graph drawing. Every obstacle belongs to a single face in the straight-line drawing of G according to this obstacle representation. In view of Theorem 2.4.4, the complexity of every face is $O(n \log n)$. Replacing each obstacle by a slightly shrunk copy of the face containing it, we can achieve that every obstacle *is* a polygonal region with $O(n \log n)$ sides.

Let S be the point sequence starting with the vertices of G, followed by the vertices of every obstacle in cyclic order, one entire obstacle after another. Let I be the set of the starting positions of the h obstacles in S. G is completely determined by the (labeled) order type of S, together with I. To see this, first observe that I tells us which points in S are graph vertices and which pairs in S define a side of a polygon. Now, notice that a given segment uv among graph vertices is blocked if and only if it meets some side ab of some polygon, for which a necessary and sufficient condition is that the ordered triples uav, avb, vbu, and bua have the same orientation.

Denote by N the length of S. The number of possibilities for I is at most $\binom{N}{h} \leq N^h$. Since $N \leq n + c_1 hn \log n$ for some absolute constant $c_1 > 0$, according to Theorem 2.4.3

and Remark 2.4.1, the number of graphs with obstacle number at most h is at most

$$N^h \cdot 2^{O(N \log N)} = 2^{O(N \log N)},$$

in other words, it is less than $2^{chn \log^2 n}$ for a suitable constant $c > 0$. $\square$

This is a generous upper bound due to overcounting, and also because most pairs (S,I) do not encode obstacle representations.

2.5 A bound for bipartite graphs via binary space partitions obtained by ham sandwich cuts

We present an alternate proof of the fact that the family of bipartite graphs has unbounded obstacle number without invoking results from extremal graph theory. The proof can be generalized to the family of split graphs and those with bipartite complement. The method used in this proof can be applied more generally in the study of obstacle representations.

Alternate proof of Corollary 2.1.3. Denote by $G(b)$ the bipartite graph consisting of an independent set $B(b)$ of b blue vertices and another independent set $R(b)$ of 2^b red vertices such that each vertex has a distinct set of neighbors. Recall that the bipartite graph G_1' obtained from $K_{5,5}$ by removing a maximum matching has obstacle number *two*. Since $G(5)$ has G_1' as an induced subgraph, it also has obstacle number greater than *one*. We define a function $k : \mathbb{N} \to \mathbb{N}$ such that the graph $G(k(h))$ has no h-obstacle representation.

Let $k(0) = k(1) = 5$. To define k on $h > 1$, let

$$a = \frac{|R(5)|}{|R(5)| - 1} = \frac{2^5}{2^5 - 1} = \frac{32}{31},$$

let $h' = h'(h) = 2^{\lceil \log_2 h \rceil}$, let $\ell = \ell(h) = \lfloor 1 + \log_a h' \rfloor$ and let $k(h) = 5\ell h'$.

For $h > 1$, we will show that $G(k(h))$ cannot be represented with $h' \geq h$ obstacles. Given an obstacle representation of $G(k(h))$ for $h > 1$, denote by B and R the sets of blue and red vertices. Decompose the plane into h' convex cells by taking ham sandwich cuts of $\{B, R\}$ recursively down to depth $\log_2 h'$. See Fig. 2.17 for an illustration. Since h' is a power of 2, it divides both $|B|$ and $|R|$, so by the general position assumption, no vertex is on a cell boundary and every cell has *exactly* $1/h'$ of B as well as of R.

We will show that every cell has an induced copy of $G(5)$. Given a cell $\mathscr{C}$, let $B' = B \cap \mathscr{C}$ and let $R' = R \cap \mathscr{C}$. Notice that $|B'| = 5\ell$. Partition B' into ℓ blocks $B_1, B_2, \ldots, B_\ell$ each of size 5. Let I_j be the subscript set of B_j. If for some block $j \in [\ell]$ the condition $\{A \cap I_j \mid r_A \in R'\} = 2^{I_j}$ held, then there would be an induced $G(5)$ in $\mathscr{C}$. Assume for contradiction that this is not the case, so for every $j \in [\ell]$, we have

$$\left|\{A \cap I_j \mid r_A \in R'\}\right| \leq \left|2^{I_j}\right| - 1 = 2^5 - 1.$$

But this will give us

$$\frac{|R|}{|R'|} \geq \left(\frac{2^5}{2^5 - 1}\right)^{\ell} = a^{\ell} > a^{\log_a h'} = h',$$

which is contrary to the definition of R' due to the strictness of the inequality.

This means each of the h' cells has an induced copy of $G(5)$. Every such $G(5)$ copy requires an internal obstacle, confined to a face within its cell, in addition to an obstacle outside that face. Hence h' internal obstacles and at least 1 other obstacle are necessary. Therefore, the obstacle number of $G(k(h))$ is greater than h'. □

Notice that in the above proof, the function $k(h)$ is $\Theta(h \log h)$. This means that $G(h)$

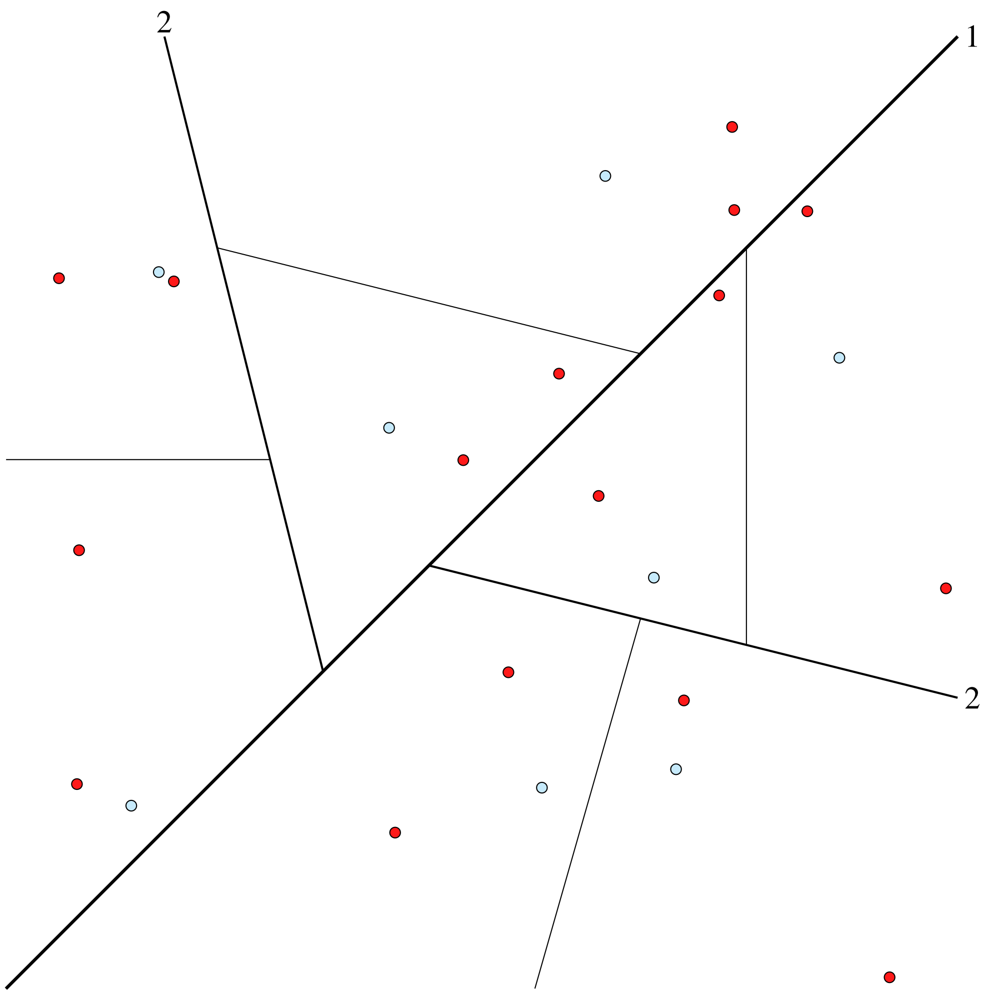

Figure 2.17 – Illustration of the construction in the alternate proof of Corollary 2.1.3. For an arbitrary set B of 8 (light) blue points and an arbitrary set R of 16 (dark) red points that taken together are in general position, a binary space partition is obtained by recursive ham sandwich cuts down to depth 3. Separating line of depth 1 and rays of depth 2 are labeled with their respective depths. Each of the resulting 8 disjoint open convex cells has exactly 1 blue point and 2 red points.

has at most $2^{O(h\log h)}$ vertices. In other words, there are bipartite graphs on n vertices with obstacle numbers at least $\Omega\left(\frac{\log n}{\log\log n}\right)$. Despite the progress that we made using order types, this remains the best lower bound we know of on the obstacle numbers of bipartite graphs. This statement holds also for split graphs and graphs with bipartite complement.

Chapter 3

Obstacle numbers of graphs: Convex, segment, and three dimensional variants

3.1 Bounds via order types for convex and segment obstacle numbers

In this section, we apply the techniques in Section 2.4 to prove slightly better bounds for convex obstacle numbers, and significantly better bounds for segment obstacle numbers. Denote by $\mathrm{o}bs_c(G)$ the convex obstacle number of a graph G, and by $\mathrm{o}bs_s(G)$ its segment obstacle number.

Theorem 3.1.1. *There exist graphs G on n vertices with convex obstacle numbers*

$$\mathrm{o}bs_c(G) \geq \Omega\left(n/\log n\right).$$

Proof. It is enough to bound the number of graphs that admit an obstacle representation

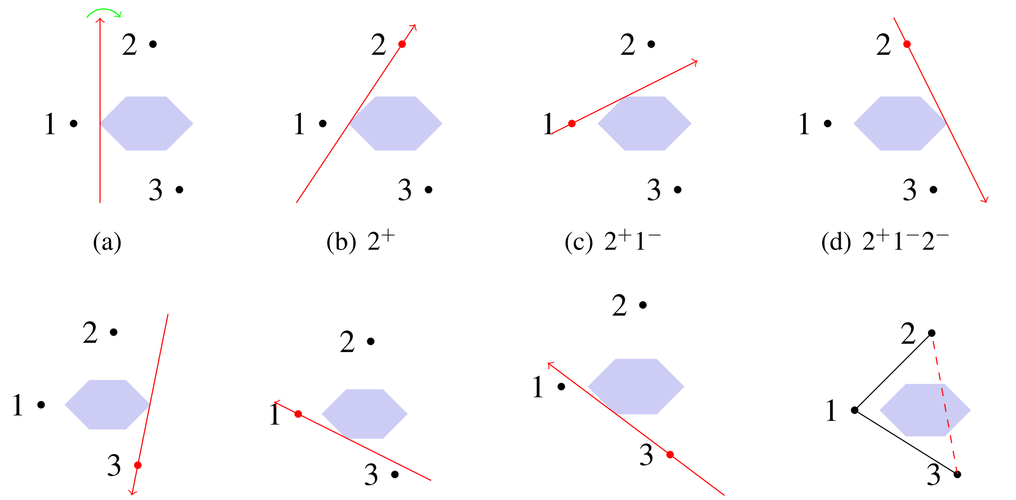

(a) (b) 2^+ (c) 2^+1^- (d) $2^+1^-2^-$

(e) $2^+1^-2^-3^+$ (f) $2^+1^-2^-3^+1^+$ (g) $2^+1^-2^-3^+1^+3^-$ (h) $2^+1^-2^-3^+1^+3^-$

Figure 3.1 – Subfigures (a) to (g) show the construction of the sequence and (h) shows the visibilities. The arrow on the tangent line indicates the direction from the point of tangency in which we assign $+$ as a label to the vertex. The additional arrow in (a) indicates that the tangent line is rotated clockwise around the obstacle.

with at most h convex obstacles. Let us fix such a graph G, together with a representation. Let V be the set of points representing the vertices, and let $O_1, \ldots, O_h$ be the convex obstacles. For any obstacle O_i, rotate an oriented tangent line ℓ along its boundary in the clockwise direction. We can assume without loss of generality that ℓ never passes through two points of V. Let us record the sequence of points met by ℓ. If $v \in V$ is met at the right side of ℓ, we add the symbol v^+ to the sequence, otherwise we add v^- (Fig. 3.1). When ℓ returns to its initial position, we stop. The resulting sequence consists of $2n$ characters. From this sequence, it is easy to reconstruct which pairs of vertices are visible in the presence of the single obstacle O_i. Observe that O_i blocks uv if and only if the subsequence induced on u and v has no consecutive pair with the same superscript. Hence, knowing these sequences for every obstacle O_i, completely determines the visibility graph G. The number of distinct sequences assigned to a single obstacle is at most $(2n)!$, so that the number of graphs with convex obstacle number at most h cannot exceed $((2n)!)^h/h! < (2n)^{2hn}$. As long as this number is smaller than $2^{\binom{n}{2}}$, there is a graph with convex obstacle number larger than h. □

Theorem 3.1.2. *There exist graphs G on n vertices with segment obstacle numbers*

$$\mathrm{obs}_s(G) \geq \Omega\left(n^2/\log n\right).$$

Proof. The proof of Theorem 2.4.1 yields the following result when s is the average number of sides among the obstacles.

Lemma 3.1.3. *The number of graphs admitting an obstacle representation with at most h obstacles, having a total of at most hs sides, is at most*

$$2^{O(n\log n + hs\log(hs))}.$$

When all obstacles are segments ($s = 2$), Lemma 3.1.3 immediately implies Theorem 3.1.2. Indeed, since there are $2^{\binom{n}{2}} = 2^{\Theta(n^2)}$ labeled graphs on n vertices, as long as $h = h(n)$ satisfies

$$n\log n + hs\log(hs) = O(n^2),$$

we can argue that there are graphs with segment obstacle number at least $\Omega(h(n))$. □

3.2 Convex obstacle numbers of outerplanar graphs and bipartite permutation graphs

A *disjoint convex obstacle representation* is one in which the obstacles are pairwise disjoint, and we define *disjoint convex obstacle number* and *h-disjoint convex obstacle representation* similarly. The convex obstacle number of a graph is at most its disjoint convex obstacle number.

Recall from Section 2.4 that the number of graphs on n vertices with obstacle number at most h is at most $2^{O(hn\log^2 n)}$. From this, it follows that every graph class with $2^{\omega(n\log^2 n)}$ members on n vertices has unbounded obstacle number. Also recall that the number of graphs on n vertices with convex obstacle number at most h is at most $2^{O(hn\log n)}$. Since the number of planar graphs on n vertices is $2^{\Theta(n\log n)}$ (see [23] for exact asymptotics), these upper bounds are inconclusive regarding the obstacle numbers or convex obstacle numbers of the class of planar graphs or a subclass.

Nonetheless, it was shown by Alpert, Koch, and Laison [4] that every outerplanar graph admits a 1-obstacle representation in which the obstacle is in the unbounded face. In the same paper the question, whether the convex obstacle number of an outerplanar graph can be arbitrarily large, was raised. We settle this question in the negatory. In particular, we prove the following two results regarding outerplanar graphs in Subsections 3.2.1 and 3.2.2.

Theorem 3.2.1. *The convex (and disjoint convex) obstacle number of every outerplanar graph is at most* five.

Theorem 3.2.2. *There are outerplanar graphs having disjoint convex obstacle number at least* four.

In Section 3.2.3, we prove the following regarding bipartite permutation graphs.

Theorem 3.2.3. *The convex (and disjoint convex) obstacle number of every bipartite permutation graph is most* four.

3.2.1 Upper bound on convex obstacle number of outerplanar graphs

Proof of Theorem 3.2.1. We shall show that the convex obstacle number of every outerplanar graph is at most *five*, by giving a method to generate five convex obstacles that can represent any outerplanar graph. For a given connected outerplanar graph G, we first construct a digraph $\overrightarrow{G'}$ with certain properties, whose underlying graph is a subgraph of G. We call $\overrightarrow{G'}$ the BFS-digraph of G. We show an obstacle representation using *five* convex obstacles for the BFS-digraph, and then modify the representation without changing the number of obstacles to represent the graph G. We finally discuss how to accommodate the disconnected case, still with *five* obstacles.

Constructing the BFS-digraph and its properties

Let G be a connected outerplanar graph. Perform the breadth-first search based Algorithm 1 on G that outputs a digraph which we call the BFS-digraph of G, and denote by $\overrightarrow{G'}$. We say that a vertex of a BFS digraph has depth i if its distance from the BFS root is i.

Lemma 3.2.4. *A BFS-digraph $\overrightarrow{G'}$ of a connected outerplanar graph G has a straight-line drawing such that*

1. *each vertex at depth i lies on the line $y = -i$;*
2. *two edges are disjoint except possibly at their endpoints; and*
3. *a vertical downward ray starting at a vertex v meets the graph only at v.*

Input: A connected graph $G = G(V,E)$
Output: The digraph $\overrightarrow{G}'$ called the BFS digraph of G

$V' := V_0 :=$ singleton set with an arbitrarily chosen vertex of G (the BFS *root*)
$\overrightarrow{E'} := \varnothing$
$i := 0$
while $V' \neq V$ **do**
  $V_{i+1} := \{v \mid u \in V_i, (u,v) \in E\} \setminus V'$
  $V' := V' \cup V_{i+1}$
  $\overrightarrow{E'} := \overrightarrow{E'} \cup \{\overrightarrow{(u,v)} \mid u \in V_i, v \in V_{i+1}, (u,v) \in E\}$
  $i := i+1$
end while
return $\overrightarrow{G}'(V, \overrightarrow{E'})$

Algorithm 1: Algorithm to compute a BFS-digraph of a connected graph

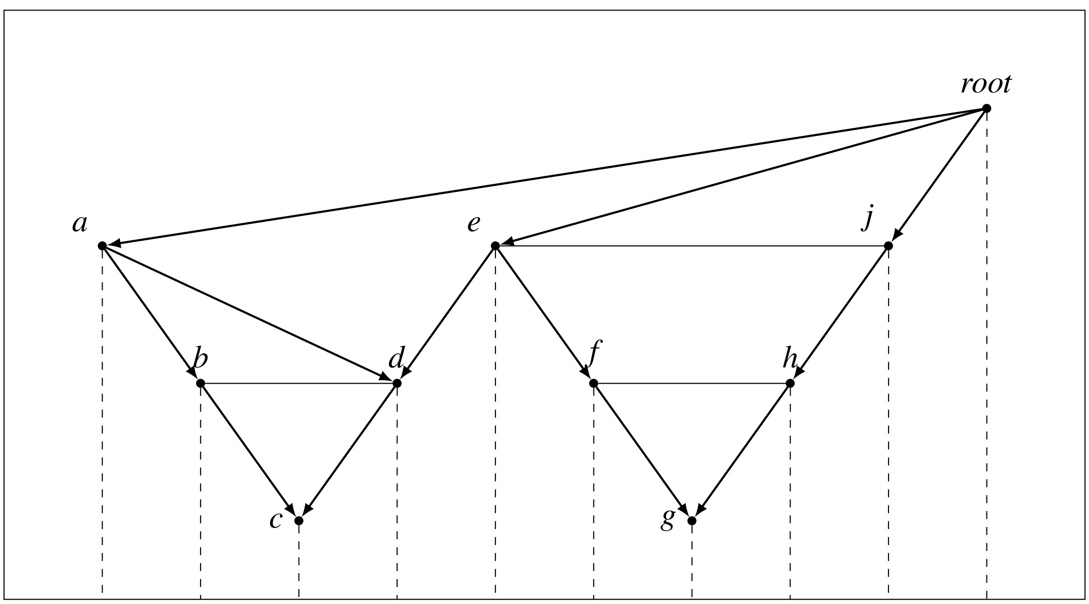

Figure 3.2 – A BFS-digraph of an outerplanar graph G drawn to exhibit the three properties in Lemma 3.2.4. The edges without arrows correspond to edges of G that are not in the digraph. For a given outerplanar graph G, regardless of the choice of the BFS root, there is a drawing of the resulting BFS-digraph that satisfies the three properties and induces a straight-line outerplanar drawing of G.

Proof. Let $\overrightarrow{G'_i}$ denote the subgraph of $\overrightarrow{G'}$ induced on vertices at depth less than or equal to i. We show the existence of such a drawing by constructing it. We will proceed by induction on i.

Consider a planar embedding of the outerplanar graph G in which every vertex meets the outer face, with all vertices on a circle having *root* as its topmost point. From now on, we do not distinguish between a graph and its embedding. Draw *root* on the line $y = 0$. Then draw all its neighbors on the line $y = -1$ and to its left, preserving their order in G. All arcs are oriented downward. So far we have $\overrightarrow{G'_1}$, which satisfies the desired properties. We now show how to extend an embedding of $\overrightarrow{G'_i}$ with the desired properties to an embedding of $\overrightarrow{G'_{i+1}}$ with the desired properties. Let $v_{i,1}, \ldots, v_{i,\ell}$ denote the members of V_i in left-to-right order. In G, the depth $i+1$ neighbors of $v_{i,k}$, for $k = 1, \ldots, \ell$, lie either on the clockwise arc from $v_{i,k}$ to $v_{i,k-1}$ (or to the parent of $v_{i,k}$ if $v_{i,k-1}$ does not exist), or on the counterclockwise arc from $v_{i,k}$ to $v_{i,k+1}$ (or to the parent of $v_{i,k}$ if $v_{i,k+1}$ does not exist). Otherwise, G is not planar, or its vertices are not in convex position. We denote the depth $i+1$ neighbors of $v_{i,k}$ on the clockwise arc from $v_{i,k}$ to $v_{i,k-1}$ (or to the parent of $v_{i,k}$ if $v_{i,k-1}$ does not exist) as the *left* children of $v_{i,k}$, and those on the counterclockwise arc from $v_{i,k}$ to $v_{i,k+1}$ (or to the parent of $v_{i,k}$ if $v_{i,k+1}$ does not exist) as the *right* children of $v_{i,k}$. Note that for vertices $v_{i,j}$ and $v_{i,j+1}$, the rightmost child of $v_{i,j}$ lies before or at the same place as the leftmost child of $v_{i,j+1}$. We apply the following steps for each $v_{i,k}$:

- put the left children of $v_{i,k}$, preserving their order with regard to their distance from $v_{i,k}$, on the line $y = -(i+1)$ so that they are to the left of $v_{i,k}$, and to the right of $v_{i,k-1}$ and the nearest left ancestor of $v_{i,k}$ (if they exist);

- put the right children of $v_{i,k}$, preserving their order with regard to their distance from $v_{i,k}$, on the line $y = -(i+1)$ so that they are to the right of $v_{i,k}$, and are to the left of

$v_{i,k+1}$ and the nearest right ancestor of $v_{i,k}$ (if they exist);

- make sure that for every pair of vertices $v_{i,j}$ and $v_{i,j+1}$, the rightmost child of $v_{i,j}$ and the leftmost child of $v_{i,j+1}$ preserve their ordering in G.

Note that due to the outerplanarity of G, a right descendent and a left descendant of a vertex have no common descendants, rendering the last step possible. Therefore, the extended embedding represents $\overrightarrow{G'_{i+1}}$, and satisfies all three conditions. □

According to this embedding, we say that two vertices are *consecutive* if they are on the same horizontal line and there is no vertex between them.

Corollary 3.2.5. *A vertex has at most two parents. Moreover, if a vertex v has two parents, the parents are consecutive; and v is the rightmost child of its left parent, and the leftmost child of its right parent.*

Proof. If any of the conditions above does not hold, property 3 of Lemma 3.2.4 is violated. □

By Corollary 3.2.5, we also know two vertices at depth i have a common child only if they are consecutive.

Corollary 3.2.6. *Two consecutive vertices such that one is a left child and the other is a right child of the same parent, do not have a common child.*

Proof. It directly follows from the third property of Lemma 3.2.4. □

5-convex obstacle representation of the BFS-digraph of a connected outerplanar graph

We demonstrate a set of five convex obstacles and describe how to place vertices of $\overrightarrow{G'}$ to obtain a 5-convex obstacle representation for $\overrightarrow{G'}$. We first describe the arrangement of the set of obstacles. We have two disjoint convex arcs symmetric about a horizontal line, such that both arcs curve toward the line of symmetry. We consider the arcs to be parts of large circles, so that they behave like lines, except that they block visibilities among vertices put sufficiently near them. In the region bounded by the two arcs, we put three line obstacles, which form an S-shape with perpendicular joints, so that the S-shape is equally far from either arc, and the projection of the S-shape onto either arc covers the whole arc. We then disconnect the line obstacles by creating a small (and similar) aperture at each joint. The arrangement of the set of obstacles is shown in Figure 3.3.

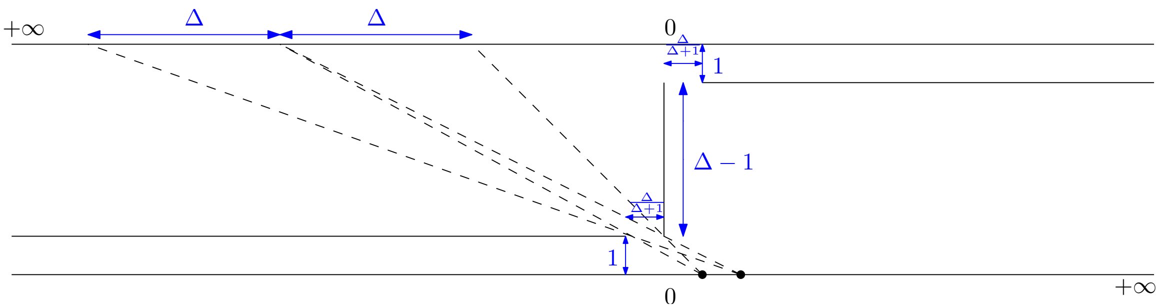

Figure 3.3 – The arrangement of the set of five convex obstacles for outerplanar graphs

The key idea is to place all vertices of the graph sufficiently close to either of the arcs, and control the visibilities through the created apertures. For the sake of simplicity of exposition, from now on, we say a vertex is placed on an arc if it is sufficiently close to an arc. For each arc, the nearby and distant apertures are respectively called the *outgoing-aperture* and the *incoming-aperture*. For each vertex on an arc, we draw the outgoing edges through the outgoing-aperture of its underlying arc. We parameterize the arcs such that the intersection points of the extended vertical line segment of the S-shape set at the arcs mark the zeros, and the positive axes of the lower arc and the upper arc point respectively to the right and to the left. We show if the S-shape is constructed so that

1. two positive points unit distance apart on one arc see parts of the opposite arc that share a single point, and

2. any point on an arc sees, through the outgoing aperture, an interval of length Δ of the other arc;

then this obstacle set can represent BFS-digraphs of all connected outerplanar graphs.

We first investigate the structure of the set of obstacles to fulfill the conditions above. Let $\Delta \geq 2$ be at least the maximum outdegree in $\overrightarrow{G'}$. We define the *distance of two sets* to be the minimum of distances between any two of their respective points. Denote by w the aperture's width. Let s represent the vertical segment's length (in the S-shape), and let x represent the S-shape's distance from either arc.

Considering the arcs as lines, the first condition manifests if and only if $\frac{w}{1} = \frac{s+x}{s+2x}$, and the second condition holds if and only if $\frac{w}{\Delta} = \frac{x}{s+2x}$. These two equations require that $w = \frac{\Delta}{\Delta+1}$ and $s = (\Delta-1)x$, and we choose $x = 1$ to make things simple. We next show that the depicted set of obstacles represents any BFS-digraph. (Surely, the obstacle set depends on Δ which is conditioned on $\overrightarrow{G'}$, and to list vertex coordinates of the polygonal obstacles we would also need to know the maximum depth in $\overrightarrow{G'}$ as we will discuss, so strictly speaking we have an obstacle set *template*.)

Proposition 3.2.7. *The arrangement of five convex obstacles shown in Figure 3.3, represents BFS-digraphs of all connected outerplanar graphs.*

Proof. We give an algorithm to place the vertices of a connected BFS-digraph $\overrightarrow{G'}$ so that, together with the set of obstacles, they form an obstacle representation of $\overrightarrow{G'}$. We consider the two arcs in the obstacle set as lines; after all vertices are placed, we curve them a bit—just as much that they block visibilities among vertices on them. This way, we ignore visibilities among vertices on the same arc (when considered as a line) and show that the set of obstacles represents $\overrightarrow{G'}$.

Consider a drawing of $\overrightarrow{G'}$ that satisfies the conditions in Lemma 3.2.4. From now on, by $\overrightarrow{G'}$ we refer to this embedding. Place the root of $\overrightarrow{G'}$ at coordinate 1 of the lower arc. We get a representation of $\overrightarrow{G'_0}$, where $\overrightarrow{G'_i}$ denotes the induced subgraph of $\overrightarrow{G'}$ containing all vertices at depth at most i. Suppose $\overrightarrow{G'_i}$ is represented such that

1. all vertices at an even depth are placed on the lower arc, and all vertices at an odd depth are placed on the upper arc;

2. on each arc, vertices at different depths are *well separated*, i.e., arc intervals containing all vertices at the same depth are disjoint;

3. vertices of each depth preserve their ordering in $\overrightarrow{G'}$; and

4. every two consecutive vertices are one unit apart, with the possible exception of the rightmost left child and the leftmost right child of the same parent.

Note that by preserving the order, we mean if a vertex is to the left of some other vertex v in $\overrightarrow{G'}$, it gets a smaller coordinate than v when put on an arc.

Now, we describe how to add vertices at depth $i+1$ to obtain a representation of $\overrightarrow{G'_{i+1}}$ satisfying the conditions above. Let $v_{i,j}$ denote the j-th vertex at depth i and let $[a_{i,j}, b_{i,j}]$ denote the interval of the opposite arc that is visible from $v_{i,j}$ through the outgoing-aperture. For each vertex $v_{i,j}$ at depth i in the representation of $\overrightarrow{G'_i}$, we add its children on the opposite arc as follows:

- If $v_{i,j}$ has a common child with its immediate preceding sibling, put its leftmost child at $a_{i,j}$; otherwise, put the leftmost child at $a_{i,j} + \frac{1}{2}$.

- If $v_{i,j}$ has a common child with its immediate next sibling, put its rightmost child at $b_{i,j}$; otherwise, put the rightmost child at $b_{i,j} - \frac{1}{2}$.
- Put the remaining left children, preserving their ordering, after the leftmost one so that all left children are one unit apart.
- Put the remaining right children, preserving their ordering, before the rightmost one so that all right children are one unit apart.

Since every point sees an interval of length Δ, we know $b_{i,j} = a_{i,j} + \Delta$. Thus, as each vertex has at most Δ children, by performing the above algorithm, the rightmost child is placed after the leftmost one, and they are at least one unit apart. Therefore, all consecutive pairs of vertices are of distance one, except for the rightmost left child and leftmost right child of the same parent that can be farther apart. Moreover, we know every two points, which are one unit apart, have a common point-of-sight; that is, the greatest point-of-sight of the smaller point equals the smallest point-of-sight of the greater one. By Corollaries 3.2.5 and 3.2.6, we know that if two vertices have a common child, then they are consecutive; and they are not right and left children of the same parent. Therefore, the presented algorithm put vertices so that two vertices at depth i and $i+1$ are visible in the representation, if and only if they are connected in $\overrightarrow{G'}$. Conditions 1, 3, and 4 are surely satisfied after performing the algorithm. Since a vertex sees no other vertex except through the apertures, to complete the proof, what remains to be shown is that a vertex sees only its children through its outgoing aperture (and only its parent(s) through its incoming aperture). To that end, next we prove that Condition 2 is satisfied, namely that vertices at different depths are well separated: they lie in pairwise disjoint intervals.

Let I_0 denote the "interval" $[1,1]$ wherein the root is placed, and for every $i \geq 0$ let I_{i+1} denote the interval visible from I_i through the outgoing aperture. Since every vertex at depth i is in I_i, and I_i and I_{i+1} belong to different arcs, to prove Condition 2, it suffices to show that $I_i < I_{i+2}$ (i.e., every point in I_i has a smaller coordinate than every point in I_{i+2}) for every $i \geq 0$. If $I_i = [a,b]$, the structure of the obstacle set yields $I_{i+1} = [\Delta \times a, \Delta \times b + \Delta]$. Since $\Delta \geq 2$, this gives $I_0 < I_2$. By induction, we obtain that $I_i = [\Delta^i, 2\Delta^i + \sum_{j=1}^{i-1} \Delta^j]$ for every $i \geq 1$. Since $\Delta \geq 2$, for every $i \geq 1$ we have $2\Delta^i + \sum_{j=1}^{i-1} \Delta^j < 3\Delta^i < \Delta^{i+2}$, therefore, $I_i < I_{i+2}$.

Since we have previously ensured that a vertex v at depth i sees only its children through the outgoing aperture among all vertices at depth $i+1$, the well ordering of the intervals implies that v cannot see any other vertices through the outgoing aperture. By symmetry of sight, this implies that no vertex can see through its incoming aperture any vertex other than its parent(s).

This concludes the proof that we gave an obstacle representation of $\overrightarrow{G'}$. □

Adjusting the representation for general outerplanar graphs

We first show how to modify the representation of a connected BFS-digraph $\overrightarrow{G'}$ to accommodate its corresponding outerplanar graph G. We know that the underlying graph of $\overrightarrow{G'}$ and G are the same, except that $\overrightarrow{G'}$ has no edge connecting two vertices at the same depth. Since G is an outerplanar graph, the extra edges of G, if any, are such that they connect two consecutive vertices. Therefore, to allow existence of extra edges in the representation, we simply shave off the portion of the arc between their endpoints.

Now, we adapt this idea for disconnected outerplanar graphs. Let $C_1, C_2, \ldots, C_n$ be the components of a given outerplanar graph. Let $\overrightarrow{C'_i}$ be a BFS-digraph of C_i, as defined in Subsection 3.2.1. Let $\Delta \geq 2$ be at least the maximum outdegree among all BFS-digraphs, and construct the obstacle set template as before. Now, let L denote the maximum depth among all $\overrightarrow{C_i'}$. We declare I_0 to be the interval $[1,1]$ on one arc, and for every $i > 0$, we let I_i be the interval $[\Delta^i, 2\Delta^i + \sum_{j=1}^{i-1} \Delta^j]$ on the arc opposite to interval I_{i-1}. The modified algorithm for representing a disconnected outerplanar graph is as follows. For each $\overrightarrow{C_i'}$, put its root at an arbitrary place in $I_{(i-1)(L+2)}$. Then carry out the algorithm described in Subsection 3.2.1 to place all vertices of $\overrightarrow{C_i'}$ for every i. This ensures that no vertex in C_i' can see a vertex of C_j' for any $i \neq j$. We then shave off the arcs as necessary to provide visibility among vertices at the same depth where desired.

We obtain a representation for an arbitrary outerplanar graph, concluding the proof of Theorem 3.2.1. □

3.2.2 Lower bound on disjoint convex obstacle number of outerplanar graphs

For a rooted tree, we use the standard terminology—the depth of a vertex is its topological distance to the root, and the height of the tree is the maximum depth over all its vertices.

Proof of Theorem 3.2.2. Denote by $T_{k,h}$ the full complete k-ary tree with height h rooted at r. We will show that the disjoint convex obstacle number of $T_{k,3}$ is at least *four*, for k to be specified later. We say that two edges form a *crossing* if they meet at an internal point of both.

Lemma 3.2.8. *For every $m \in \mathbb{Z}^+$, there is a value of k such that $T_{k,2}$ has no m-convex obstacle representation without edge crossings.*

Proof. Denote by V_1 the set of vertices at depth 1, which comprises an independent set. Note that $|V_1| = k$. For any given s, we can find a subset $V' \subseteq V_1$ of size s (provided large enough $k = k(s)$) such that every non-edge with both endpoints in V' is blocked by

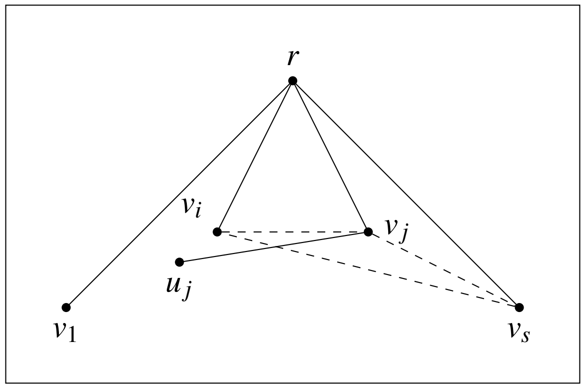

(a) To have $u_j < v_i$ without edge crossings, $v_j u_j$ must meet $v_i v_s$ as shown. But then, the three sides of the triangle $v_i v_j v_s$ cannot be blocked by the same convex obstacle O_1, which is a contradiction.

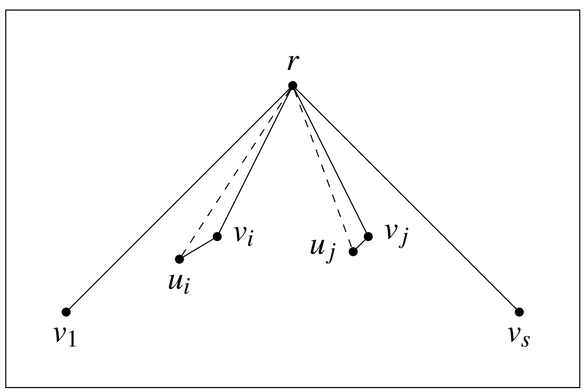

(b) Since $u_i < v_i < u_j < v_j$, no convex obstacle can meet both ru_i and ru_j without crossing an edge, so we have a contradiction to the assumption that *three* non-edges of the form ru_i can be blocked by the same obstacle O'.

Figure 3.4 – For the proof of Lemma 3.2.8. Since all non-edges among $v_1, v_2, \ldots, v_s$ are blocked by a single convex obstacle O_1, these vertices are in convex position and below r in the manner shown in both subfigures.

a common obstacle O_1. This is because we can assign every non-edge among V_1 to a single obstacle that blocks it to obtain an m-edge-coloring of a K_k induced on V_1, which by Ramsey's Theorem has a monochromatic clique of size s for large enough k. The set V' lies in some half-plane having r on its boundary, without loss of generality, below a horizontal line; otherwise, r would be inside a triangle with vertices in V', yet no single convex obstacle could block all three sides of it without meeting an edge of $T_{k,2}$. Let us write $u < v$ whenever the triple ruv is counterclockwise. Let $v_1 < v_2 < \ldots < v_s$ denote the vertices in V'. For each $i : 1 < i < s$, let u_i denote a child of v_i. We claim that at least $(s-2)/2$ (not necessarily disjoint) convex obstacles are required.

To prove the claim, assume for contradiction that some obstacle O' blocks three non-edges of the form ru_i. Then without loss of generality, for some pair $i < j$ such that ru_i and ru_j are blocked by O', both $u_i < v_i$ and $u_j < v_j$ hold. It must be that $v_i < u_j$; otherwise, $v_j u_j$ would cross an edge or meet O_1 which blocks both $v_i v_j$ and $v_j v_{j+1}$. See Figure 3.4(a). Choose two points $p_i \in ru_i \cap O'$ and $p_j \in ru_j \cap O'$. Then the segment $p_i p_j$ must intersect the union of the edges rv_i and $v_i u_i$. See Figure 3.4(b). By the convexity of O', we have a contradiction.

Thus, for $s \geq 2m+4$, at least $m+1$ convex obstacles are required if no edges cross. □

Assume for contradiction that we have a representation of $T_{k,3}$ with *three* pairwise disjoint convex obstacles O_1, O_2 and O_3.

If the endpoints of a crossing induced only the two edges (that is, an "X" type crossing), at least four convex obstacles would be needed to block the non-edges, since any convex set that intersects two non-edges must meet an edge. However, no more than three edges can be induced by four vertices without forcing a cycle. Therefore, the four endpoints of every crossing induce a path with three edges.

By Lemma 3.2.8, we know that for large enough k, there are crossings within each subtree of $T_{k,3}$ isomorphic to $T_{k,2}$. Pick three crossings c_1, c_2, and c_3 in $T_{k,3}$, each in a subtree rooted at a different neighbor of the root vertex. For each $i \in \{1,2,3\}$, denote by $u_i v_i$ and $w_i z_i$ the edges of the crossing c_i, with the corresponding induced path on four vertices $P_i = u_i v_i w_i z_i$.

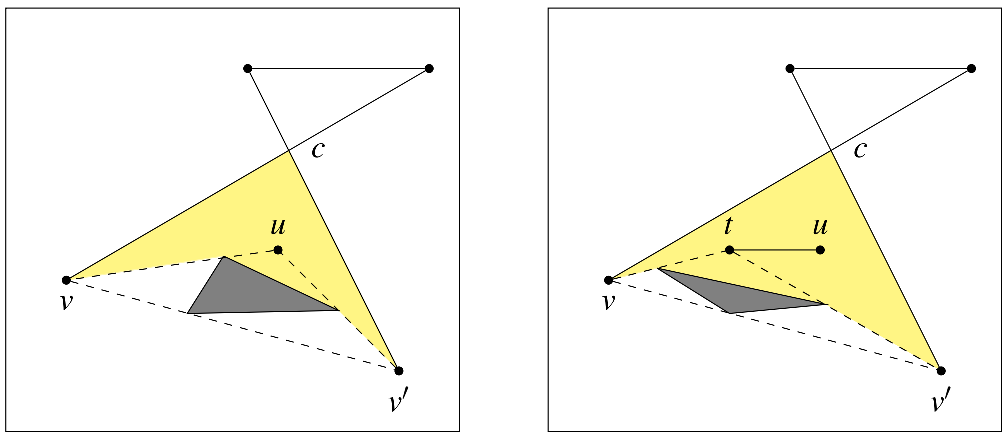

(a) Since the obstacle blocking vv' must also be responsible for blocking uv and uv', every neighbor of u must be inside the lightly shaded region inside triangle vcv'.

(b) But every neighbor t of u is subject to the same conditions as t! Hence, the neighbors of t must be inside the lightly shaded region inside triangle vcv'...

Figure 3.5 – Proof of Theorem 3.2.2: Case of P_1 and P_2 with non-disjoint convex hulls. Non-edges shown in each subfigure imply a respective minimal portion (dark gray) of an obstacle. The third edge of the path could have been incident on v or v' but this makes no difference. Only the obstacle that blocks vv' can be inside the convex angle $v'cv$.

Let us first consider the case where the convex hulls of two of these paths, say P_1 and P_2, meet. If this is the case, with no vertex of P_1 being inside the convex hull of P_2 or vice versa, then some edge of P_1 must intersect some edge of P_2, inducing an "X" type crossing which requires four obstacles. Hence, without loss of generality, some vertex u of P_1 is in the convex hull of P_2. Let c be the point of intersection of the two edges of P_2. Then, u

is inside some triangle vcv' where $v, v' \in P_2$. If vv' is an edge, then vcv' induces a bounded face, so uv would require an obstacle in addition to the three required by P_2. Now, since u is inside a triangle vcv', the obstacle blocking vv' must also block uv and uv', but this forces all neighbors of u to be inside vcv'. Applying this argument to the neighbors of u, which satisfy the same conditions as u, we see that P_1 must be completely inside vcv' (see Figure 3.5). But every non-edge of P_1 requires a distinct obstacle, at most one of which may coincide with one blocking vv' while none among them may coincide with any other obstacle, so five obstacles are required, a contradiction.

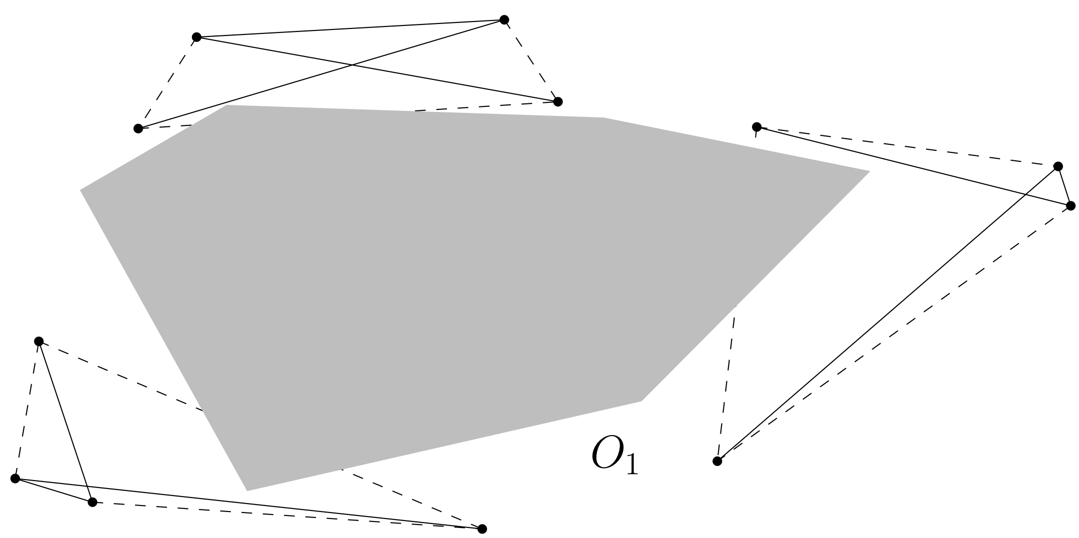

Figure 3.6 – Proof of Theorem 3.2.2: Case of P_1, P_2 and P_3 with pairwise disjoint convex hulls

This means that the convex hulls of P_1, P_2 and P_3 are pairwise disjoint. Recall that for each of these paths P_i, each of the three non-edges of P_i must be blocked by a unique obstacle among three pairwise disjoint obstacles. Hence by the Jordan Curve Theorem we get a contradiction (see Figure 3.6). □

3.2.3 Convex obstacle number of bipartite permutation graphs

A *permutation graph* is a graph on $[n]$ according to a permutation $(\sigma_1, \sigma_2, \ldots, \sigma_n)$ of $[n]$ such that there is an edge between two elements $\sigma_i > \sigma_j$ whenever $i < j$. We show that the idea of having a small aperture between two classes of vertices, which are placed close to two convex obstacles, is readily extended to the class of bipartite permutation graphs.

Proof of Theorem 3.2.3. By a result from [43], a bipartite graph $G(V,E)$ is a permutation graph if and only if its two independent vertex classes V_1 and V_2 can be ordered such that the neighborhood of every vertex $u_i \in V_1$ forms an interval $[a_i, b_i]$ in V_2, and if $u_i < u_j$ for two vertices in V then $a_i \leq a_j$ and $b_i \leq b_j$.

We illustrate in Figure 3.7 a set $\mathscr{C}$ of *four* disjoint convex obstacles allowing an obstacle representation of G. $\mathscr{C}$ consists of two convex arcs C_1 and C_2, and two vertical line segments (labeled $\mathscr{A}$) which form an aperture between C_1 and C_2.

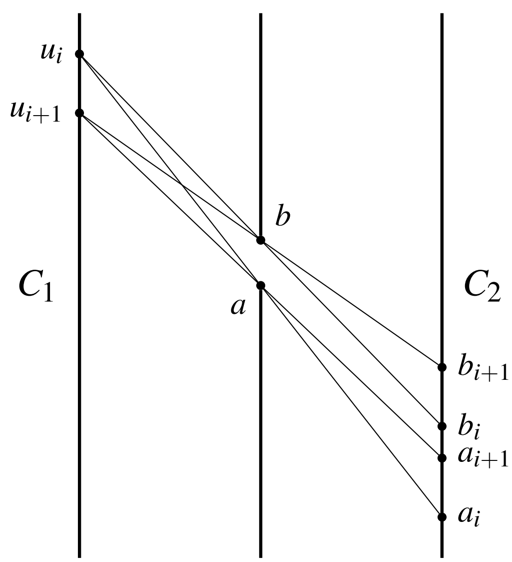

Figure 3.7 – Four obstacles allowing a representation of a bipartite permutation graph

Similar to the treatment of the arcs in Subsection 3.2.1, we regard C_1 and C_2 as line segments, except that they block visibilities among graph vertices placed near them. For convenience, we shall speak of placing vertices of G *on* these arcs.

We put vertices of V_1 and V_2 on C_1 and C_2 respectively. Let $u_1, u_2, \dots, u_n$ and $v_1, v_2, \dots, v_n$ be the ordering of the vertices in V_1 and V_2 guaranteed by the aforementioned result in [43]. We place the vertices, in order, inductively. In the basis step, we place u_1 arbitrarily on the relative interior of C_1. Let a_i and b_i denote the endpoints of the segment of C_2 that u_i can see through the aperture (see Figure 3.7). We place neighbors of u_1 in the relative interior of segment a_1b_1 on C_2 so that the order of their y-coordinates corresponds to their order in V_2.

At an inductive step $i+1$, where $i \geq 1$, we place the $(i+1)$-th vertex of V_1 together with its children, on the corresponding arcs as follows. We first find a consistent place for a_{i+1}. If the first neighbor w of the $(i+1)$-th vertex (with regards to the order in V_2) is already placed on C_2, we pick a_{i+1} so that it precedes w (with regards to y-coordinate) and succeeds a_i and any other point already placed on C_2. Otherwise, we pick a_{i+1} so that it succeeds b_i. We place u_{i+1} at the intersection of C_1 and the line through a_{i+1} and a (see Figure 3.7). The line through u_{i+1} and b intersects C_2 at point b_{i+1}, which has a higher y-coordinate than b_i. Therefore we can place neighbors of u_{i+1} that are not neighbors of u_i on the non-empty line segment b_ib_{i+1}. This concludes the proof of Theorem 3.2.3. □

3.3 Segment obstacle numbers of some graph families: disjoint and not-necessarily-disjoint cases

We define a *disjoint segment obstacle representation* of a graph G as a segment obstacle representation of G in which the obstacles are pairwise disjoint. Accordingly, we define the *disjoint segment obstacle number of* G as the fewest number of obstacles in a disjoint segment obstacle representation of G.

We compute the exact values of the segment obstacle number and disjoint segment obstacle number parameters for some simple graph families. For some families they turn out to be arbitrarily far apart, while for others they are—somewhat surprisingly—within a constant factor of one another.

3.3.1 Disjoint segment obstacle numbers

Theorem 3.3.1. *The null graph on* n *vertices (for every* $n \geq 1$*) has disjoint segment obstacle number* $n-1$*. More generally, this is the disjoint segment obstacle number of every disjoint union of any* n *complete graphs, and a lower bound for the disjoint segment obstacle number of every graph with an independent set of size* n*.*

Theorem 3.3.2. *The disjoint segment obstacle number of* C_n *is* $\lceil n/2 \rceil - 1$ *for* $n > 3$*, and* 0 *for* $n = 3$*.*

Theorem 3.3.3. *The disjoint segment obstacle number of* P_n*, the path graph on* n *vertices, is* $\lceil n/2 \rceil - 1$ *for every* $n \geq 1$*.*

For all of the above, we need the following lemma.

Lemma 3.3.4. *Every* m *disjoint segments (no two collinear) are part of the boundary in some decomposition of the plane into* exactly $m+1$ *convex cells.*

Proof of Lemma 3.3.4. Given any m disjoint segments comprising an obstacle set $\mathscr{S}$, we build a plane graph H that contains these segments and has $m+1$ convex faces. For simplicity of exposition, we consider the plane to include the point at infinity. Let H_0 be the graph on $2m+1$ vertices, the endpoints in $\mathscr{S}$ and the point at infinity; with m edges: the segments in $\mathscr{S}$. Initialize the plane graph H to H_0. For every segment $ab \in \mathscr{S}$, extend it until it meets a vertex or an internal point of an edge of H, so that the four points induce the path $a'abb'$ in H (which could be a cycle if a' and b' are both the point at infinity).

At the end of this process, H is connected—one way of seeing this is that the point at infinity is reachable **in** H from every vertex in $V(H)$ by following directions in $(-\pi/2, \pi/2]$. Further, every face of H is an intersection of half-planes and hence convex. In the iterative stage, $2m$ edges are added to H, in addition to as many additional edges as new vertices, since a new vertex is created if and only if an existing edge is subdivided. In other words, H has exactly $2m$ more new edges than it has new vertices. Therefore, by Euler's Polyhedral Formula, the number of faces of H is $|E(H)| - |V(H)| + 2 = |E(H_0)| - |V(H_0)| + 2m + 2 = m - (2m+1) + 2m + 2 = m+1$. □

Proof of Theorem 3.3.1. First we show that the disjoint segment obstacle number of every graph with an independent set of size n is at least $n-1$. For every n graph vertices together with fewer than $n-1$ obstacles in the plane in a disjoint segment obstacle representation, by Lemma 3.3.4 the obstacles can be extended to obtain a decomposition of the plane into fewer than n convex cells, so at least two of the graph vertices are inside the same convex cell. Hence, no obstacle could have been blocking their visibility path.

To represent the disjoint union of any n complete graphs, draw each complete graph arbitrarily near a distinct lattice point on the x-axis. Separate them with $n-1$ sufficiently tall vertical segment obstacles. The graph vertices together with the segment endpoints can be perturbed to attain simple position if desired. (That is, every connecting line has a distinct slope and every three lines have empty intersection.) A special case of such a graph is the null graph on n vertices. □

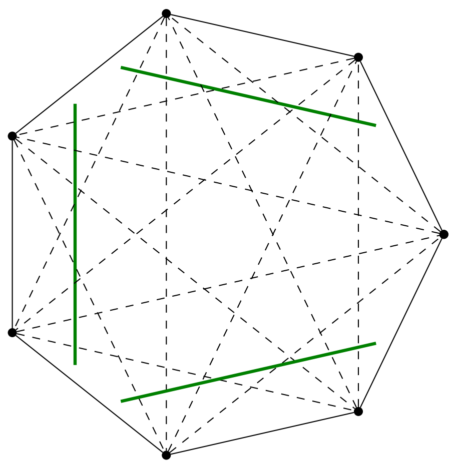

(a) C_7

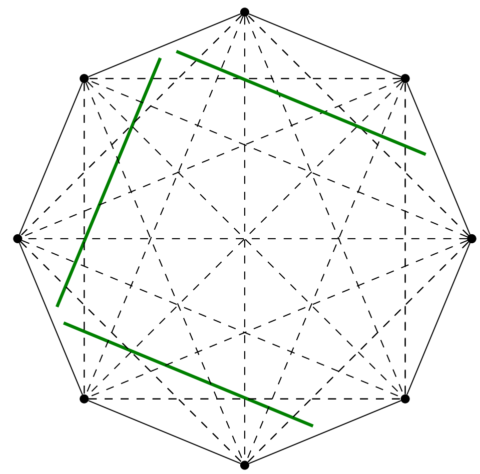

(b) C_8

Figure 3.8 – 3-disjoint segment obstacle representations of C_7 and C_8

Proof of Theorem 3.3.2. C_3 is a complete graph, so it has obstacle number 0.

For $n > 3$, since C_n has no C_3 subgraph, there can be at most two vertices in a cell in every convex decomposition of the plane with the obstacles on the boundary. By the pidgeonhole principle, the plane cannot be decomposed into fewer than $\lceil n/2 \rceil$ convex cells with the obstacles on the boundary, so by Lemma 3.3.4, every disjoint segment obstacle representation of C_n has at least $\lceil n/2 \rceil - 1$ obstacles.

Now we show that this is sufficient. Draw C_n on the plane as a regular n-gon with vertices $v_0, v_1, \ldots, v_{n-1}$. For every v_i, all the segments corresponding to non-edges having v_{i-1} or v_i as an endpoint can be blocked simultaneously by a sufficiently long segment obstacle inside the n-gon arbitrarily close to the edge $v_{i-1}v_i$ and parallel to it. For every positive even i less than n, place such a segment obstacle. See Fig. 3.8. This will suffice for blocking all non-edges of C_n. Since C_n has exactly $\lceil n/2 \rceil - 1$ vertices with non-zero even subscript, we have succeeded in using exactly $\lceil n/2 \rceil - 1$ disjoint segment obstacles. $\square$

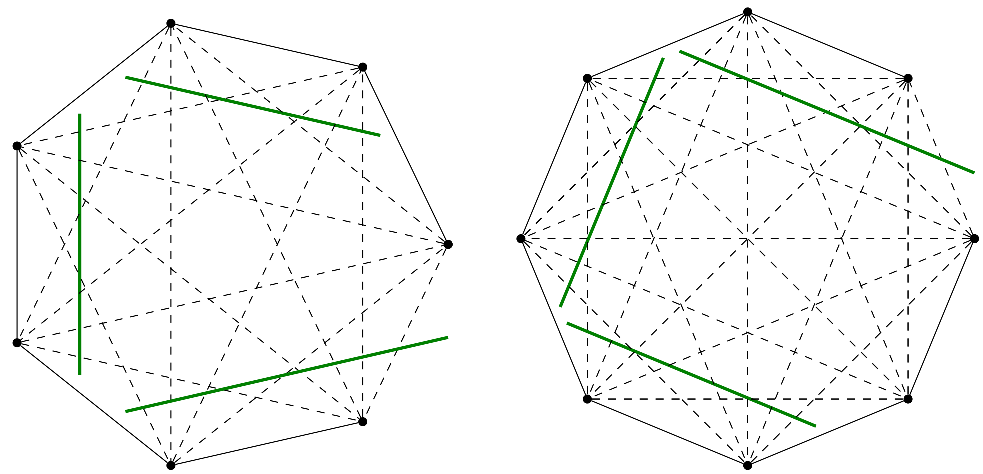

(a) P_7 (b) P_8

Figure 3.9 – 3-disjoint segment obstacle representations of P_7 and P_8

Proof of Theorem 3.3.3. The path graph P_n on n vertices has an independent set of size $\lceil n/2 \rceil$, so by Theorem 3.3.1 has disjoint segment obstacle number at least $\lceil n/2 \rceil - 1$.

To see that this is attainable, take an obstacle representation of C_n with $\lceil n/2 \rceil - 1$ disjoint segment obstacles as in the proof of Theorem 3.3.2. Now extend any one of the obstacles in a single direction, until it crosses an edge of the n-gon, giving us an obstacle representation of P_n. See Fig. 3.9.

Therefore, for every $n \geq 1$, the obstacle number of P_n is $\lceil n/2 \rceil - 1$. □

3.3.2 Segment obstacle numbers (not necessarily disjoint)

We now explore the case in which segments are allowed to intersect one another, using techniques similar to those in the more restricted case.

Theorem 3.3.5. *The null graph on n vertices (for every $n \geq 1$) has segment obstacle number $\lceil (\sqrt{8n-7} - 1)/2 \rceil$. More generally, this is the segment obstacle number of every disjoint* *union of any n complete graphs, and a lower bound for the segment obstacle number of every graph with an independent set of size n.*

Theorem 3.3.6. *The segment obstacle number of C_n is $\lceil n/4 \rceil$ for $n \geq 4$, and 0 for $n = 3$.*

Theorem 3.3.7. *The segment obstacle number of P_n is $\lceil n/4 \rceil$ for $n \geq 3$, and 0 for $n \leq 2$.*

First, we state the analog of Lemma 3.3.4 in this context that will be useful in the proof of Theorem 3.3.5.

Lemma 3.3.8. *Every m line segments in the plane can be extended to decompose the plane into at most $m(m+1)/2+1$ convex cells (they cannot be extended to decompose the plane into more convex cells). Equivalently, every arrangement of line segments that can be extended to decompose the plane into at least n convex cells consists of at least $s(n) := \lceil (\sqrt{8n-7}-1)/2 \rceil$ segments.*

Lemma 3.3.8 follows from the well-known result that m lines partition the plane into at most $m(m+1)/2+1$ convex cells and that this upper bound is attained for every simple arrangement of lines (one in which every pair of lines meet at a unique point, and every three lines have empty intersection). Lemma 3.3.9 complements this and is pertinent to how a graph component needs to be embedded in a segment obstacle representation.

Lemma 3.3.9. *For every plane arrangement of $m \geq 1$ line segments with k cells, there exists a decomposition of the plane into at most k' open convex cells disjoint from the segments of the arrangement, where $k' \leq k+2m-1$. Every cell in an arrangement of line segments can be decomposed into at most $2m$ such convex cells.*

The following corollary to Lemma 3.3.9 is almost immediate.

Corollary 3.3.10. *For a graph with an m-segment obstacle representation, the largest independent set chosen from the vertices of a single component has size at most* $2m$*. In particular, a* connected *graph with independence number* α *has segment obstacle number at least* $\alpha/2$*.*

Proof. Given a graph and an m-segment obstacle representation of it, consider any connected component of the graph. This component must be in a single cell of the arrangement of segment obstacles, otherwise some obstacle would meet the graph drawing. By Lemma 3.3.9, this cell can be decomposed into at most $2m$ convex cells by extending the segments, and by the pidgeonhole principle, there can be no independent set of size $2m+1$ in the cell. □

Proof of Theorem 3.3.5. First we show that the segment obstacle number of every graph with an independent set of size n is at least $s(n)$. For every n graph vertices together with fewer than $s(n)$ segment obstacles in the plane in a segment obstacle representation, by Lemma 3.3.8 and the pidgeonhole principle, at least two of the vertices are inside the same convex cell obtained by extending the segments. Hence, no obstacle could have been blocking their visibility path.

Now we give a segment obstacle representation for the disjoint union of any n complete graphs with $s(n)$ segment obstacles. Start with a simple arrangement of $s(n)$ line obstacles, so that they partition the plane into n convex cells. Embed each of the n complete graphs inside a distinct convex cell. Crop the line obstacles to segment obstacles by using a compact disk with all graph vertices inside as an intersection mask. Due to the convexity of the mask, every line obstacle corresponds to exactly one segment obstacle, furthermore, every non-edge uv continues to meet obstacles at exactly the same points that it did before. Therefore, we have a segment obstacle representation as desired. A special case of such a graph is the null graph on n vertices. □

Proof of Lemma 3.3.9. Given any m segments with endpoints in general position comprising an obstacle set $\mathscr{S}$ and defining k cells, we build a plane graph H containing these segments as in the proof of Lemma 3.3.4 and show that it has $k+2m-1$ convex faces. Denote by H_0 the graph of the arrangement of $\mathscr{S}$ together with an isolated vertex corresponding to the point at infinity. Initialize H to H_0. The iterative stage is similar to that of the proof of Lemma 3.3.4, based on extending every original segment until it meets a vertex or edge in H.

The final plane graph H is connected and every face of it is convex. Every new edge in the iterative stage must either join two components, or subdivide a face. Since H_0 has at least two connected components (some component including a segment in $\mathscr{S}$ since $m \geq 1$, and another consisting of the point at infinity) and H has exactly one component, this means that H has at most $2m-1$ more faces than H_0 does. In particular, a face of H_0 may become at most $2m$ cells of H. □

Remarks (pertaining to Lemma 3.3.9): There are arrangements of m line segments that permit the outside face to be decomposed into exactly $2m$ convex faces, i.e., connected arrangements of line segments in which every endpoint is on the outside face, such as any arrangement of pairwise intersecting line segments. For every bounded face f of an arrangement of line segments, at least three points are on the convex hull boundary of f and each is a point of intersection of pairs of line segments. By drawing tangents to the convex hull of f at such three points it can be seen that at least six segment endpoints are outside f, since each such tangent line to an extremal point of the convex hull of f will separate two distinct endpoints defining the extremal point from the convex hull of f. Hence, the maximum number of convex cells that a bounded face can be decomposed into by extending the segments is $2m-6$. This is attainable for certain arrangements of m segments in which three segments have pairwise intersection and the triangle defined by them contains all remaining segments.

Proof of Theorem 3.3.6. The disjoint segment obstacle number of C_n in the cases of $n=3$ and $n=4$ carry over verbatim from the proof of Theorem 3.3.2, since C_n has disjoint segment obstacle number at most 1 in these cases.

C_n has no C_3 subgraph, so no three vertices can be inside the same convex cell in any convex decomposition of the plane obtained by extending the segments in an obstacle representation of C_n. Since C_n is connected, in a segment obstacle representation, all its vertices must be in the same face defined by the obstacles, so by Lemma 3.3.9 and the generalized pidgeonhole principle, no m line segments can comprise the obstacle set in any obstacle representation of C_n for any $n>4m$. In other words, the segment obstacle number of C_n is at least $n/4$.

Now that we have established this lower bound, we explain how to obtain an $\lceil n/4 \rceil$-segment obstacle representation of C_n. To obtain a 1-segment obstacle representation of C_4, draw C_4 as a convex quadrilateral and place a segment obstacle inside the quadrilateral that meets both diagonals. Henceforth, we focus exclusively on the case of $n>4$.

Set $m := \lceil n/4 \rceil$ and $\theta := \pi/m$. We shall first give an m-segment obstacle representation of C_{4m}, and then argue that it can be modified to be an obstacle representation of C_n. Place m line segments $s_0, s_1, \ldots, s_{m-1}$ of length 2 on the plane, each symmetric about the origin, in counterclockwise order, starting with s_0 on the x-axis, such that the segment endpoints taken together constitute the vertices of a regular $2m$-gon. Denote by p_0 and p_1 the endpoints of m_0 and m_1 with coordinates $(1,0)$ and $(\sin\theta, \cos\theta)$, respectively. Pick ε small enough that the point $(1,\varepsilon)$ is strictly below the line through p_1 and perpendicular to m_1.

Let $u_0 := (1+\delta, \varepsilon)$ where $\delta > 0$ is a parameter to be determined later. Let v_0 be obtained by reflecting u_0 about the x-axis. For every $k \in \{1,2,\ldots,2m-1\}$, let u_k be obtained by rotating u_0 through an angle of θk about the origin, and let v_k be obtained by rotating v_0 through an angle of θk about the origin. Intuitively, u_k and v_k flank the same endpoint of some segment. For every $\delta > 0$ and every k, the segment $u_k v_k$ will not meet any obstacle (setting $\delta := 0$ is forbidden because it would put the pertinent segment endpoint between the vertices), and the segment $v_k u_{k+1}$ will not meet any obstacle either. The only problem, then, is that we may have too many undesirable visibilities. But δ can be chosen small enough that v_0 does not see u_1. Since $q := (1, \varepsilon)$ is strictly below the perpendicular line ℓ through p_1, the point q can be continuously moved to the right while applying a continuous clockwise rotation to ℓ while keeping q below ℓ. Setting δ thus will clearly ensure for every vertex u_k that the only two points visible from u_k are v_k and v_{k+1}. This is an m-segment obstacle representation of C_{4m}.

For the case of $n = 4m - r$ given any $m \geq 2$ and $r \in \{1,2,3\}$, we now explain how to obtain an m-obstacle representation of C_n from this m-segment obstacle representation of C_{4m}. Remove the first r vertices among u_0, u_1, u_2 and slightly shorten the first r obstacles among s_0, s_1, s_2 on their relevant sides to permit visibilities between the first r pairs among $v_0 v_1$, $v_1 v_2$ and $v_2 v_3$, but without introducing any additional visibilities. □

Proof of Theorem 3.3.7. For $n \geq 4$, the disjoint segment obstacle number of P_n has been computed to be at most one in the proof of Theorem 3.3.3, hence the corresponding values apply here as well. For $n > 4$, the lower bound for C_n applies for P_n as well, since P_n also contains no C_3 subgraph either and is yet connected. As for the upper bound, by extending one of the segments in an $\lceil n/4 \rceil$-obstacle representation of C_n in some direction, one of the edges of the cycle can be broken, resulting in an obstacle representation of P_n. □

3.4 Three-dimensional variants

We conclude this chapter with brief comments about three-dimensional variants.

Proposition 3.4.1. *In* three *dimensions, every graph can be represented with one obstacle that is a polygonal chain.*

Proof. Embed in $\mathbb{R}^3$ the vertices of a given graph G such that they are in general position, that is, no *four* are coplanar. This ensures that no edge of G is on the same plane as a non-edge of G. Label the non-edges of G as $s_1, s_2, \ldots, s_k$. Pick points $p_1, p_2, \ldots, p_k$ uniformly at random inside corresponding non-edges, and denote by C the polygonal chain $p_1 p_2 \ldots p_k$. By the general position condition, C will avoid all edges and vertices of G with probability 1 while meeting all non-edges by design. Therefore, $V(G)$ together with C (which can be perturbed to attain general position) constitute an obstacle representation of G in $\mathbb{R}^3$. □

Proposition 3.4.2. *In* three *dimensions, every* planar *graph can be represented with one convex obstacle.*

Proof. Given a planar graph G, triangulate a planar embedding of it to obtain the graph T. Now take a convex polyhedron C (no four vertices coplanar) with graph T. Let O be the convex hull of the set of midpoints of all pairs in $V(C)$ that do not correspond to edges in G. The point set $V(C)$ together with O (which can be perturbed to attain general position) constitute a 1-convex obstacle representation of G in three dimensions. □

Chapter 4

Obstacle representations: Algorithms and complexity

There are various interesting questions of algorithmic importance regarding obstacle representations of graphs. All the computations in the previous chapter exploited symmetries in the graphs and were for the most part ad hoc. Is it possible to devise computational procedures to compute, or hope to approximate, a minimum obstacle representation of a graph? What if we fix a drawing for the given graph? What if we further restrict these drawings to plane graphs? In this chapter, we give some answers to these questions.

As a negative result, we have that the problem of computing a minimum obstacle representation for a given graph drawing is NP-complete, even for plane graphs. As for positive results, we give two approximation algorithms to minimize the number of obstacles for a drawing of any graph, with approximation factors $O(\log n)$ and $O(\log OPT)$ for the general case: The former is made possible by showing a reduction to a finite hypergraph transversal problem, and the latter via bounding the Vapnik-Chervonenkis dimension of the relevant hypergraph family. We also show that the restriction of this problem to plane graphs is

Fixed Parameter Tractable (FPT) and even with weights admits a Polynomial-Time Approximation Scheme (PTAS). We conclude with some remarks about the complexity of the original minimization problem which is quantified over all drawings of an abstract graph.

4.1 Preliminaries for computing an obstacle representation of a graph drawing (ORGD)

Let D be a straight-line drawing of a (not necessarily planar) graph G on n vertices in the Euclidean plane with no three graph vertices on the same line. We refer to the open line segment between a pair of non-adjacent graph vertices as a *non-edge* of D. An *obstacle representation of* D is the set of vertices of G identified with their positions in D together with a collection of polygons (not necessarily convex) called *obstacles*, such that:

1. D does not meet any obstacle, and
2. every non-edge of D meets at least one obstacle.

The obstacle number of D is the least number of obstacles over all obstacle representations of D. We denote by *ORGD* the decision version of the *obstacle representation of a graph drawing* problem, namely: Given a graph drawing D with no three vertices on a line, and an integer k, do there exist s polygons that complete D to an obstacle representation that satisfies $s \leq k$?

4.1.1 Intersection hypergraphs and their transversals

A *hypergraph* is a pair $(X, \mathscr{F})$ in which X is a set of ground elements, and $\mathscr{F}$ is a collection of subsets of X. We introduce the following notation and terminology for intersection hypergraphs. Let X and Y be collections of sets. For each $y \in Y$, let $N(y) = \{x \in X \mid x \cap y \neq \varnothing\}$, and say that y *generates* $N(y)$. Let $\mathscr{F} = \{N(y) \mid y \in Y\}$. Then $(X, \mathscr{F})$ is an *intersection hypergraph*, which we shall denote by $H(X,Y)$ whenever convenient.[1] Similarly, for each $x \in X$, let $N(x) = \{y \in Y \mid x \cap y \neq \varnothing\}$, and say that x *generates* $N(x)$. Let $\mathscr{F}' = \{N(x) \mid x \in X\}$. The hypergraph $(Y, \mathscr{F}')$, which we shall denote by $H(Y,X)$ when convenient, is known as the *dual* of the hypergraph $H(X,Y)$.

A *transversal* of an intersection hypergraph $H(X,Y)$ is a subset $T \subseteq X$ such that every member of Y—that meets some member of X—meets some member of T. Let τ denote the minimum cardinality of a transversal of $H(X,Y)$. The (optimization version of) the hypergraph transversal problem is to compute τ exactly, and this has an equivalent formulation as the *set cover* problem. The decision version of the hypergraph transversal problem is NP-complete; indeed, the restriction to the case in which every member of Y meets exactly two members of X corresponds to a canonical NP-complete problem, "Vertex Cover."

4.1.2 Reducing ORGD to a $\mathrm{poly}(n)$ sized transversal problem

Theorem 4.1.1. *ORGD is in NP.*

Proof. Recall that for a given graph drawing D, we refer to a connected component of

[1] In many well-studied geometric hypergraphs $H(X,Y)$, each set in X is a singleton. The intersection of a member of X with a member of Y thus corresponds to the inclusion of the former in the latter. The members of Y are commonly referred to as *ranges*, especially in hypergraphs in which Y is a natural feature of the geometric space that the "points" of X belong to, e.g., the set of all half-spaces, all balls, or all axis-parallel boxes. We eschew the use of the term *range* since this is not the case for problems we are interested in, and also because our hypergraphs are defined by intersection not limited to inclusion: A set in X can meet two disjoint sets in Y and vice versa.

the complement of D as a *face* of the drawing. To rephrase an observation in [4] in our context, in an obstacle representation of D with cardinality $obs(D)$, there can be at most one obstacle per face, for otherwise obstacles in the same face could be merged, contradicting the minimality of $obs(D)$. Hence for any given graph drawing, each polygonal obstacle to be included in a minimal obstacle representation can be considered to be a face of the drawing. If need be, one can compute for each face a representative simple polygon that lies inside the face and meets every non-edge that the face meets. (This is not always a simple matter of perturbing the boundary of a face to lie inside the face, since a face may have holes and so its boundary may be a disconnected set.)

Since an n-vertex graph has fewer than n^2 edges (with $\Omega(n^2)$ edges attainable), its drawing must have fewer than n^4 faces (with $\Omega(n^4)$ faces attainable). Computing $obs(D)$ is therefore a matter of computing a transversal for the finite intersection hypergraph $H(X,Y)$ where X is the face set of D and Y is the non-edge set of D. Observe that $|X| < n^4$ and $Y < n^2$, with $|X| = \Omega(n^4)$ attainable simultaneously with $|Y| = \Omega(n^2)$. Using a representation of D with integer coordinates represented as signed integers, an incidence matrix representation of $H(X,Y)$ with fewer than n^8 bits (and possibly $\Omega(n^8)$) can be computed using standard techniques [10] in time polynomial in the number of input bits, and in time $\text{poly}(n)$ in the RAM model with unit-cost arithmetic operations. □

From now on, we will assume that the obstacles will be chosen from the $O(n^4)$ faces of the given drawing, which we shall refer to as canonical obstacles.

4.2 NP-hardness via reduction from planar vertex cover

Here we prove the following.

Corollary 4.2.1. *ORGD is NP-complete.*

A *plane graph* is a planar graph drawn on the plane without edge crossings and with no vertex on an edge. Denote by ORPeG the restriction of ORGD to plane graphs. Corollary 4.2.1 follows from Theorem 4.1.1 and the following.

Theorem 4.2.2. *ORPeG is NP-Hard.*

Recall that given a graph $G = (V, E)$, a vertex cover for G is a subset $U \subseteq V$ that contains at least one endpoint from every edge. We denote by *P-VC* the decision version of the vertex cover for planar graphs problem, namely: Given a planar graph $G = (V, E)$ and an integer k, does there exist a vertex cover U for G satisfying $|U| \leq k$? Garey, Johnson, and Stockmeyer have shown that P-VC is NP-complete [21]. We intend to show that ORGD is NP-hard by giving a polynomial-time reduction from P-VC to ORGD. Our reduction is inspired by the reduction given by Garey and Johnson [20] from P-VC to P-VC-3, the restriction of itself to maximum degree 3.

Proof. We reduce from planar vertex cover. We are given an abstract planar graph G having (without loss of generality) no isolated vertex, and a number k. Let $n = |V(G)|$, $m = |E(G)|$, and denote by f the number of faces in a crossing-free planar drawing of G. We will transform G in polynomial time into a plane graph $\tilde{G}$ in such a way that G has a vertex cover of size at most k if and only if $\tilde{G}$ has an obstacle representation of size at most $\tilde{k}$ (for $\tilde{k}$ defined below).

First, we construct from the planar vertex cover instance G a planar vertex cover problem instance G^3 of maximum degree 3, adapting and extending the construction of [20].

We show that G^3 admits a vertex cover of size at most $\tilde{k}$ if and only if G admits a vertex cover of size at most k. Second, we construct an ORPeG instance $\tilde{G}$ in such a way that an obstacle representation of $\tilde{G}$ will correspond to a vertex cover of G^3 of the same size, and vice versa.

Constructing the maximum degree 3 vertex cover instance G^3. The planar graph G^3 is constructed in the following way. We transform each vertex v_i of G into a cycle C^i of length $2b_i$, with $b_i \in \deg(v_i) + \{0,1,2\}$ (with the exact value decided below). We color the vertices of C^i alternating between blue and red. We then create a single leaf vertex z_i adjacent to some arbitrary red vertex of C^i. We transform each edge (v_i, v_j) of G into a path P_{ij} with *three* edges whose endpoints are *distinct* blue vertices of C^i and C^j. We finally create f copies of the 3-vertex path graph P_3, each constituting a component of G^3.

Now we show that G has a vertex cover of size at most k if and only if G^3 has a vertex cover of size at most $k' = k + f + m + \sum_i b_i$.

($\Rightarrow$): For each vertex v_i in a given vertex cover for G of size k, we select z_i and all the blue vertices of C^i, thus including an endpoint of each path P_{ij}; and for each v_i not in the cover, we select all the red vertices of C^i (a total so far of $k + \sum_i b_i$ vertices). Since for every path P_{ij} at least one of the cycles C^i and C^j will have all its blue vertices chosen, thus including at least one endpoint of P_{ij}, choosing one internal vertex from each P_{ij} (m more), and the central vertex of each P_3 (f more) suffices to complete a size k' vertex cover for G^3.

($\Leftarrow$): Given a vertex cover for G^3 of size k', we obtain a canonical vertex cover for G^3 of size $k'' \leq k'$ in the following way. Each copy of P_3 contributes at least one vertex to a cover, so have it contribute exactly its central vertex, for a total of f vertices. Each path P_{ij} contributes at least one of its internal vertices to cover its central edge. If both internal vertices of a path P_{ij} are in the given cover, take one internal vertex out and ensure that its blue neighbor is in, which makes for m internal vertices from these paths. Note that every cycle C^i contributes at least b_i vertices, lest some edge of the cycle be uncovered. This holds with equality only if C^i contributes (including 'its' z_i) exactly its red vertices—in that case, do nothing. Otherwise, ensure that C^i contributes exactly $1+b_i$ vertices: 'its' z_i and its blue vertices. Since every remaining vertex of C^i is adjacent in G^3 only to blue vertices of C^i and possibly z_i, this swap will not ruin the cover. Denote by k'' the size of this resulting canonical vertex cover. The cycles in G^3 contributing blue vertices therefore correspond to a vertex cover for G of size $k'' - f - m - \sum_i b_i \leq k' - f - m - \sum_i b_i = k$.

Constructing the ORGD instance $\tilde{G}$. In the remainder of the proof, we show how to "implement" the graph G^3 as an equivalent ORPG problem instance. The basic building blocks of the construction are *empty triangles* and *diamonds*. An *empty triangle* is a face of a plane graph that is surrounded by three edges and has no vertex inside. A *diamond* consists of two empty triangles sharing an edge and having their *four* vertices in convex position. Observe that a diamond contains a non-edge between two of its vertices. Hence at least one empty triangle of every diamond must be chosen in an obstacle representation. The f copies of P_3 in G^3 will match the faces of $\tilde{G}$ besides empty triangles, all of which must be chosen. The remaining vertices of G^3 will match the empty triangles of $\tilde{G}$, such that the edges among them match the diamonds of $\tilde{G}$. Hence there will be a natural bijection between minimum vertex covers of G^3 and minimum obstacle representations of $\tilde{G}$.

To begin the construction, we use the linear-time algorithm of de Fraysseix, Pach, and Pollack [11] to obtain a planar imbedding of G on a $O(n) \times O(n)$ portion of the integer lattice and then perturb the coordinates to obtain general position. (We do not distinguish between G and this imbedding.) We first visualize $\tilde{G}$ as a *bold drawing* [46] of G, whose vertices are represented by small disks and edges by solid rectangles: we draw each vertex

u_i of G as a disk D_i about u_i (with boundary $\tilde{C}^i$), and every edge u_iu_j as a solid rectangle R_{ij}. See Fig. 4.1(a). Each R_{ij} has two vertices t_{ij},v_{ij} on $\tilde{C}^i$ and two vertices t_{ji}, v_{ji} on $\tilde{C}^j$ such that the line u_iu_j is a midline of R_{ij}, and $t_{ij}u_iv_{ij}t_{ji}u_jv_{ji}$ is a counterclockwise ordering of the vertices of a convex hexagon.

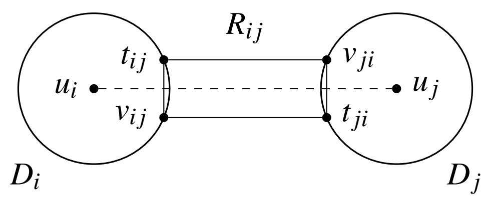

(a) Bold drawing of vertices u_i, u_j and edge u_iu_j

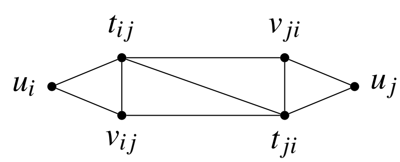

(b) Edge gadget in $\tilde{G}$ for edge u_iu_j of G

Figure 4.1 – Bold drawing and edge gadget for an edge of G

We draw the disks small enough to ensure that they are well-separated from one another. We set the radius r of every disk to the smaller of 1/4 and half of the minimum distance between a vertex u_i and an edge u_ju_k ($j \neq i \neq k$) of G. To fix a single width for all rectangles (i.e., $||t_{ij} - v_{ij}||$), we set a global angle measure α to the smaller of 45° and half of the smallest angle between two edges of $E(G)$ incident on the same vertex of $V(G)$.

$\tilde{G}$ is modeled on the bold drawing, by implementing each edge of G (path P_{ij} of G^3) with an edge gadget and each vertex of G (cycle C^i of G^3) with a vertex gadget. The edge gadget, consisting of four triangles forming three diamonds, is shown in Fig. 4.1(b). (Note that each pair $v_{ij}v_{ji}$ defines a non-edge.)

The vertex gadget is a modified wheel graph whose triangles correspond to the vertices of cycles C^i in G^3 (see Fig. 4.2(c)). On every circle $\tilde{C}^i$, for every edge u_iu_j in G, we color blue the arc of measure α centered about the intersection of circle $\tilde{C}^i$ with u_iu_j. We place t_{ij} and v_{ij} at the endpoints of this arc so that $t_{ij}u_iv_{ij}$ is a counterclockwise triple. By the choice of α, all blue arcs are well-separated, and hence the rectangles are well-separated from one another and from other disks, by the choice of r. We color the remaining arcs red to obtain a red-blue striped pattern on each circle $\tilde{C}^i$, corresponding in color to the vertices of the corresponding C^i in G^3.

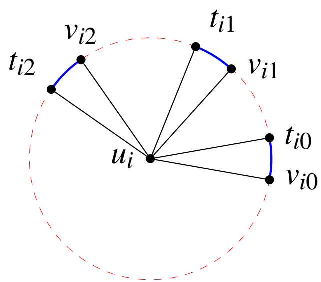

(a) Initial circle with blue (solid) arcs of measure α and red (dashed) arcs has a large red arc of measure in $[180° - \alpha, 270°)$.

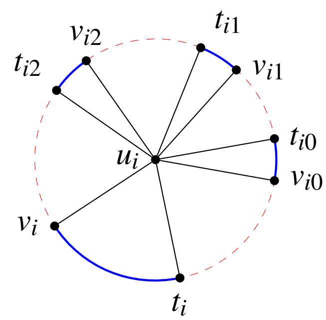

(b) After subdividing the large red arc into *three*, coloring its middle part blue, and adding dummy vertices v_i and t_i

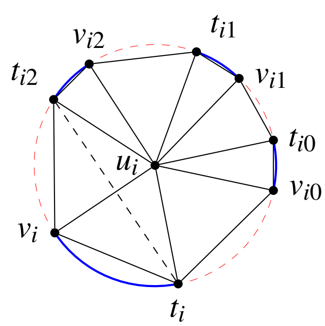

(c) In the resulting wheel graph, each pair of triangles sharing an edge induce a diamond. One diamond's non-edge is shown dashed.

Figure 4.2 – Constructing the wheel graph drawing in the case of a large red arc

On every circle $\tilde{C}^i$, we intend to add the remaining edges between consecutive vertices on $\tilde{C}^i$ to complete the union of the triangles $t_{ij}u_iv_{ij}$, forming a wheel graph on hub u_i, such

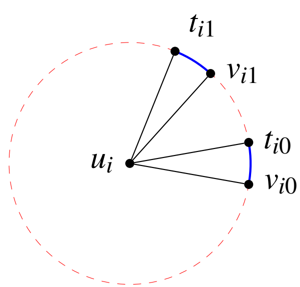

(a) An initial circle with blue (solid) and red (dashed) arcs has a very large red arc, of measure at least 270°.

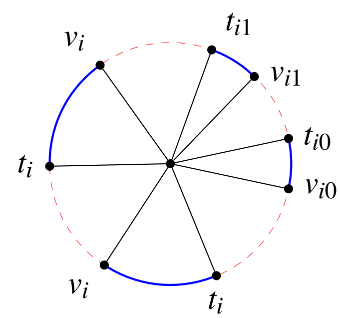

(b) After subdividing the very large red arc into *five*, coloring its second and fourth parts blue, and adding v_i and t_i dummy vertices

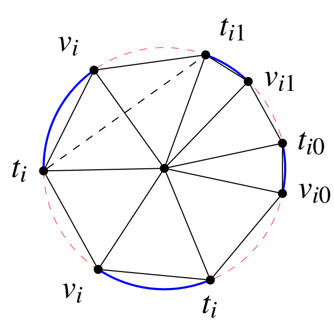

(c) In the resulting wheel graph, each pair of triangles sharing an edge induce a diamond. One diamond's non-edge is shown dashed.

Figure 4.3 – Constructing the wheel graph drawing in the case of a *very* large red arc

that every pair of triangles that share a spoke form a diamond. If a red arc has measure at least $180° - \alpha$, however, we must add additional spokes. By the general position assumption, at most one red arc per wheel can have such great measure. If such a red arc has measure less than $270°$, we divide it evenly into *three* parts and color the middle part blue (see Fig. 4.2); otherwise, we divide it evenly into *five* parts and color the second and fourth parts blue (see Fig. 4.3), maintaining the striped pattern in both cases. We place dummy vertices t_i and v_i [2] at the newly created (*zero*, *two*, or *four*) arc endpoints. We add the requisite edges to complete the wheel graph.

We place a vertex $\tilde{z}_i$ on an arbitrary red arc of $\tilde{C}^i$, and connect it in $\tilde{G}$ to the end vertices (say t_{ij} and v_{ik}) of that arc. Thus an empty triangle $t_{ij}z_iv_{ik}$ is formed in $\tilde{G}$ as part of a diamond with u_i, corresponding to z_i and its incident edge in G^3.

In the unbounded face of $\tilde{G}$ we place two isolated vertices inducing a non-edge inside the unbounded face, thus requiring this face to be chosen in any solution. Every non-triangular face of $\tilde{G}$ must be selected as an obstacle, since every simple polygon with at least 4 vertices has an internal diagonal (i.e., a non-edge). The selection of these faces are forced moves and correspond to the selection, in a vertex cover for G^3, of the central vertex of each P_3.

This completes the construction of $\tilde{G}$. Since each pair of neighboring triangles in $\tilde{G}$ indeed form a diamond and every non-triangular face is indeed a forced move, the result follows. □

Remark 4.2.1. To represent coordinates *exactly* as described would require a very permissive unit-cost RAM model of computation in which it is possible to represent real numbers and perform arithmetic and trigonometric functions in unit time. The reduction above can

[2] Dummy vertices have no adjacencies with any vertices outside of D_i.

be modified in such a way that each vertex position of $\tilde{G}$ is represented using $O(\log n)$ bits.

4.3 Approximation algorithms for ORGD

It is well-known that the greedy algorithm for the hypergraph transversal problem, which iteratively adds to an initially empty set T a member $x \in X$ that meets the largest number of sets $y \in Y$ that do not already meet some $x' \in T$, provides a $O(\log |Y|)$-factor approximation [48]. A simple example shows that the approximation ratio is tight for this greedy algorithm [48]. Thus we have a natural $O(\log n)$-factor approximation algorithm for our problem of computing the obstacle number for a given drawing.

4.3.1 The Vapnik-Chervonenkis dimension of a hypergraph

How about doing better? Lund and Yannakakis [29] showed that the general hypergraph transversal problem cannot be approximated in polynomial time within ratio $c \lg |Y|$ for any constant $c < 1/4$ unless NP is contained in $DTIME(n^{\text{poly } \log n})$. Raz and Safra [41] showed that unless P = NP, it cannot be approximated in polynomial time within ratio $c \ln |Y|$ for any constant $c < 1/8$. Alon, Moshkovitz and Safra [3] improved the hardness factor to more than $0.2267 \ln |Y|$ under this weaker complexity theoretic assumption.

But it is well-known [35, 48] that if every member of X meets at most Δ members of Y, then the greedy algorithm attains an approximation ratio of $O(\log \Delta)$. Unfortunately, this does not make our task easier, since it is seen that a face could meet $\Omega(n^2)$ non-edges by considering any drawing of the null graph on n vertices. Nonetheless, many families of hypergraphs arising in geometric settings lend themselves to algorithms with approximation ratios that do not depend on either $|X|$ or $|Y|$ using the following ideas.

In the context of an intersection hypergraph $H(X,Y)$, a set $S \subseteq X$ is said to be *shattered* if for every $A \subseteq S$ there is some $y \in Y$ such that $S \cap N(y) = A$. The size of a largest shattered set is called the *V-C dimension* of $H(X,Y)$, after Vapnik and Chervonenkis who first defined it in [47]. For some hypergraphs in which X and Y are both infinite, the V-C dimension is undefined and said to be infinite. Furthermore, for a family of hypergraphs of the form $H(X,Y)$, even if each hypergraph has finite V-C dimension, there may exist no absolute constant upper bound for the V-C dimension. If there an integer d such that every hypergraph in that family has V-C dimension at most d, we say that the family has bounded V-C dimension.

Let $w : 2^X \to [0,1]$ be an additive weight function with $w(X) = 1$. For a given $\varepsilon > 0$, an ε-net (with respect to w) is a set $S \subseteq X$ that is a transversal of $H(X,Y_\varepsilon)$, where $Y_\varepsilon \subseteq Y$ consists exactly of those members of Y each of which generates a subset of X with weight at least ε. Haussler and Welzl have shown [27] that if the V-C dimension of $H(X,Y)$ is some integer d, then for every $\varepsilon > 0$ there is an ε-net of size at most $cd\varepsilon^{-1}\ln\varepsilon^{-1}$, where c is a small constant. This is remarkable because the size of an ε-net bears no relation to the sizes of X or Y. See also [30,35].

Based on this observation, various—deterministic as well as randomized—efficient algorithms have been presented [8,12,13,17] to compute a transversal of size within a tiny constant factor of $d\tau\ln\tau$, where τ denotes the size of an optimum transversal. In Subsection 4.3.3, we formulate a specific approximation algorithm for ORGD. For now, we suffice it to say that bounding the V-C dimension for the corresponding hypergraph implies an efficient algorithm with approximation ratio independent of $|X|$ or $|Y|$ (and hence n).

Before we proceed, we state an important fact that we make immediate use of. It is well-known [30] that if the V-C dimension of $H(Y,X)$ is d, then the V-C dimension of $H(X,Y)$

is at most 2^{d+1}. The V-C dimension of a family of hypergraphs is therefore bounded if and only if the V-C dimension of the family of dual hypergraphs is bounded.

4.3.2 Bounding the V-C dimension

We show that the V-C dimension is bounded for the family of hypergraphs of the form $H(Y,X)$ where Y is the set of non-edges of a graph drawing and X is the set of faces of that drawing. This implies that the V-C dimension of $H(X,Y)$ (the hypergraph for the transversal problem at hand) is also bounded.

Theorem 4.3.1. *The V-C dimension is bounded for the family of hypergraphs of the form $H(Y,X)$ in which Y is the set of non-edges in a straight-line drawing D of a graph, and X is the set of faces of D.*

We can replace each face $x \in X$ by a simple path x' inside x that meets every non-edge that x meets and does not meet any non-edges that x does not. This substitution will not alter the hypergraph structure, and the resulting paths will be pairwise disjoint. Hence Theorem 4.3.1 is implied by the following result.

Theorem 4.3.2. *The V-C dimension is bounded for the family of hypergraphs of the form $H(Y,X)$ in which Y is a set of line segments (with every pair meeting at a single point or not at all) and X is a set of pairwise disjoint simple paths.*

Proof. Assume for contradiction that for every m, there is a hypergraph $H(Y,X)$ such that Y is a set of m line segments (with every pair meeting at a single point or not at all), X is a set of pairwise disjoint paths, and Y is shattered.

Given m, and a pair (Y,X) such that $|Y| = m$ and X shatters Y, let $X_3 \subseteq X$ be a minimal set of paths that generate all the $\binom{m}{3}$ *triples* in Y. That is, every path in X_3 meets exactly three segments in Y, and for every three segments $i, j, k \in Y$ exactly one path $\pi_{ijk} \in X_3$ meets all three. To keep the following argument simple, without loss of generality, no $\pi \in X_3$ meets any intersection points among the segments, of which there are at most $\binom{m}{2} = O(m^2)$. If there were such paths in X_3, we could remove them from X_3 and still be left with at least $\binom{m}{3} - \binom{m}{2} = \Omega(m^3)$ paths. We will now charge each path $\pi \in X_3$ to a line segment in Y,

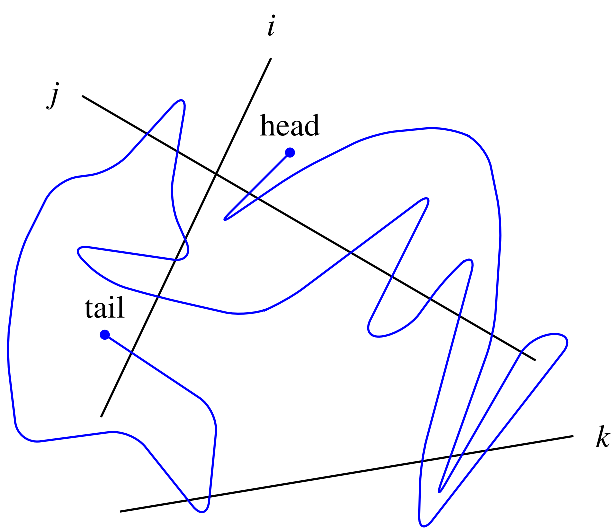

Figure 4.4 – Example of an original path π_{ijk} meeting segments i, j, and k

intuitively, "the middle segment" that it meets. Nothing prevents such a path from going back and forth between three segments, so we need to define this more carefully. For a given path $\pi = \pi_{ijk}$ that meets segments $i, j, k \in Y$ (see Fig. 4.4), arbitrarily label one end of the path as the tail and the other as the head. Starting from the tail, erase π as long as it still meets all three segments, and stop erasing when erasing any longer would cause the remaining path to intersect fewer segments. The tail is now on one of the three segments, without loss of generality, i (see Fig. 4.5). Note that the path does not intersect i anywhere else but the tail. Now start erasing π from the head in a similar fashion, and stop erasing when erasing any more would cause the remaining path to intersect fewer than

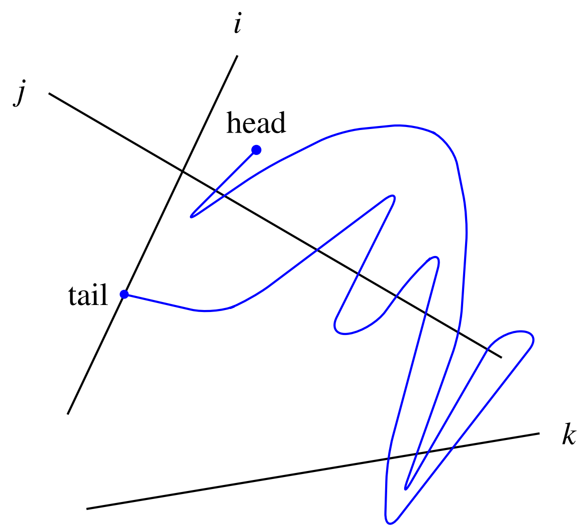

Figure 4.5 – The interim path obtained by erasing from the tail

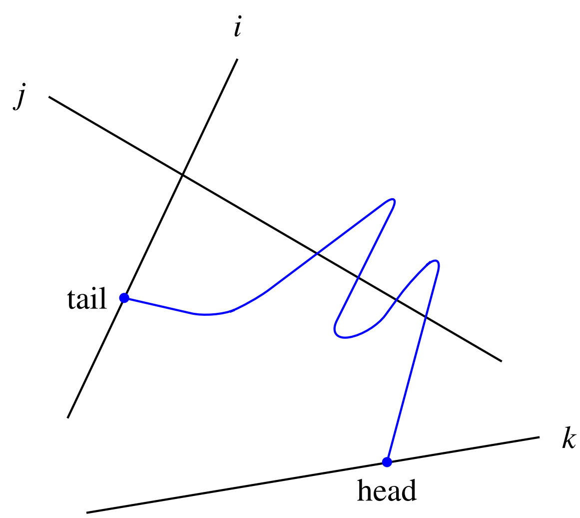

Figure 4.6 – The final path obtained by erasing from the head too, which will be charged to segment j

three segments. Again, the head must be on some segment when we stop, but it could not be on i by the above observation (see Fig. 4.6). Without loss of generality, the head is on the segment k. Now notice that the path does not meet k anywhere else but the head. Without changing the shatter property, let us replace π_{ijk} with this shorter version of its former self: starting at i, meeting j but no other segment one or more times, before it ends at k. We charge π_{ijk} to j.

Let $\hat{s}$ be a segment that accumulated the greatest charge at the end of this process. Since at least $\Omega(m^3)$ paths were charged to at most m segments, $\hat{s}$ was charged by at least $\Omega(m^2)$ paths. Let X' denote the set of paths in X_3 that were charged to $\hat{s}$.

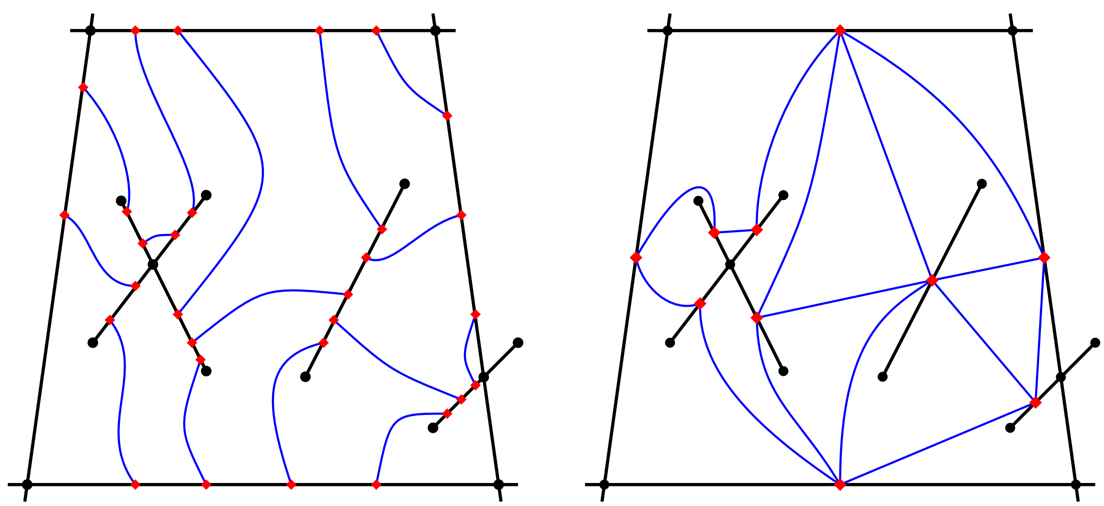

(a) With clipped paths from the original X' (b) With path endpoints on each edge of $\mathscr{A}$ merged

Figure 4.7 – Example of a cell of $\mathscr{A}$: vertices of $\mathscr{A}$ indicated as black disks, paths from X' drawn blue and their endpoints as red diamonds

Denote by $\mathscr{A}$ the arrangement of the line segments in $Y \setminus \{\hat{s}\}$. Note that every path in X' starts at an edge of $\mathscr{A}$, ends at another edge of $\mathscr{A}$, and its interior is fully contained in a cell of $\mathscr{A}$ (see Fig. 4.7(a)). Let $G(X')$ be the graph with the endpoints of the paths as the vertex set, and the interiors of the paths as edges. Since the paths are pairwise disjoint, clearly, $G(X')$ is planar. Now for each edge of the arrangement, merge the path endpoints on that edge at the midpoint of the edge while making sure that the paths remain pairwise interior-disjoint (see Fig. 4.7(b)). Recall that by Euler's formula a planar graph on n vertices has at most $3n-6$ edges. It is not clear that we have a contradiction yet, since $\mathscr{A}$ may have up to m^2 edges (attained when every pair among the m segments cross). Hence it seems that $G(X')$ may have $\Theta(m^2)$ vertices, so it is plausible that $G(X')$ has $\Theta(m^2)$ edges.

However, each edge of $G(X')$ must be contained in a single cell that meets $\hat{s}$. Might this mean that $G(X')$ has $o(m^2)$ vertices? We know that every vertex of $G(X')$ corresponds to a distinct edge of $\mathscr{A}$ in a cell of $\mathscr{A}$ that meets $\hat{s}$. Hence, the complexity of the zone[3] of $\hat{s}$ is an upper bound on $|V(G(X'))|$. We present a lemma regarding the complexity of the zone of a line segment in an arrangement of line segments.

Lemma 4.3.1 (B. Aronov, personal communication)**.** *Let $\mathscr{A}$ be an arrangement of n line segments, and let s be another line segment. The zone of s has complexity $O(n\alpha(n))$ where α denotes the very slow growing inverse Ackerman function.*

Proof. Let the shape s' be obtained by enlarging s (e.g. taking the Minkowski sum of s with a small enough disk) such that s' meets no vertex of $\mathscr{A}$ that s does not. Obtain a new arrangement $\mathscr{A}'$ of line segments by erasing s' from $\mathscr{A}$. Doing this will possibly disconnect some original line segments that define $\mathscr{A}$ into two, ending up with an arrangement $\mathscr{A}'$ of at most $2n$ line segments. Every point of s' is in the same cell of this new arrangement. Every cell in an arrangement of m line segments has complexity $O(m\alpha(m))$ [37]. Hence, every cell of $\mathscr{A}'$ has complexity at most $O(2n\alpha(2n))$, i.e., $O(n\alpha(n))$, including the cell

[3] Recall that the *complexity* of a cell of an arrangement is the number of edges of the arrangement that are incident to it. The *zone* of a segment is the set of cells that it meets, and the complexity of the zone of a segment is the number of edges incident to all the cells that it meets.

that s is in. Since every edge bordering a cell of $\mathscr{A}$ that s meets corresponds to one or two edges bordering the cell of $\mathscr{A}'$ that s is in, the complexity of the zone of s in $\mathscr{A}$ is at most $O(n\alpha(n))$. □

By Lemma 4.3.1, the zone of $\hat{s}$ has complexity $o(m^2)$, hence the number of vertices of the planar graph $G(X')$ is $o(m^2)$. But since $G(X')$ has $\Omega(m^2)$ edges, a contradiction is reached due to Euler's Formula for a large enough value of m. Therefore, the V-C dimension is bounded for the family of hypergraphs $H(Y,X)$ in which Y is a set of line segments and X is a set of pairwise disjoint paths. □

4.3.3 A concrete randomized $O(\log \tau)$-approximation algorithm

We continue using the terminology of having been given a drawing D of an n-vertex graph, with X as the set of faces in D and Y as the set of non-edges in D. For the rest of this section, we assume that D has been processed into the incidence matrix of the hypergraph $H(X,Y)$ with $O(n^8)$ entries in memory, as discussed in Section 4.1.

We present a randomized algorithm, Algorithm 2, modeled on those presented by Efrat and Har-Peled [12] and Efrat et al. [13], which rely on the analyses of Clarkson [9] and Brönniman-Goodrich [8]. Their randomized algorithms for computing transversals of unrelated hypergraphs run in low polynomial time in n and return a transversal of size at most $O(\tau \log \tau)$ with high probability. Algorithm 2 uses dovetailing to vary not only k (essentially a guess for τ) but also d (essentially a guess for the V-C dimension), whereas the aforementioned papers implicitly or explicitly use a known upper bound on the V-C dimension for d. Algorithm 2 simulates a pass through the algorithm of [12] for each possible value of d rounded to the nearest power of two, and similarly compresses the space of k for added efficiency.

```
ComputeObstacleRepresentation(set_of_faces X, set_of_nonedges Y) {
  for (deekay = 2; ; deekay = 2 * deekay) {
    for (k = 2;  k ≤ deekay;  k = 2 * k) {
      d = deekay / k;
      Assign weight 1 to each face in X;
      for (i = 1;  i ≤ numRounds(k, |X|);  i = i + 1) {

          • Pick randomly a set S of sampleSize(d,k) obstacles, choosing each
            obstacle randomly and independently from the face set X
            according to their weights.

          • Check if the obstacles in S together meet all of the non-edges in Y;
            if so, return the transversal S.

          • Else, find a non-edge y that does not meet any obstacle in S, and let
            N(y) be the set of faces in X that the non-edge y meets.

          • Compute ω, the sum of weights of faces in N(y).
            If 2kω ≤ the sum of weights of all faces in X,
            double the weight of every face in N(y).

      } // end for i
    } // end for k
  } // end for deekay (dovetailing diagonal)
}
```

Algorithm 2: A randomized $O(\log \tau)$-approximation algorithm for ORGD stated using C-style pseudocode. The expression sampleSize(d, k) stands for $\lceil 2\text{dk}\lg \text{k} \rceil$, and the expression numRounds(k, $|X|$) stands for $\lceil 8\text{k}\lg |X| \rceil$.

We claim that Algorithm 2 returns a transversal of size at most $O(\tau \log \tau)$ with high probability. We show this by means of arguing backward compatibility with the one in [12]. Imagine a matrix with the value of d doubling as we go one column to the right, and the value of k doubling as we go one row down. Algorithm 2 processes one diagonal at a time, starting from the top left, going through each diagonal in the down and left direction.

Denote by $\hat{d}$ the V-C dimension of the given hypergraph, and by $\hat{k}$ be the minimum value of k for which the algorithm in [12] provides an approximation guarantee of $O(\tau \log \tau)$ with high probability, terminating in polynomial time. Without loss of generality, $\hat{d}$ and $\hat{k}$ are powers of two. Suppose that Algorithm 2 terminates for values of d' and k' for which $\mathit{sampleSize}(d', k') > \mathit{sampleSize}(\hat{d}, \hat{k})$. With high probability, the algorithm did not yet consider the pair $(\hat{d}, \hat{k})$, so it terminated before it even got to the diagonal of the pair $(\hat{d}, \hat{k})$ predicated on the product $\hat{d}\hat{k}$. Hence we have $d'k' < \hat{d}\hat{k}$, and since $\mathit{sampleSize}(d', k') >$ $\mathit{sampleSize}(\hat{d}, \hat{k})$ we have $k' > \hat{k}$. Visually, (d', k') is somewhere above the diagonal for $(\hat{d}, \hat{k})$ but below the horizontal line for $\hat{k}$.

Denote $d'k'/\hat{k}$ by d'' such that $(d'', \hat{k})$ is on the same diagonal as (d', k') and directly to the left of $(\hat{d}, \hat{k})$. We have $d'' < \hat{d}$, consequently, $\mathit{sampleSize}(d'', \hat{k}) < \mathit{sampleSize}(\hat{d}, \hat{k})$. The diagonal segment from $(d'', \hat{k})$ to (d', k') has already been processed, each entry defining a sample size of no more than double the previous. Let (d^*, k^*) be the first pair that was processed in that diagonal segment with $\mathit{sampleSize}(d^*, k^*) \geq \mathit{sampleSize}(\hat{d}, \hat{k})$. This guarantees $\mathit{sampleSize}(d^*, k^*) < 2 \cdot \mathit{sampleSize}(\hat{d}, \hat{k})$. Since $k^* > \hat{k}$, we must have that $\mathit{numRounds}(d^*, k^*) > \mathit{numRounds}(\hat{d}, \hat{k})$, so Algorithm 2 should have terminated for the pair (d^*, k^*) with high probability, preserving the approximation ratio.

This analysis also leads us to see that the expected running time of Algorithm 2 is polynomial in n.

4.4 For plane graphs: FPT algorithm and PTAS via reduction to P-VC-3

Here we reduce the problem of computing a minimum obstacle representation for a plane graph (ORPeG) to maximum degree 3 planar vertex cover (P-VC-3).

Theorem 4.4.1. *Weighted ORPeG is no harder to approximate than planar vertex cover.*

Proof. We are given a plane graph G on n vertices in general position and an additive weight function w on its face set. We will construct a graph $\hat{G}$ that admits a vertex cover of cost k if and only if G admits an obstacle representation of cost $k + w(F_0)$, where F_0 is defined in the following paragraph.

Every bounded face of G that is not an empty triangle[4] must be selected as an obstacle; moreover, the unbounded face must be chosen if and only if its convex hull boundary contains a non-edge. Denote by F_0 the set of faces in G corresponding the these forced moves, which can be determined in polynomial time using standard techniques [10]. From now on, we only consider the non-edges of G that are not blocked by one of these faces. That is, non-edges that meet no faces other than empty triangles.

We now show that for every remaining non-edge s of G, there is a diamond in G such that s meets both triangles in that diamond. Assume for contradiction that s never crosses the diagonal edge of a diamond. Denote by u and v the endpoints of s, and orient the plane such that u is directly below v. Obtain a sequence of empty triangles $(f_0, f_1, \ldots, f_k)$ by tracing s from u (a vertex on f_0) to v (a vertex on f_k). Denote by v_i (for $1 \leq i \leq k$) the unique vertex in face f_k that is not a vertex of f_{i-1} (so that $v_k = v$). Without loss of generality, the reflex angle of f_0 and f_1 is to the right of s, which implies that v_1 is to the

[4] Recall that an *empty triangle* is a bounded face on three vertices not containing any other vertices, and that a *diamond* consists of two empty triangles that share an edge and have their vertices in convex position.

right of s. In order for f_2 to be the next face in this sequence, v_2 must be to the left of s. In general, in order for f_i to be the next face in this sequence, v_i must be on the other side of s from v_{i-1}. This pattern must continue indefinitely, lest two consecutive triangles form a diamond. The indefinite continuation of this pattern implies an infinite sequence of faces defined by s, and hence a contradiction.

We define $\hat{G}$, which is a subgraph of the dual of G, based on diamonds: each edge of $\hat{G}$ corresponds to diamond of G. The graph $\hat{G}$ is induced by these edges and the weights on its vertices are the weights on the corresponding empty triangles of G. At least one of the two triangles of every diamond must be chosen in any obstacle representation due to the diagonal non-edge. But by the previous paragraph, this is also sufficient to block all non-edges in G that only meet faces that are empty triangles. Therefore, we have a natural bijection that takes vertex covers of $\hat{G}$ with cost k to obstacle representations of G with cost $k + w(F_0)$. □

Remark 4.4.1. We may wish to adopt the more realistic bit model, since a plane graph drawing may have been expressed using a number of bits super-polynomial in $|V(G)|$ for vertex coordinates. In this model, the reduction could require time super-polynomial in $|V(G)|$ but nonetheless polynomial in the number of input bits.

From Theorem 4.4.1 we obtain the following results almost immediately.

Corollary 4.4.2. *Weighted ORPeG admits a polynomial-time approximation scheme (PTAS).*

Proof. Pre-process a given plane graph G on n vertices in general position (with an additive weight function w on its face set) in polynomial time to compute F_0 and $\hat{G}$ as in the proof of Theorem 4.4.1. Denote $|V(\hat{G})|$ by $\hat{n}$.

Denote by τ the size of a minimum obstacle representation for G and denote by τ' the size of a minimum vertex cover for $\hat{G}$. The well-known PTAS for Vertex Cover for

Planar Graphs [6] can be used to obtain for any given $\varepsilon > 0$ a vertex cover for $\hat{G}$ of size at most $(1+\varepsilon)\tau'$ in additional time $2^{O(1/\varepsilon)}\hat{n}^{O(1)} = 2^{O(1/\varepsilon)}n^{O(1)}$. Since $\tau = w(F_0) + \tau'$, we have $w(F_0) + (1+\varepsilon)\tau' \leq (1+\varepsilon)\tau$, which means that we have a PTAS for weighted ORPeG. □

A problem is called fixed parameter tractable (FPT) if for an input instance of size n together with an integer k, it is possible to determine whether or not there is a solution instance of size k in time $f(k)n^{O(1)}$ where f is an arbitrary function depending only on k. An algorithm that achieves such a running time is usually referred to as an FPT algorithm.

Corollary 4.4.3. *(Cardinality) ORPeG is fixed parameter tractable (FPT).*

Proof. Pre-process a given plane graph G with n vertices in general position in polynomial time (see Remark 4.4.1) to compute F_0 and $\hat{G}$ as in the proof of Theorem 4.4.1. Denote $\left|V(\hat{G})\right|$ by $\hat{n}$.

Using the FPT algorithm by Xiao [49] for Vertex Cover for Graphs with Maximum Degree 3 on $\hat{G}$, we can compute an obstacle representation for G with k obstacles in additional time at most $1.1616^{k-|F_0|}\hat{n}^{O(1)} = 1.1616^{k-|F_0|}n^{O(1)} = 1.1616^{k}n^{O(1)}$.

Alternatively, we can use the FPT algorithm by Alber et al. [1] for Vertex Cover for Planar Graphs on $\hat{G}$ to compute an obstacle representation for G with k obstacles in additional time at most $O(2^{4\sqrt{3(k-|F_0|)}}\hat{n}) = O(2^{4\sqrt{3(k-|F_0|)}}n) = O(2^{4\sqrt{3k}}n)$. □

4.5 Concluding remarks

Since computing the obstacle number for a graph drawing D exactly can be reduced to a problem of hitting non-edges with $O(n^4)$ faces, a naive deterministic algorithm can compute $obs(D)$ in time $2^{O(n^4)}$. In fact, such an algorithm uses time merely $n^{O(obs(D))}$, by trying every k-face combination for every value of k from 0 up to $obs(D)$.

What about the original problem of determining the obstacle number of a given abstract graph on n vertices? If all drawings of a graph could be enumerated up to the incidence matrix of faces versus non-edges, then by using our approximation algorithm in the "inner loop," we could obtain a $O(\log OPT)$-approximation to the original problem. While this may be viable for small instances (perhaps in conjunction with a distributed approach), we conjecture that this problem lies outside of NP and believe it to be intractable in a centralized model of computation. Our rationale follows.

For some simple order types of n-point configurations on the plane, a coordinate representation on an integer lattice needs exponentially many bits in n in order to allow the order type to be inferred [26]. Further, we know that a particular labeled graph has two drawings with different obstacle numbers but vertex sets of the same simple labeled order type. (The dual labeled order type of the $\binom{n}{2}$ connecting lines of n vertices in a drawing appears to be sufficient to determine the obstacle number for the drawing.) Hence, coordinate representations of some drawings for the present purpose seem to require at least exponential storage in n. Some drawing-based certificates will in turn have sizes super-polynomial in the number of bits that represent the abstract graph. It may be tempting to think that a certificate could instead be based on the poly(n) sized incidence matrix of faces versus non-edges, but it seems unlikely that one can decide in polynomial time whether or not the given graph has *some* drawing corresponding to a given incidence matrix.

Bibliography

[1] J. Alber, H. Fernau, and R. Niedermeier. Parameterized complexity: Exponential speed-up for planar graph problems. In *Electronic Colloquium on Computational Complexity (ECCC)*, pages 261–272. Springer, 2001.

[2] N. Alon. The number of polytopes, configurations and real matroids. *Mathematika*, 33(1):62–71, 1986.

[3] N. Alon, D. Moshkovitz, and S. Safra. Algorithmic construction of sets for k-restrictions. *ACM Trans. Algorithms*, 2:153–177, April 2006.

[4] H. Alpert, C. Koch, and J. Laison. Obstacle numbers of graphs. *Discrete & Computational Geometry*, 44:223–244, July 2010. www.springerlink.com/content/45038g67t22463g5.

[5] E. M. Arkin, D. Halperin, K. Kedem, J. S. B. Mitchell, and N. Naor. Arrangements of segments that share endpoints: single face results. *Discrete & Computational Geometry*, 13(3-4):257–270, 1995.

[6] B. S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. ACM*, 41:153–180, January 1994.

[7] B. Bollobás and A. Thomason. Hereditary and monotone properties of graphs. In R. L. Graham and J. Nešetřil, editors, *The Mathematics of Paul Erdős vol. 2, Algorithms and Combinatorics 14*, pages 70–78. Springer, Berlin - New York, 1997.

[8] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995.

[9] K. L. Clarkson. Algorithms for polytope covering and approximation. In *Proc. Algorithms and Data Structures, 3rd Workshop, WADS '93, Montréal, Canada, August 11–13, 1993*, LNCS 709, pages 246–252. Springer, 1993.

[10] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications (2nd ed.).* Springer-Verlag, Berlin, 2000.

[11] H. de Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41–51, 1990.

[12] A. Efrat and S. Har-Peled. Guarding galleries and terrains. *Info. Proc. Letters*, 100:238–245, December 2006.

[13] A. Efrat, F. Hoffmann, C. Knauer, K. Kriegel, G. Rote, and C. Wenk. Covering with ellipses. *Algorithmica*, 38:145–160, 2003.

[14] P. Erdős, P. Frankl, and V. Rödl. The asymptotic number of graphs not containing a fixed subgraph and a problem for hypergraphs having no exponent. *Graphs and Combinatorics*, 2:113–121, 1986.

[15] P. Erdős, D. J. Kleitman, and B. L. Rothschild. Asymptotic enumeration of K_n-free graphs. In *Colloq. Int. Teorie Comb., Roma, Tomo II*, pages 19–27, 1976.

[16] P. Erdős and G. Szekeres. A combinatorial problem in geometry. *Compositio Math.*, 2:463–470, 1935.

[17] G. Even, D. Rawitz, and S. Shahar. Hitting sets when the VC-dimension is small. *Info. Proc. Letters*, 95(2):358–362, 2005.

[18] S. Foldes and P. L. Hammer. Split graphs having Dilworth number 2. *Canadian Journal of Mathematics - Journal Canadien de Mathematiques*, 29(3):666–672, 1977.

[19] R. Fulek, N. Saeedi, and D. Sarıöz. Convex obstacle numbers of outerplanar graphs and bipartite permutation graphs. In J. Pach, editor, *Thirty Essays in Geometric Graph Theory*, Algorithms & Combinatorics Series. Springer, 2012. In press. Pre-print: arXiv:1104.4656v3 [cs.DM], 11p., Sep 2011.

[20] M. R. Garey and D. S. Johnson. The rectilinear Steiner tree problem is NP-complete. *SIAM Journal of Applied Mathematics*, 32:826–834, 1977.

[21] M. R. Garey, D. S. Johnson, and L. J. Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1(3):237–267, 1976.

[22] S. K. Ghosh. *Visibility Algorithms in the Plane*. Cambridge University Press, Cambridge, 2007.

[23] O. Giménez and M. Noy. The number of planar graphs and properties of random planar graphs. In *International Conference on Analysis of Algorithms, Conrado Martínez (ed.), Discrete Mathematics and Theoretical Computer Science Proceedings AD*, pages 147–156, 2005.

[24] J. E. Goodman and R. Pollack. Upper bounds for configurations and polytopes in $\mathbb{R}^d$. *Discrete & Computational Geometry*, 1(3):219–227, 1986.

[25] J. E. Goodman and R. Pollack. Allowable sequences and order types in discrete and computational geometry. In *New Trends in Discrete and Computational Geometry*, volume 10 of *Algorithms Combin.*, pages 103–134. Springer, Berlin, 1993.

[26] J. E. Goodman, R. Pollack, and B. Sturmfels. Coordinate representation of order types requires exponential storage. In *Proc. 21st Annual ACM Symposium on Theory of Computing*, STOC '89, pages 405–410, New York, NY, USA, 1989. ACM.

[27] D. Haussler and E. Welzl. ε-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987.

[28] M. P. Johnson and D. Sarıöz. Computing the obstacle number of a plane graph, August 2011. arXiv:1107.4624v2 [cs.CG].

[29] C. Lund and M. Yannakakis. On the hardness of approximating minimization problems. *Journal of the ACM*, 41:960–981, September 1994.

[30] J. Matoušek. *Lectures on Discrete Geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.

[31] J. Matoušek and P. Valtr. The complexity of lower envelope of segments with h endpoints. *Intuitive Geometry, Bolyai Society of Math. Studies*, 6:407–411, 1997.

[32] P. Mukkamala, J. Pach, and D. Sarıöz. Graphs with large obstacle numbers. In *Graph Theoretic Concepts in Computer Science*, LNCS 6410, pages 292–303. Springer, 2010.

[33] J. O'Rourke. Visibility. In *Handbook of Discrete and Computational Geometry*, CRC Press Ser. Discrete Math. Appl., pages 467–479. CRC, 1997.

[34] J. O'Rourke. Open problems in the combinatorics of visibility and illumination. In *Advances in Discrete and Computational Geometry (South Hadley, MA, 1996)*, volume 223 of *Contemp. Math.*, pages 237–243. Amer. Math. Soc., Providence, RI, 1999.

[35] J. Pach and P. K. Agarwal. *Combinatorial Geometry*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons Inc., 1995.

[36] J. Pach and D. Sarıöz. On the structure of graphs with low obstacle number. *Graphs and Combinatorics*, 27:465–473, 2011.

[37] R. Pollack, M. Sharir, and S. Sifrony. Separating two simple polygons by a sequence of translations. *Discrete & Computational Geometry*, 3:123–136, January 1988.

[38] H. J. Prömel and A. Steger. Excluding induced subgraphs: Quadrilaterals. *Random Structures and Algorithms*, 2(1):55–71, 1991.

[39] H. J. Prömel and A. Steger. Excluding induced subgraphs III: A general asymptotic. *Random Structures and Algorithms*, 3(1):19–31, 1992.

[40] H. J. Prömel and A. Steger. Excluding induced subgraphs II: extremal graphs. *Discrete Applied Mathematics*, 44:283–294, 1993.

[41] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, STOC '97, pages 475–484, New York, NY, USA, 1997. ACM.

[42] D. Sarıöz. Approximating the obstacle number for a graph drawing efficiently. In *Proc. 23rd Canadian Conference on Computational Geometry (CCCG '11)*, pages 297–302, August 2011.

[43] J. Spinrad, A. Brandstädt, and L. Stewart. Bipartite permutation graphs. *Discrete Applied Mathematics*, 18(3):279–292, 1987.

[44] R. I. Tyškevič and A. A. Černjak. Canonical decomposition of a graph determined by the degrees of its vertices (in Russian). *Vestsī Akad. Navuk BSSR Ser. Fīz.-Mat. Navuk*, 5(5):14–26, 138, 1979.

[45] J. Urrutia. Art gallery and illumination problems. In *Handbook of Computational Geometry*, pages 973–1027. North-Holland, Amsterdam, 2000.

[46] M. J. van Kreveld. Bold graph drawings. In *Proc. 21st Canadian Conference on Computational Geometry (CCCG '09)*, pages 119–122, 2010.

[47] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies to their probabilities. *Theory Probab. Appl.*, 16(2):264–280, 1971.

[48] V. V. Vazirani. *Approximation Algorithms*. Springer, 2004.

[49] M. Xiao. A note on vertex cover in graphs with maximum degree 3. In M. Thai and S. Sahni, editors, *Computing and Combinatorics*, volume 6196 of *Lecture Notes in Computer Science*, pages 150–159. Springer Berlin / Heidelberg, 2010.