Irene Parada

On straight-line and topological drawings of graphs in the plane

DOCTORAL THESIS

to achieve the university degree of

Doktorin der Naturwissenschaften

submitted to

Graz University of Technology

Supervisor

Assoc.-Prof. Dipl-Ing. Dr.techn. Oswin Aichholzer

Institute of Software Technology

Graz, October 2019

Affidavit

I declare that I have authored this thesis independently, that I have not used other than the declared sources/resources, and that I have explicitly indicated all material which has been quoted either literally or by content from the sources used. The text document uploaded to TUGRAZonline is identical to the present doctoral thesis.

Date

Signature

Abstract

This thesis deals with drawings of graphs in the plane, both from a structural and an algorithmic point of view. We focus on non-planar graphs and different common restrictions on the drawings. The first part of the thesis considers straight-line drawings. For these drawings, we study a variant of the crossing number problem in which the edges of the drawing are 2-colored and only crossings between edges of the same color are counted. We denote this variant the 2-colored crossing number. For the complete graph K_n we obtain asymptotic upper and lower bounds on the 2-colored crossing number by showing that they can be derived from optimal and near-optimal instances with few vertices. Moreover, for any fixed straight-line drawing of K_n, we improve the upper bound on the ratio between the same-color crossings in the best 2-coloring of the edges and the total number of crossings in the drawing.

In a straight-line drawing of a graph, the placement of the vertices in the plane determines the drawing. Thus, we extend our study to point sets in the plane and combinatorial properties of them. Order types are equivalence classes of point sets relevant for many combinatorial problems. We investigate straight-line drawings with few edges that unequivocally display the order type of the point set they are drawn on. For that, we introduce the concept of exit edges, which prevent the order type from changing under continuous motion of vertices. Exit edges have a natural dual characterization, which allows us to efficiently compute them and to bound their number.

Problems for which order types are relevant include Erdős-Szekeres-type questions. In this thesis we study the variant on the number $h_5(n)$ of empty convex pentagons in any set of n points in general position in the plane. Despite many efforts made in the last 30 years, the best upper bound was quadratic while the best lower bound was linear. We show that $h_5(n) = \Omega(n \log^{4/5} n)$, thus proving that there is always a superlinear number of empty convex pentagons.

In the second part of this thesis we focus on (not necessarily straight-line) drawings of graphs in the plane with restrictions on how the edges can cross. We first consider simple drawings, in which two edges can share at most one point, either a common endpoint or a crossing point. Given a simple drawing of a graph, we prove that, surprisingly, deciding if an edge of the complement graph can be inserted into the drawing such that the result is still a simple drawing is NP-hard. We also study the computational complexity of inserting edges of the complement graph into a particular class of simple drawings, namely 1-plane drawings. For simple drawings of the complete bipartite graph $K_{m,n}$, we prove the existence of plane spanning trees as subdrawings, shedding light on one of the basic aspects of these drawings.

Finally, we consider semi-simple drawings, in which two incident edges cannot cross, but two independent edges can cross multiple times. We study the question on deciding whether a given rotation system can be realized in a semi-simple drawing and show that, in contrast to the case of simple drawings, it is not enough to check the sub-rotation systems of five vertices. This contrast between simple and semi-simple drawings is also demonstrated for other questions.

Kurzfassung

In dieser Dissertation beschäftigen wir uns mit Zeichnungen von Graphen in der Ebene. Dabei werden sowohl strukturelle, als auch algorithmische Gesichtspunkte betrachtet. Der Schwerpunkt der Arbeit liegt dabei auf nicht-planaren Graphen und verschiedenen, in der Literatur üblichen, Klassen von Zeichnungen.

Im ersten Teil dieser Arbeit werden geradlinige Zeichnungen betrachtet. Wir untersuchen eine Variante des Kreuzungszahlen-Problems geradliniger Zeichnungen, bei der die Kanten zweigefärbt sind und nur Kreuzungen zwischen Kanten gleicher Farbe für die Kreuzungszahl berücksichtigt werden. Wir bezeichnen diese Variante als die 2-gefärbte Kreuzungszahl. Für den vollständigen Graphen K_n erhalten wir asymptotische obere und untere Schranken für die 2-gefärbte Kreuzungszahl, indem wir zeigen, dass diese von optimalen und nahezu optimalen Instanzen mit wenigen Knoten abgeleitet werden können. Weiters betrachten wir das Verhältnis zwischen der 2-gefärbten Kreuzungszahl und der Gesamtanzahl an Kreuzungen für beliebige geradlinige Zeichnungen und zeigen eine verbesserte obere Schranke für dieses Verhältnis.

Im Falle geradliniger Zeichnungen eines Graphen bestimmt die Platzierung der Knoten in der Ebene die gesamte Zeichnung. Daher untersuchen wir Punktmengen in der Ebene und ihre kombinatorischen Eigenschaften. Ordnungstypen sind Äquivalenzklassen von Punktmengen, die für viele kombinatorische Probleme relevant sind. Wir untersuchen geradlinige Zeichnungen mit wenigen Kanten, die den Ordnungstypus ihrer Knotenmenge eindeutig darstellen. Dafür führen wir das Konzept von sogenannten Exitkanten ein. Diese verhindern eine Veränderung des Ordnungstypus unter stetiger Bewegung. Exitkanten haben eine natürliche duale Charakterisierung, die es uns erlaubt, die Exitkanten effizient zu berechnen und ihre Anzahl zu beschränken.

Zu den Problemen, für die Ordnungstypen relevant sind, gehört auch die Klasse der Erdős-Szekeres-Probleme über die Existenz und Anzahl von Vielecken in Punktmengen. In dieser Arbeit behandeln wir die Anzahl $h_5(n)$ von leeren, konvexen Fünfecken in einer beliebigen Menge von n Punkten in allgemeiner Lage in der Ebene. Trotz vieler Bemühungen in den letzten 30 Jahren war die beste bekannte obere Schranke quadratisch und die beste untere Schranke linear. Wir zeigen, dass $h_5(n) = \Omega(n \log^{4/5} n)$ und beweisen damit, dass es immer eine superlineare Anzahl an leeren, konvexen Fünfecken gibt.

Im zweiten Teil dieser Arbeit konzentrieren wir uns auf allgemeine (nicht zwingend geradlinige) Zeichnungen von Graphen in der Ebene, in denen es Beschränkungen gibt, inwiefern ihre Kanten sich kreuzen dürfen. Wir betrachten zunächst simple Zeichnungen, in denen zwei Kanten maximal einen Punkt gemeinsam haben. Dies kann ein gemeinsamer Endpunkt oder ein Kreuzungspunkt sein. Wir zeigen, dass es in einer gegebenen simplen Zeichnung eines Graphen NP-schwer ist zu entscheiden, ob der Zeichnung eine Kante des vollständigen Graphen so hinzugefügt werden kann, dass das Resultat immer noch eine simple Zeichnung ist. Für spezielle Klassen von simplen Zeichnungen, nämlich 1-planare Zeichnungen, untersuchen wir ebenso die algorithmische Komplexität des Hinzufügens von Kanten des Komplementgraphen. Für simple Zeichnungen des vollständigen, bipartiten Graphen zeigen wir die Existenz kreuzungsfreier Spannbäume als Teilzeichnungen und gewinnen so mehr Einsicht in die grundlegende Struktur dieser Zeichnungen.

Schließlich behandeln wir sogenannte semisimple Zeichnungen. Das sind Zeichnungen, in denen sich zwei inzidente Kanten nicht kreuzen dürfen, während zwei unabhängige Kanten sich auch mehrfach kreuzen können. Dabei gehen wir der Frage nach, ob ein gegebenes Rotationssystem als semisimple Zeichnung realisiert werden kann und zeigen, dass es, anders als bei simplen Zeichnungen, nicht hinreichend ist, die Teilrotationssysteme von fünf Knoten zu überprüfen. Dieser Unterschied zwischen simplen und semisimplen Zeichnungen wird auch für weitere Fragen demonstriert.

Acknowledgements

I want to start by thanking my advisor Oswin Aichholzer for his continuous support during my PhD studies. I am especially grateful for him being always approachable and for the many discussions and conversations from which I learned a lot. Research with Oswin during these years was always not just productive, but also fun.

This gratitude also extends to the other past and present members of the computational geometry group in Graz. Very especially, I want to thank Birgit Vogtenhuber for being a wonderful tutor during my PhD. Her support, insightful advice, and the abundant discussions with her on both scientific and non-scientific topics made a big difference in my PhD and beyond. I would also like to thank all the members of the Doctoral Program "Discrete Mathematics", who made my PhD and my time in Graz a much better experience.

If Oswin Aichholzer is my academic father, Vera Sacristán is my academic mother. She introduced me to the field of computational geometry and to its amazing community of researchers. I am very grateful for that and for her honest advice and support during all these years. It is always a pleasure to work with Vera and to learn from and with her.

During my studies I had the opportunity to work with many researchers of the discrete and computational geometry community. Collaborating with them was an exciting and inspiring experience that also broadened my understanding of the field. I particularly want to thank my numerous co-authors for that. Not only did I meet wonderful researchers, but I also made good friends. Moreover, I had the privilege to make several research visits. I want to thank my hosts Ruy Fabila (with Carlos Hidalgo), Gelasio Salazar (with Carolina Medina), Jorge Urrutia, and Pavel Valtr (with Martin Balko and Jan Kynčl) for the great moments of research and their hospitality.

Furthermore, since my undergraduate studies, I have been enrolled as a student in five different universities (UCM, UNED, KTH, UPC, and TU Graz) in three different countries (Spain, Sweden, and Austria). All these institutions are publicly funded, and I want to acknowledge the opportunity that I was given by them. I also want to mention the public health system in both Spain and Austria, and I want to thank the professionals working in them for their assistance when it was most needed.

Last, but not least, I want to thank my family and friends. I very especially want to thank my parents, Gonzalo and Vega, for their constant support. ¡Muchas gracias a los dos! (Thanks a lot to the two of you!). Finally, I want to deeply thank Fabian, for everything.

To women in science, who have paved the way for my generation of female scientists.

Contents

1. Introduction

Questions on how a graph can be drawn constitute a central part of graph theory and discrete and computational geometry. In particular, the study of how non-planar graphs can be drawn in the plane while respecting certain properties is a highly active area of research. Probably the oldest and most relevant open problem in that direction is the crossing number problem, which asks for the minimum number of crossings that a drawing of a given graph can have.

During World War II, the mathematician Paul Turán was forced to work in a brick factory pushing wagons from the kilns to the storage yards. In that situation, the crossings in the rails were the main source of trouble. This inspired Turán to state the question of what is the minimum number of crossings that the rails connecting every kiln with every storage yard can possibly have [190]. The mathematical reformulation of this question asks for the crossing number of the complete bipartite graph $K_{m,n}$. In the 1950s, both Urbanik and Zarankiewicz published a solution with a tight bound of $Z(m,n) = \lfloor \frac{n}{2} \rfloor \lfloor \frac{n-1}{2} \rfloor \lfloor \frac{m}{2} \rfloor \lfloor \frac{m-1}{2} \rfloor$ crossings, attained by a family of drawings like the one depicted in Figure 1.1(a). However, in 1965 and 1966 Kainen and Ringel, independently, found a flaw in the arguments. More on the early history of the crossing number problem can be found in [111]. The conjecture that the crossing number of $K_{m,n}$ is $Z(m,n)$ is still open and nowadays known as Zarankiewicz's conjecture. In 1971 Kleitman [129] proved the conjecture for drawings of $K_{5,n}$. Since then, the main progress on the lower bound for particular cases has relied on computer aid. In 1993 Woodall [200] showed that the conjecture is true for drawings of $K_{m,n}$ with $m \in \{7,8\}$ and $n \in \{7,8,9,10\}$. The conjecture was tackled recently using flag algebras [43, 153] and the claimed result[1] is that, for a large enough value of n, $K_{n,n} \geq 0.973 \cdot Z(n,n)$. This is an improvement over several previous results [74, 73].

A similar open conjecture exists for the crossing number of the complete graph K_n. Harary–Hill conjecture states that the crossing number of K_n is $Z(n) = \frac{1}{4} \lfloor \frac{n}{2} \rfloor \lfloor \frac{n-1}{2} \rfloor \lfloor \frac{n-2}{2} \rfloor \lfloor \frac{n-3}{2} \rfloor$. Actually, as for Zarankiewicz's conjecture, there are drawings of K_n conjectured to be crossing-optimal [54, 113]. Among them, there are cylindrical drawings that can be drawn in the plane with the vertices split between two concentric circles. Moreover, the edges between vertices on the inner (outer) circle are drawn inside (outside) it, and the edges between vertices on different circles are drawn in the region between the circles; see Figure 1.1(b). Harary–Hill conjecture has been proven for $n \leq 12$ [110, 164]. In the last years, the conjecture has also been proven for some relevant classes of drawings of increasing generality, and thus containing previous classes, in a series of papers [2, 3, 4, 5, 42, 149].

In crossing-optimal drawings two adjacent edges do not cross, and independent edges can cross at most once. The drawings with these properties are known as simple drawings, and also

[1] Details have not been yet published.

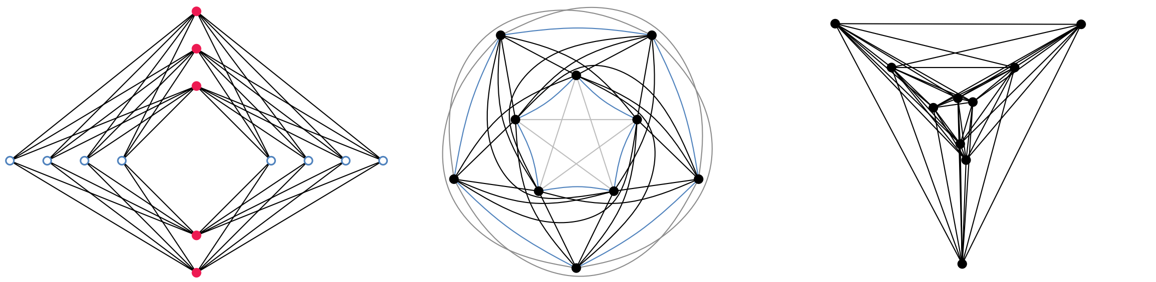

(a) Zarankiewicz's drawing of $K_{5,8}$ with 48 crossings. (b) Cylindrical drawing of K_{10} with 60 crossings. (c) Optimal straight-line drawing of K_{10} with 62 crossings.

Figure 1.1.: Examples of crossing-optimal drawings. The families of drawings represented in (a) and (b) are conjectured to be crossing-optimal also in examples with more vertices.

denoted good drawings or simple topological graphs in the literature. Due to their connection with Harary–Hill conjecture, simple drawings of the complete graph have received considerable attention and several aspects of them have been studied [24, 38, 104, 164, 178]. In addition, there is an efficient way of encoding all different weak isomorphism classes of simple drawings of K_n, where two simple drawings are weakly isomorphic if the same pairs of edges cross. The rotation of a vertex in a drawing is the (clockwise) cyclic order of all the edges incident to it. The rotations of all the vertices of a drawing form its rotation system. Two simple drawings of K_n have the same rotation system if and only if the same pairs of edges cross [104, 162]. This allows for all classes of simple drawings of K_n with $n \leq 9$ to be enumerated [1]. Moreover, it can be decided in polynomial time if a given set of rotations corresponds to the rotation system of a simple drawing of K_n [137, 138]. In contrast, this realizability question is NP-hard for general graphs. For general graphs, the crossing number problem is also NP-hard [102], even when the rotation system is prescribed [169].

In the second part of this thesis we explore central aspects of simple drawings of general graphs from an algorithmic perspective and of the complete bipartite graph from a structural perspective. On the algorithmic side, we focus on the problem of extending a partial simple drawing of a graph to a simple drawing of the full graph. In the last decade, questions on the extendibility of partial drawings have been considered for several classes of graphs [33, 40, 57, 72, 123, 146, 167]. These questions fit in the broader context of extendibility of partial representations of graphs, where these representations might not be drawings of a graph [64, 63, 65, 125, 128, 126, 127]. We also extend the algorithmic study on the extendibility of simple drawings to 1-plane drawings. These drawings, considered for the first time by Ringel in 1965 [174], are simple drawings in which each edge participates in at most one crossing. We also study semi-simple drawings of K_n, in which incident edges do not cross, but independent edges are allowed to cross an arbitrary number of times.

A natural restricted version of (simple) drawings are straight-line drawings, also known as geometric graphs. In this kind of drawings the edges are drawn as straight-line segments connecting their two endpoints. Thus, the placement of the vertices in the plane entirely determines the full drawing. The crossing number problem has also been studied in this setting, and it is known to be $\exists\mathbb{R}$-complete for general graphs [119]. Of particular interest is the so-called rectilinear crossing number of K_n, that asks for the minimum number of crossings in a straight-line drawing of K_n. This value is in general different to the crossing number without the restriction of the drawing being a straight-line one; see Figures 1.1(b) and 1.1(c) for an illustration. Actually, the rectilinear crossing number of K_n and the crossing number of K_n, though both are asymptotically $\Theta(n^4)$, differ in the asymptotically relevant term [144]. The question asking for the rectilinear crossing number of K_n is also an open problem, but, in contrast to the case of general drawings, there is no formula conjectured to give the optimal value for every n. Both the upper and the lower bound have been object of abundant work, including very recent one [18, 88]; see the survey [8] for a compilation of previous progress. In Chapter 2 we consider a variant of the rectilinear crossing number of K_n in which edges are 2-colored and only monochromatic crossings are counted. For this variant we give both lower and upper asymptotic bounds. These bounds require the exploration of instances with few vertices. In this case, as in a wide variety of combinatorial problems on point sets in the plane, there is no need to deal with all precise placements of the points; instead, it is possible to just consider equivalent classes that have a combinatorial description.

Especially relevant examples of equivalence classes of point sets are order types [105]. Two sets of n points in general position in the plane have the same order type if their points can be labeled with integers from 1 to n such that the corresponding triples of points in both point sets have the same orientation (clockwise or counterclockwise). Among other things, the order type determines which sets of four points are in convex position. Thus, if we connect every pair of points with a line segment, that we can see as the straight-line drawing of the complete graph on the point set, the order type determines which edges cross. For this reason, order types appear ubiquitously in the study of extremal problems on point sets and straight-line drawings. There exists an order type database [12, 26], containing all the order types realizable as sets of up to eleven points in the plane. The order types realizable as sets of up to ten points are available online [11] and the (over two thousand million) ones realizable as sets of eleven points are available upon request from Aichholzer (they require almost 100GB of space). In the next three chapters of this thesis this database is used to obtain several computer-assisted results. Apart from order types, there are also other definitions of equivalence classes of point sets that can be described combinatorially. For example, circular sequences define a finer classification of point sets in the plane while radial systems provide a coarser one [27, 89].

Combinatorial problems on point sets that do not depend on the precise coordinates of the points but on their order type, include Erdős-Szekeres-type questions. In 1933 Esther Klein asked whether every large enough point set in general position contains a convex k-gon. The affirmative answer published in 1935 is known as the Erdős-Szekeres theorem[2] [85], a classic result in discrete geometry and Ramsey theory. Let $ES(k)$ be the smallest integer such that every set of at least $ES(k)$ points in general position contains a convex k-gon. The original paper [85] contains two proofs. The first one is based on Ramsey's theorem and the second one gives the (better) upper bound $ES(k) \leq \binom{2k-4}{k-2} + 1$. More than 25 years later, Erdős and Szekeres published a construction giving the lower bound $2^{k-2} + 1 \leq ES(k)$ and conjectured it

[2]There is another Erdős-Szekeres theorem on the existence of monotonically increasing or decreasing subsequences of certain length in any large enough sequence.

to be tight. Recently, after many years of efforts and progress that did however not improve the order of magnitude [66, 130, 150, 154, 189, 196], Suk [187] showed that $ES(k) = 2^{k+o(k)}$.

In the 1970s, Erdős [83] asked whether in every large enough point set in general position in the plane there is always an empty convex k-gon. Harborth [115] proved that there is an empty convex pentagon in every set of 10 points in general position and showed that 9 points are not enough. The answer was conjectured to be affirmative for every k. However, Horton [122] gave a construction, for every value of n, of a set of n points in general position with no empty convex heptagon. His construction was later extended to so-called Horton sets and squared Horton sets [193] and to higher dimensions [195]. More than 20 years later, Gerken [103] and Nicolás [151] showed independently that there is an empty convex hexagon in every sufficiently large point set in general position. Erdős [84] also asked for the minimum number of empty convex k-gons in a set of n points in general position. Horton's construction implies that this number is 0 for $k \geq 7$. Abundant research has been made for $k \in \{3, 4, 5, 6\}$; see [10] for a survey. In Chapter 4 we study this question for $k = 5$ and give the first superlinear lower bound.

1.1. Outline of the thesis

The first part of the thesis studies both point sets and straight-line drawings of graphs drawn on them. In Chapter 2 we consider a variant of the rectilinear crossing number in which we only count the monochromatic crossings of the best 2-edge-coloring. We give both upper and lower asymptotic bounds for the complete graph K_n. In Chapter 3, we study compact visualizations of order types in the plane. We introduce the concept of exit edges, which prevent the order type from changing under continuous motion of vertices. Exit edges have a natural dual characterization, which allows us to efficiently compute them and to bound their number. In Chapter 4 we show that, asymptotically, the number of empty convex pentagons in every set of n points in general position in the plane is at least superlinear.

The second part of the thesis deals with topological drawings of graphs. Given a simple drawing $D(G)$ of a graph G, in Chapter 5 we prove that it is NP-complete to decide whether an edge connecting two non-adjacent vertices can be inserted into $D(G)$ such that the result is a simple drawing. Moreover, we present a polynomial algorithm for the particular case in which the two vertices that we aim to connect are a dominating set for G. We also consider inserting edges into a 1-plane drawing such that the result is still a 1-plane drawing. In Chapter 6 we show that the problem of deciding whether k edges, each connecting two non-adjacent vertices, can be inserted into a 1-plane drawing is fixed-parameter tractable with respect to k. In Chapter 7 we study simple drawings of the complete bipartite graph and show that all of them contain plane spanning trees of a particular kind. Finally, in Chapter 8 we consider semi-simple drawings the complete graph K_n, especially in relation to rotation systems and to long-standing conjectures for simple drawings of K_n.

Several parts of the contents of this thesis correspond to work that has been presented or is accepted for publication [13, 14, 15, 16, 17, 19, 20, 21, 29, 35, 36, 131]. Moreover, the content of Sections 5.5 and 7.5 is currently being prepared for publication [23, 37]. We remark that the work in all these papers was the product of a collaborative effort between the coauthors.

1.2. Note on notation

We use standard graph-theoretic terminology. Following [78], a *graph* G is an ordered pair of (disjoint) sets (V, E) such that $E \subseteq [V]^2$, that is, the elements of E are 2-element subsets of V. The elements of V are the *vertices* of the graph G and the elements of E are its edges. In this thesis we consider only finite graphs where V and E are finite sets. A graph with vertex set V is said to be a graph *on* V. A vertex $v \in V$ is *incident* to an edge $e \in E$ if $v \in e$, and in that case, e is also *incident* to v. The *degree* of a vertex is the number of edges incident to it. An edge $\{u, v\}$ is denoted by uv. If $uv \in E$, then u and v are said to be *adjacent* and two (distinct) edges are *adjacent* if there is a vertex such that both edges are incident to it. A set of vertices or of edges is *independent* if no two of its elements are adjacent. A graph $G' = (V', E')$ is a *subgraph* of $G = (V, E)$ if $V' \subseteq V$ and $E' \subseteq E$. If $G' = (V', E')$ is a *subgraph* of $G = (V, E)$ and E' contains all the edges $uv \in E$ with $u, v \in V'$, then G' is an *induced subgraph* of G. In this case we say that G' is *induced* by V' and denote it by $G[V']$. A subgraph $G' = (V', E')$ of $G = (V, E)$ is *spanning* if $V' = V$. The *complement* $\overline{G}$ of $G = (V, E)$ is the graph on V with edge set $[V]^2 \setminus E$.

In a drawing of a graph, vertices are represented by distinct points in the plane and edges are represented by Jordan arcs with their incident vertices as endpoints. We identify the vertices and edges of the underlying abstract graph with the corresponding ones in the drawing. Thus, in a slight abuse of notation, vertices and edges also refer to the points and arcs that represent them. We consider drawings in which edges do not contain any other vertices, no three edges intersect in the same point, and any point shared by two edges is either a proper crossing or a common endpoint (tangencies are not allowed). A drawing is *plane* if there are no crossings.

When considering point sets in the plane in this thesis, being in general position means that no three points are collinear. Given two points p and q in the plane, pq denotes the line segment spanned by p and q. If p and q are points representing vertices of a graph, pq also denotes the edge incident to the corresponding vertices.

Throughout this thesis we use O-notation to describe asymptotic behavior. We also make use without definition of basic concepts in computational complexity, such as P, NP, and NP-hardness. A basic introduction to these notions can be found in [68]. More specific concepts are described in the corresponding chapters. The author of this thesis is not aware of a particular textbook on computational complexity presenting the basic notions together with (in)approximability, parameterized complexity and concepts more used in geometry like $\exists\mathbb{R}$-hardness, but an entertaining introduction to most of these topics oriented to obtaining lower bounds is given in Demaine's course [76] (available online). The interested reader can find an introduction to computational complexity in various textbooks, including [100, 165, 184, 197]. For more advanced topics and current trends we refer to [34]. Finally, there are several textbooks on parameterized complexity [69, 79, 80, 92, 94, 152].

Specific concepts and notation are introduced in the corresponding chapters.

Part I.

Point sets and straight-line drawings of graphs

2. Bounding the rectilinear 2-colored crossing number

Part of the results presented in this chapter have been accepted for publication [20]. It is planned that these results also appear in the thesis of the coauthor Carlos Hidalgo Toscano.

2.1. Introduction

For a drawing of a non-planar graph G in the plane it is of interest from both a theoretical and practical point of view, to minimize the number of crossings. The minimum such number is known as the *crossing number* $\mathrm{cr}(G)$ of G. There are many variants on crossing numbers, see the comprehensive dynamic survey of Schaefer [179]. In this chapter we focus on a version combining two of them: the *k-planar crossing number* and the *rectilinear crossing number*.

The *k-planar crossing number* $\mathrm{cr}_k(G)$ of a graph G is the minimum of $\mathrm{cr}(G_1) + \cdots + \mathrm{cr}(G_k)$ over all sets of k graphs $\{G_1, \ldots, G_k\}$ whose union is G. For $k = 2$, it was introduced by Owens [156] who called it the *biplanar crossing number*; see [70, 71] for a survey on biplanar crossing numbers. Shahrokhi et al. [183] introduced the generalization to $k \geq 2$.

The *rectilinear crossing number* of G, $\overline{\mathrm{cr}}(G)$, is the minimum number of pairs of edges that cross in any straight-line drawing of G. Of special relevance is $\overline{\mathrm{cr}}(K_n)$, the rectilinear crossing number of the complete graph on n vertices. The current best published bounds on $\overline{\mathrm{cr}}(K_n)$ are $0.379972\binom{n}{4} < \overline{\mathrm{cr}}(K_n) < 0.380473\binom{n}{4} + \Theta(n^3)$ [7, 88]. The upper bound was achieved using a duplication process and has been improved in [18] to $\overline{\mathrm{cr}}(K_n) < 0.38044921\binom{n}{4} + \Theta(n^3)$.

A *k-edge-coloring* of a drawing D of a graph is an assignment of one of k possible colors to every edge of D. The *rectilinear k-colored crossing number* of a graph G, $\overline{\mathrm{cr}}_k(G)$, is the minimum number of monochromatic crossings (pairs of edges of the same color that cross) in any k-edge-colored straight-line drawing of G. This parameter was introduced before and called the *geometric k-planar crossing number* [159]. In the same paper, as well as in [183], also the *rectilinear k-planar crossing number* was considered, which asks for the minimum of $\overline{\mathrm{cr}}(G_1) + \ldots + \overline{\mathrm{cr}}(G_k)$ over all sets of k graphs $\{G_1, \ldots, G_k\}$ whose union is G. We prefer our terminology because the terms geometric and rectilinear are very often used interchangeably in the context of crossing numbers and because the term k-planar is extensively used in graph drawing with a different meaning; see for example [77, 120, 132]. We remark that in graph drawing, *rectilinear* sometimes also refers to orthogonal grid drawings (which is not the case here).

In this chapter we focus on the case where G is the complete graph K_n, and we prove the following lower and upper bounds on $\overline{\text{cr}}_2(K_n)$:

$$0.03\binom{n}{4} + \Theta(n^3) < \overline{\text{cr}}_2(K_n) < 0.11798016\binom{n}{4} + \Theta(n^3).$$

Our approach is based on theoretical results that guarantee asymptotic bounds from the information of small point sets. Thus, it implies computationally dealing with small sets, both to guarantee a minimum amount of monochromatic crossings (for the lower bound) and to find examples with few monochromatic crossings and some other desired properties (for the upper bound).

From an algorithmic point of view, our challenge is twofold. On the one hand, we need to optimize the point configuration (order type) to obtain a small number of crossings. On the other hand, we need to determine a coloring of the edges of K_n that minimizes the colored crossing number for a fixed point set.

For the first problem there is not even a conjecture of point configurations that minimize the rectilinear crossing number of K_n for any n. The latter problem corresponds to finding a maximum cut in a segment intersection graph, which in general is NP-complete [28]. Moreover, these two problems are not independent. There exist examples where a point set with a non-minimal number of uncolored crossings allows for a coloring of the edges so that the resulting colored crossing number is smaller than the best colored crossing number obtained from a set minimizing the uncolored crossing number. Thus, the two optimization processes need to interleave if we want to guarantee optimality. But, as we will see in Section 2.2, even this combined optimization does not guarantee to yield the best asymptotic result. There are sets of fixed cardinality and with larger 2-colored crossing number which—due to an involved duplication process—give a better asymptotic constant than the best minimizing sets. This is in contrast to the uncolored setting [6, 7], where for any fixed cardinality, sets with a smaller crossing number always give better asymptotic constants. Also, it clearly indicates that our extended duplication process for 2-colored crossings differs essentially from the original version.

As mentioned, drawings with few crossings do not necessarily admit a coloring with few monochromatic crossings. This observation motivates the following question: given a fixed straight-line drawing D of K_n, what is the ratio between the number of monochromatic crossings for the best 2-edge-coloring of D and the number of (uncolored) crossings in D? A simple probabilistic argument shows that this ratio is less than $1/2$. In Section 2.4, we improve that bound, showing that for sufficiently large n, it is less than $1/2 - c$ for some positive constant c.

In a slight abuse of notation, we denote by $\overline{\text{cr}}(D)$ the number of pairs of edges in D that cross and call it the rectilinear crossing number of D. The (rectilinear) 2-colored crossing number of a straight-line drawing D, $\overline{\text{cr}}_2(D)$, is then the minimum of $\overline{\text{cr}}(D_1) + \overline{\text{cr}}(D_2)$, over all pairs of straight-line drawings $\{D_1, D_2\}$ whose union is D. For a given 2-edge-coloring χ of D, we denote by $\overline{\text{cr}}_2(D, \chi)$ the number of monochromatic crossings in D. Thus, $\overline{\text{cr}}_2(D)$ is the minimum of $\overline{\text{cr}}_2(D, \chi)$ over all 2-edge-colorings χ of D.

Outline. In Section 2.2 we prove that, given a 2-colored straight-line drawing D of K_n, there is a duplication process that allows us to obtain a 2-colored straight-line drawing D_k of $K_{2^k n}$ for any $k \geq 1$ whose 2-colored crossing number $\overline{\mathrm{cr}}_2(D_k)$ can be easily calculated. Moreover, we can obtain the asymptotic value when $k \to \infty$. By finding good sets of constant size as a seed for the duplication process, we obtain an asymptotic upper bound for $\overline{\mathrm{cr}}_2(K_n)$. In Section 2.3 we obtain a lower bound for $\overline{\mathrm{cr}}_2(K_n)$ using the crossing lemma, and we improve it with an approach again using small drawings. For sufficiently large n, we show in Section 2.4 that for any straight-line drawing D of K_n, $\overline{\mathrm{cr}}_2(D)/\overline{\mathrm{cr}}(D) < 1/2 - c$ for a positive constant c, that is, using two colors saves more than half of the crossings. Finally, in Section 2.5 we present a summary of the chapter.

2.2. Upper bounds on $\overline{\mathrm{cr}}_2(K_n)$

For the rectilinear crossing number $\overline{\mathrm{cr}}(K_n)$, the best upper bound [18] comes from finding examples of straight-line drawings of K_n (for a small value of n) with few crossings which are then used as a seed for the duplication process in [6, 7]. To be able to apply this duplication process, the starting set P with m points has to contain a halving matching. If m is even (odd), a *halving line* of P is a line that passes exactly through two (one) points of P and leaves the same number of points of P to each side. If it is possible to match each point p of P with a halving line of P though this point in such a way that no two points are matched with the same line, P is said to have a *halving matching*. It is then shown in [6] that every point of P can be substituted by a pair of points in its close neighborhood such that the resulting set Q with $2m$ points contains again a halving line matching. Iterating this process yields the mentioned upper bound for $\overline{\mathrm{cr}}(K_n)$, where this bound depends only on m and the number of crossings of the starting set P.

In this section, we prove that a significantly more involved but similar approach can be adopted for the 2-colored case. Unlike the original approach, we cannot always get a matching which simultaneously halves both color classes. Moreover, even for sets where such a halving matching exists, it cannot be guaranteed that this property is maintained after the duplication step. We will see below that we need a more involved approach, where the matchings are related to the distribution of the colored edges around a vertex. Consequently, the number of crossings which are obtained in the duplication, and thus, the asymptotic bound we get, not only depends on the rectilinear 2-colored crossing number of the starting set, but also on the specific distribution of the colors of the edges. In that sense, both the heuristics for small drawings and the duplication process for the rectilinear 2-colored crossing number differ significantly from the uncolored case.

Throughout this section, P is a set of m points in general position in the plane, where m is even. Let p be a point in P. By slight abuse of notation, in the following we do not distinguish between a point set and the straight-line drawing of K_n it induces. Given a 2-coloring χ of the straight-line drawing of the complete graph that P induces, we denote by $L(p)$ and $S(p)$ the edges incident to $p \in P$ of the larger and smaller color class at p, respectively. An edge pq incident to $p, q \in P$ is called a *χ-halving edge of p* if the number of edges of $L(p)$ to the right of the line $\overline{qp}$ spanned by p and q (and directed from q to p) and the number of edges of $L(p)$ to the left of $\overline{qp}$ differ by at most one. A matching between the points of P and their χ-halving edges is called a *χ-halving matching for P*.

Theorem 2.1. *Let P be a set of m points in general position and let χ be a 2-coloring of the edges induced by P. If P has a χ-halving matching, then the 2-colored rectilinear crossing number of K_n can be bounded by*

$$\overline{\mathrm{cr}}_2(K_n) \leq \frac{24A}{m^4}\binom{n}{4} + \Theta(n^3)$$

where A is a rational number that depends on P, χ, and the χ-halving matching for P.

Proof. First we describe a process to obtain from P a set Q of $2m$ points, a 2-edge-coloring χ' of the edges of the straight-line drawing of K_{2m} that Q induces and a χ'-halving matching for Q. The set Q is constructed as follows. Let p be a point in P and pq its χ-halving edge in the matching. We add to Q two points p_1, p_2 placed along the line $\overline{qp}$ and in a small neighborhood of p such that:

(i) if f is an edge different from pq that is incident to p, then p_1 and p_2 lie on different sides of the line spanned by f;

(ii) if f is an edge different from pq that is not incident to p, then p_1 and p_2 lie on the same side of the line spanned by f as p; and

(iii) the point p_1 is further away from q than p_2.

The set Q has $2m$ points and the above conditions ensure that it is in general position.

Next, we define a coloring χ' and a χ'-halving matching for Q. For every edge pq induced by P, we color the four edges p_iq_j with $i, j \in \{1, 2\}$ with the same color as pq. Hence, the only edges remaining to be colored are the edges p_1p_2 between the duplicates of a point $p \in P$.

Let the point $p \in P$ be matched with the edge pq of P, and let $\overline{qp}$ be the line spanned by pq and directed from q to p. Further, let q_1 and q_2 be the points that originated from duplicating q, such that q_1 lies to the left of $\overline{qp}$ and q_2 lies to the right of $\overline{qp}$. We denote by $L_l(p)$ and $L_r(p)$ the number of edges in $L(p)$ to the left and right of $\overline{qp}$, respectively. Analogously, we denote by $S_l(p)$ and $S_r(p)$ the number edges in $S(p)$ to the left and right of $\overline{qp}$. Finally, we denote by $H_l(p)$ and $H_r(p)$ the number of edges incident to p that lie to the left and right of $\overline{qp}$, respectively, that have the same color as pq.

There are six cases in which p can fall, depending on the color of the edge pq matched with it and the numbers $L_l(p)$ and $L_r(p)$; see again Figure 2.1, where the larger color class is blue and the smaller red. In each case, the color of the edge p_1p_2 and the χ'-halving matching edges for p_1 and p_2 need to be determined.

In the first three cases pq is in the smaller color class $S(p)$ at p while in the last three cases pq is in the larger color class $L(p)$ at p. Note that in all cases, $L_l(p)$ and $L_r(p)$ differ by at most one. Furthermore, as $|P|$ is even, the degree of p is odd and hence $L(p)$ contains at least one

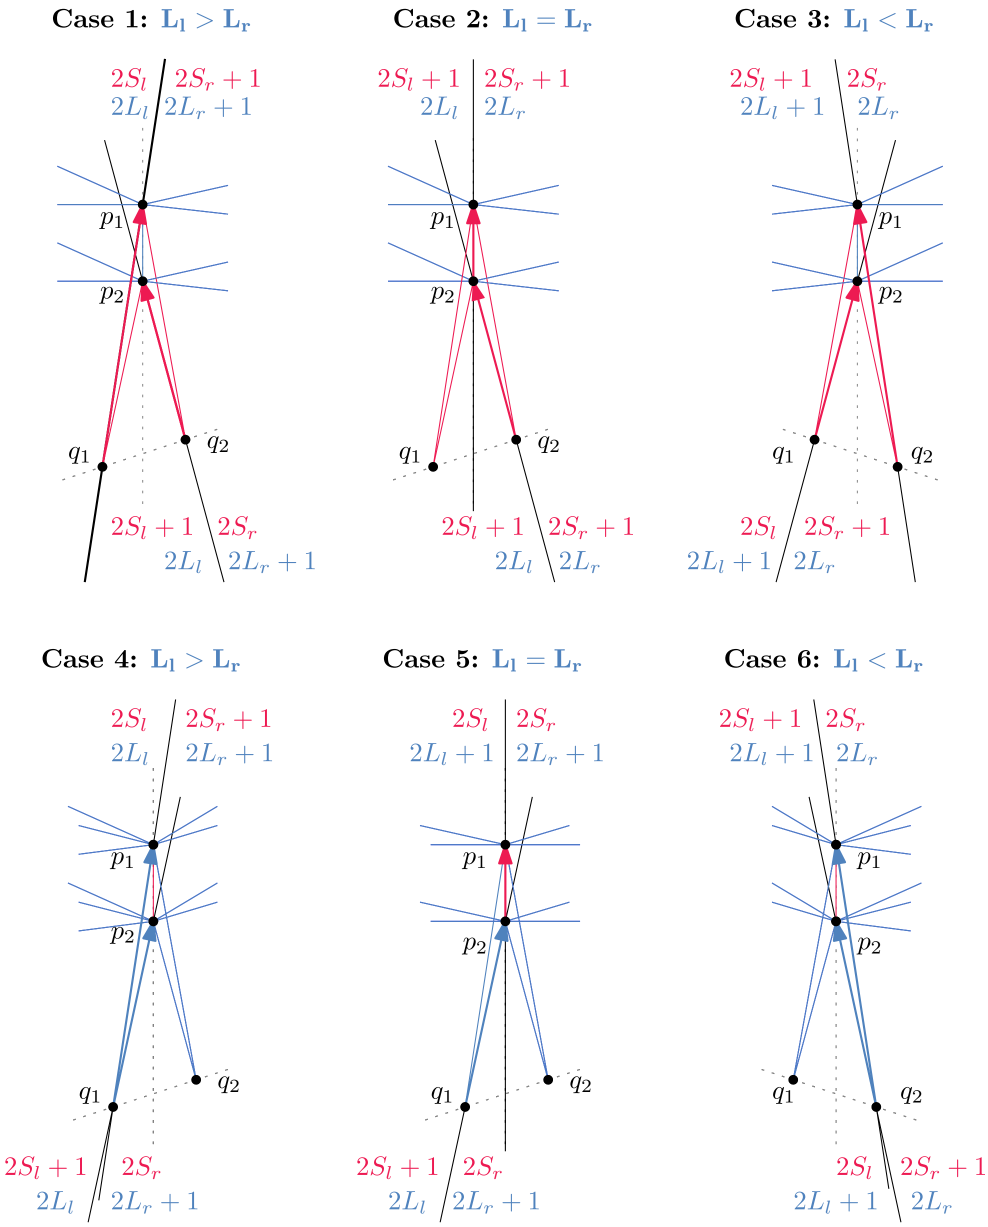

Figure 2.1.: The cases in the duplication process of Theorem 2.1 when the larger color class at p is blue. The dotted lines represent the lines spanned by the χ-halving matching edges for P. The numbers of blue (red) edges at p to the left and right of l_e, is denoted by L_l and L_r (S_l and S_r), respectively.

more edge than $S(p)$. Thus, the larger color class at p_1 and p_2 for χ' is the same as the one at p for χ.

Case 1: $pq \in S(p)$ and $L_l(p) > L_r(p)$. The edge p_1p_2 is colored with the color of $L(p)$. In the matching for Q, p_1 is matched with p_1q_1 and p_2 is matched with p_2q_2. By this we obtain $L_l(p_i) = 2L_l(p)$ and $L_r(p_i) = 2L_r(p) + 1$ for $i \in \{1, 2\}$, implying that the matched edges are indeed χ'-halving. Further, we have $S_l(p_1) = 2S_l(p)$, $S_r(p_1) = 2S_r(p) + 1$, $S_l(p_2) = 2S_l(p) + 1$, and $S_r(p_2) = 2S_r(p)$.

Case 2: $pq \in S(p)$ and $L_l(p) = L_r(p)$. The edge p_1p_2 is colored with the color of $S(p)$. In the matching for Q, p_1 is matched with p_1p_2 and p_2 is matched with p_2q_2. By this we obtain $L_l(p_i) = 2L_l(p)$ and $L_r(p_i) = 2L_r(p)$ for $i \in \{1, 2\}$, implying that the matched edges are χ'-halving. Further, $S_l(p_i) = 2S_l(p) + 1$ and $S_r(p_i) = 2S_r(p) + 1$ for $i \in \{1, 2\}$.

Case 3: $pq \in S(p)$ and $L_l(p) < L_r(p)$. The edge p_1p_2 is colored with the color of $L(p)$. In the matching for Q, p_1 is matched with p_1q_2 and p_2 is matched with p_2q_1. By this we obtain $L_l(p_i) = 2L_l(p) + 1$ and $L_r(p_i) = 2L_r(p)$ for $i \in \{1, 2\}$, implying that the matched edges are χ'-halving. Further, $S_l(p_1) = 2S_l(p) + 1$, $S_r(p_1) = 2S_r(p)$, $S_l(p_2) = 2S_l(p)$, and $S_r(p_2) = 2S_r(p) + 1$.

Case 4: $pq \in L(p)$ and $L_l(p) > L_r(p)$. The edge p_1p_2 is colored with the color of $S(p)$. In the matching for Q, p_1 is matched with p_1q_1 and p_2 is matched with p_2q_1. By this we obtain $L_l(p_i) = 2L_l(p)$ and $L_r(p_i) = 2L_r(p) + 1$ for $i \in \{1, 2\}$, implying that the matched edges are indeed χ'-halving. Further, we have $S_l(p_1) = 2S_l(p)$, $S_r(p_1) = 2S_r(p) + 1$, $S_l(p_2) = 2S_l(p) + 1$, and $S_r(p_2) = 2S_r(p)$.

Case 5: $pq \in L(p)$ and $L_l(p) = L_r(p)$. The edge p_1p_2 is colored with the color of $S(p)$. In the matching for Q, p_1 is matched with p_1p_2 and p_2 is matched with p_2q_1. By this we obtain $L_l(p_1) = 2L_l(p)+1$, $L_r(p_1) = 2L_r(p)+1$, $L_l(p_2) = 2L_l(p)$, and $L_r(p_2) = 2L_r(p)+1$. Hence, the matched edges are χ'-halving. Further, we have $S_l(p_1) = 2S_l(p)$, $S_r(p_1) = 2S_r(p)$, $S_l(p_2) = 2S_l(p) + 1$, and $S_r(p_2) = 2S_r(p)$.

Case 6: $pq \in L(p)$ and $L_l(p) < L_r(p)$. The edge p_1p_2 is colored with the color of $S(p)$. In the matching for Q, p_1 is matched with p_1q_2 and p_2 is matched with p_2q_2. By this we obtain $L_l(p_i) = 2L_l(p) + 1$ and $L_r(p_i) = 2L_r(p)$ for $i \in \{1, 2\}$, implying that the matched edges are indeed χ'-halving. Further, we have $S_l(p_1) = 2S_l(p) + 1$, $S_r(p_1) = 2S_r(p)$, $S_l(p_2) = 2S_l(p)$, and $S_r(p_2) = 2S_r(p) + 1$.

Having completed the coloring χ' for the edges induced by Q, we next consider the number of monochromatic crossings in the resulting drawing on Q.

Claim 1. *The pair (Q, χ') satisfies*

$$\begin{aligned}\overline{\mathrm{cr}}_2(Q, \chi') = 16\ \overline{\mathrm{cr}}_2(P, \chi) + \binom{m}{2} - m \\ + 4 \sum_p \left(\binom{L_l(p)}{2} + \binom{L_r(p)}{2} + \binom{S_l(p)}{2} + \binom{S_r(p)}{2} \right) + 2 \sum_p (H_l(p) + H_r(p)).\end{aligned}$$

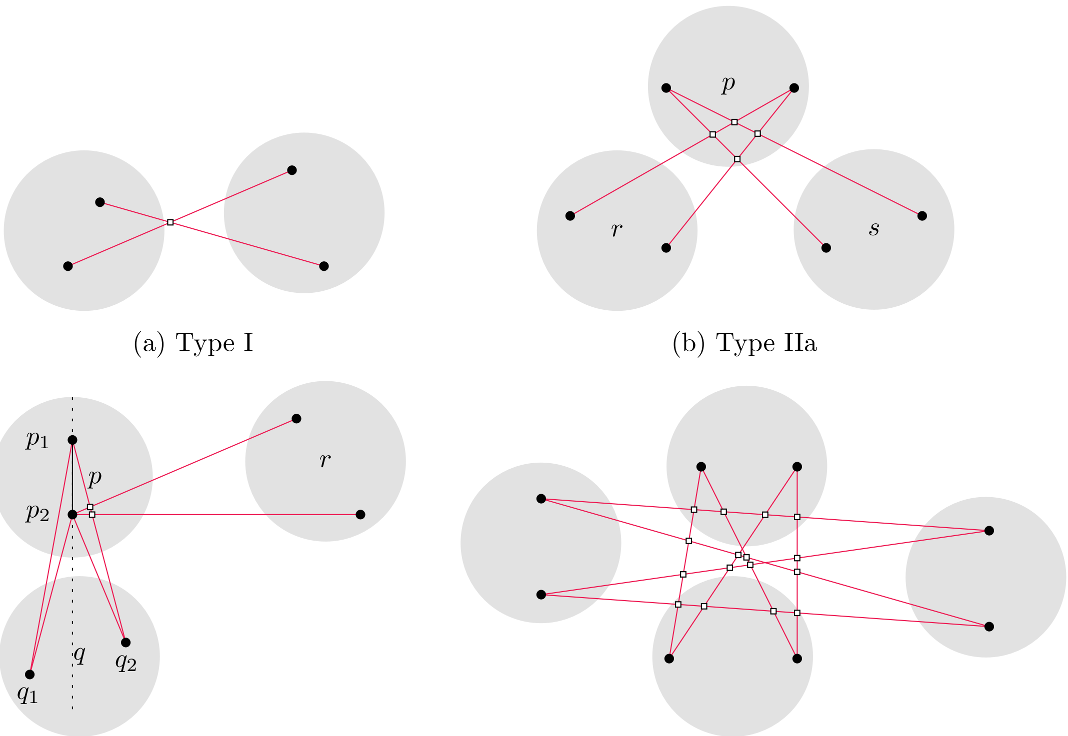

(a) Type I (b) Type IIa

(c) Type IIb (d) Type III

Figure 2.2.: Counting the crossings of different types in the duplication process. The gray circles represent the close neighborhood of the points before the duplication.

The proof of this claim follows the same counting technique used in [6].

Proof. We count the crossings in the same way as in the proof of Lemma 3 of [6]. A crossing in Q comes from four points in convex position. We classify the crossings in three types, according to the number of points in P that originated them; see Figure 2.2.

Type I: The points originate from duplicating two points in P; see Figure 2.2(a). There are $\binom{m}{2}$ ways of choosing a pair of points in P, and every such pair determines a crossing in Q unless the edge between them is a matching edge. Since we have m matching edges, there are

$$\binom{m}{2} - m$$

crossings of this type.

Type IIa: The points originate from duplicating three points p, r, and s in P and none of the edges between those points is a matching edge; see Figure 2.2(b). Without loss of generality, p_1 and p_2 are involved in the crossing. Then r and s lie on the same side of the line spanned by the matching edge e of p and both pr and ps have the same color.

Any pair (r, s) of points of P that satisfies those conditions with respect to p generates four crossings in Q. Thus, the number of Type IIa crossings for p is

$$4\left[\binom{L_l(p)}{2} + \binom{L_r(p)}{2} + \binom{S_l(p)}{2} + \binom{S_r(p)}{2}\right].$$

Type IIb: The points originate from duplicating three points p, q, and r in P and one of the edges between those points is a matching edge; see Figure 2.2(c). Without loss of generality, assume pq is the matching edge of p. Any pair of points that originated from a point $r \in P$ such that pr has the same color as pq generates two crossings with either p_1q_1 or p_1q_2. Thus, the number of Type IIb crossings for p is

$$2(H_l(p) + H_r(p)).$$

Type III: The points originate from duplicating four points p, q, r, and s in P that generate a crossing in P; see Figure 2.2(d). There are $\overline{\text{cr}}(P, \chi)$ such quadruples of points, and each one generates 16 crossings in Q. Thus, the number of Type III crossings in Q is

$$16\overline{\text{cr}}(P, \chi).$$

Summing the Type II crossings over each point p of P and adding them to the crossings of Type III and Type I gives the claimed result. □

We now apply the duplication process multiple times. To this end, consider again the six different cases for a point $p \in P$ when obtaining a coloring and a matching for Q. Note that if one of the Cases 1, 2, 3, 4 and 6 applies for p, then the same case applies for its duplicates $p_1, p_2 \in Q$ (and will apply in all further duplication iterations). If p falls in Case 5, then for p_1 and p_2 we have Case 2 and 4, respectively. As no point in Q falls in Case 5, from now on, we assume that P is such that no point of P falls in Case 5 either.

Let $k \geq 1$ be an integer and let (Q_k, χ_k) be the pair obtained by iterating the duplication process k times, with $(Q_0, \chi_0) = (P, \chi)$. We claim the following on $\overline{\text{cr}}_2(Q_k, \chi_k)$, the number of monochromatic crossings in the 2-edge-colored drawing of K_n induced by Q_k and χ_k:

Claim 2. *After k iterations of the duplication process, it holds that*

$$\overline{\text{cr}}_2(Q_k, \chi_k) = A \cdot 2^{4k} + B \cdot 2^{3k} + C \cdot 2^{2k} + D \cdot 2^k$$

where A, B, C and D are rational numbers that depend on P and its χ-halving matching.

The proof of this claim uses a careful analysis of the structure of (Q_k, χ_k) in dependence of (P, χ) and the χ-halving matching for P as well as involved calculations.

Proof. Let p be a point of P. We iteratively construct a rooted binary tree $T(p)$ of height k containing a vertex for each point q of Q_i that stems from duplicating p in the following way. The root of $T(p)$ contains the tuple $(L_l(p), L_r(p), S_l(p), S_r(p))$ representing p. For vertex v in $T(p)$ that represents a point q of Q_i with $0 \leq i \leq k-1$, its left child contains the tuple $(L_l(q_1), L_r(q_1), S_l(q_1), S_r(q_1))$ and its right child contains the tuple $(L_l(q_2), L_r(q_2), S_l(q_2), S_r(q_2))$, where $q_1, q_2 \in Q_{i+1}$ are the duplicates of q. In addition, we mark whether the matching edge of p (and hence the ones of all points originating from p) is of the larger or the smaller color class at p.

We next elaborate on the exact content of the tuple stored in the j-th vertex of the i-th level of $T(P)$ with $j \in \{1, \dots, 2^i\}$, depending on the case to be applied for p in the duplication process.

Cases 1, 3, 4 and 6: Let p be a point in P that falls in Case 1. Then in the i-th level of $T(p)$, the j-th vertex contains the tuple

$$(2^i L_l(p), 2^i L_r(p) + 2^i - 1, 2^i S_l(p) + j - 1, 2^i S_r(p) + 2^i - j).$$

We show this by induction on i. It follows directly from the duplication process that it is true when $i = 1$. Suppose that $i > 1$. From the induction hypothesis, the j-th vertex v of level i contains the tuple $(2^i L_l(p), 2^i L_r(p) + 2^i - 1, 2^i S_l(p) + j - 1, 2^i S_l(p) + 2^i - j)$. Since all the vertices of $T(p)$ represent points that fall in Case 1, the left and right children of v contain the tuples

$$(2^{i+1} L_l(p), 2^{i+1} L_r(p) + 2^{i+1} - 1, 2^{i+1} S_l(p) + (2j-1) - 1, 2^{i+1} S_r(p) + 2^{i+1} - (2j-1))$$

and

$$(2^{i+1} L_l(p), 2^{i+1} L_r(p) + 2^{i+1} - 1, 2^{i+1} S_l(p) + (2j-1), 2^{i+1} S_r(p) + 2^{i+1} - 2j),$$

respectively. These two vertices are precisely the $(2j-1)$-st and the $2j$-th vertex in level $i+1$.

Note that, if p falls in Case 4, $T(p)$ has the exact same structure as a point of Case 1. Furthermore, if p is a point that falls in Case 3 or Case 6, the structure of $T(p)$ is exactly a mirrored version of the tree from a point that falls in Case 1.

Case 2: Let p be a point in P that falls in Case 2. Then in the i-th level of $T(p)$, the j-th vertex contains the tuple

$$(2^i L_l(p), 2^i L_r(p), 2^i S_l(p) + 2^i - 1, 2^i S_r(p) + 2^i - 1).$$

We again proceed by induction on i. It follows directly from the duplication process that this happens when $i = 1$, so suppose that $i > 1$. From the induction hypothesis, the j-th vertex v of level i contains the tuple $(2^i L_l(p), 2^i L_r(p) + 2^i - 1, 2^i S_l(p) + j, 2^j S_r(p) + 2^j - i)$. Since all the vertices of $T(p)$ represent points that fall in Case 2, the left and right children of v contain the tuple

$$(2^{i+1} L_l(p), 2^{i+1} L_r(p), 2^{i+1} S_l(p) + 2^{i+1} - 1, 2^{i+1} S_r(p) + 2^{i+1} - 1).$$

Note that $T(p)$, together with the information whether the matching edges are of the smaller or the larger color class, contains all the information needed to compute the crossings of Type II in Q_{i+1} that involve points which originate from p.

Using the above observations we can now determine $\overline{\mathrm{cr}}_2(Q_k, \chi_k)$. We use the following notation: $f_i(x) = \binom{2^i x}{2}$, $g_i(x) = \binom{2^i x + 2^i - 1}{2}$, $h_{i,j}(x) = \binom{2^i x + j}{2}$, and P_C is the subset of P of points that fall in Case C.

Type III: Each crossing of Type III in P generates 16 crossings in Q. Iterating this process k times, we obtain

$$16^k \ \overline{\mathrm{cr}}_2(P, \chi)$$

crossings in Q_k.

Type I: Every set Q_i has a χ_i-halving matching and $|Q_i| = 2^i m$, thus, there are $\binom{2^i m}{2} - 2^i m$ crossings of Type I in Q_{i+1}. Moreover, each of these crossings becomes a Type III crossing in further duplication steps, that is, it produces 16 crossings per each further duplication step. Hence, adding the crossings of Type I that we get at each iteration and the according crossings of Type III that they generate later, we obtain

$$\sum_{i=0}^{k-1} 16^{k-i-1} \left[\binom{2^i m}{2} - 2^i m \right]$$

crossings in Q_k.

Type II for Case 2: Consider a point $p \in P$ that falls in Case 2 together with all the points in Q_i that originate from it (and hence fall in Case 2 as well). Using Claim 1 and the information from the i-th level of $T(p)$, we obtain that Q_{i+1} has

$$\begin{aligned} &4 \cdot 2^i \left[f_i(L_l(p)) + f_i(L_r(p)) + g_i(S_l(p)) + g_i(S_r(p)) \right] \\ &+ 2 \cdot 2^i \left[2^i (H_l(p) + H_r(p)) + 2^{i+1} - 2 \right] \end{aligned}$$

crossings of Type II that come from all points in Q_i originating from p. Moreover, each of these crossings becomes a Type III crossing for all further duplication steps. Hence, adding the crossings of Type II that we count for points originating from p at each iteration and the according crossings of Type III that they generate later, we obtain

$$\begin{aligned} &4 \sum_{i=0}^{k-1} 16^{k-i-1} 2^i \left[f_i(L_l(p)) + f_i(L_r(p)) + g_i(S_l(p)) + g_i(S_r(p)) \right] \\ &+ 2 \sum_{i=0}^{k-1} 16^{k-i-1} 2^i \left[2^i (H_l(p) + H_r(p)) + 2^{i+1} - 2 \right] \end{aligned}$$

crossings in Q_k.

Type II for Cases 1 and 3: Consider a point $p \in P$ that falls in Case 1 or 3, and with all points in Q_i that originate from it (and hence fall in Case 1 or Case 3 as well). Using Claim 1, and the information from the i-th level of $T(p)$, we obtain that Q_{i+1} has

$$\begin{aligned}&4 \cdot 2^i \left[f_i(L_l(p)) + g_i(L_r(p))\right]\\ &+ 4 \sum_{j=0}^{2^i-1} \left[h_{i,j}(S_l(p)) + h_{i,j}(S_r(p))\right]\\ &+ 2 \sum_{j=0}^{2^i-1} \left[2^i(H_l(p) + H_r(p)) + 2j\right]\end{aligned}$$

crossings of Type II that come from all points in Q_i in the tree $T(p)$. Again, each of these crossings becomes a Type III crossing for all further duplication steps. Hence, adding the crossings of Type II that we we count for points originating from p at each iteration and the according crossings of Type III that they generate later, we obtain

$$\begin{aligned}&4 \sum_{i=0}^{k-1} 16^{k-i-1} 2^i \left[f_i(L_l(p)) + g_i(L_r(p))\right]\\ &+ 4 \sum_{i=0}^{k-1} 16^{k-i-1} \sum_{j=0}^{2^i-1} \left[h_{i,j}(S_l(p)) + h_{i,j}(S_r(p))\right]\\ &+ 2 \sum_{i=0}^{k-1} 16^{k-i-1} \sum_{j=0}^{2^i-1} \left[2^i(H_l(p) + H_r(p)) + 2j\right]\end{aligned}$$

crossings in Q_k.

Type II for Cases 4 and 6: Consider a point $p \in P$ that falls in Case 4 or 6, and with all points in Q_i that originate from it (and hence fall in Case 4 or Case 6 as well). Using Claim 1, and the information from the i-th level of $T(p)$, we obtain that Q_{i+1} has

$$\begin{aligned}&4 \cdot 2^i \left[f_i(L_l(p)) + g_i(L_r(p))\right]\\ &+ 4 \sum_{j=0}^{2^i-1} \left[h_{i,j}(S_l(p)) + h_{i,j}(S_r(p))\right]\\ &+ 2 \cdot 2^i \left[2^i(H_l(p) + H_r(p)) + 2^i - 1\right]\end{aligned}$$

crossings of Type II that come from all points in Q_i in the tree $T(p)$. Again, each of these crossings becomes a Type III crossing for all further duplication steps. Hence, adding the crossings of Type II that we we count for points originating from p at each iteration and the according crossings of Type III that they generate later, we obtain

$$4 \sum_{i=0}^{k-1} 16^{k-i-1} 2^i \left[f_i(L_l(p)) + g_i(L_r(p))\right]$$

$$+ 4\sum_{i=0}^{k-1} 16^{k-i-1} \sum_{j=0}^{2^i-1} \left[h_{i,j}(S_l(p)) + h_{i,j}(S_r(p))\right]$$
$$+ 2\sum_{i=0}^{k-1} 16^{k-i-1} 2^i \left[2^i(H_l(p) + H_r(p)) + 2^i - 1\right]$$

crossings in Q_k.

Adding the number of crossings of each type, we obtain the following expression for $\overline{\mathrm{cr}}(Q_k, \chi_k)$, the number of monochromatic crossings in the straight-line 2-edge-colored drawing of the complete graph induced by Q_k and χ_k:

$$\overline{\mathrm{cr}}_2(Q_k,\chi_k) = 16^k\ \overline{\mathrm{cr}}_2(P,\chi) + \sum_{i=0}^{k-1} 16^{k-i-1}\left[f_i(m) - 2^i m\right] \tag{2.1}$$

$$+ \sum_{p\in P_2}\Bigg[4\sum_{i=0}^{k-1} 16^{k-i-1} 2^i\left[f_i(L_l(p)) + f_i(L_r(p)) + g_i(S_l(p)) + g_i(S_r(p))\right] \tag{2.2}$$
$$+ 2\sum_{i=0}^{k-1} 16^{k-i-1} 2^i [2^i(H_l(p) + H_r(p)) + 2^{i+1} - 2]\Bigg]$$

$$+ \sum_{p\in P_1\cup P_3}\Bigg[4\sum_{i=0}^{k-1} 16^{k-i-1} 2^i [f_i(L_l(p)) + g_i(L_r(p))] \tag{2.3}$$
$$+ 4\sum_{i=0}^{k-1} 16^{k-i-1} \sum_{j=0}^{2^i-1} [h_{i,j}(S_l(p)) + h_{i,j}(S_r(p))]$$
$$+ 2\sum_{i=0}^{k-1} 16^{k-i-1} \sum_{j=0}^{2^i-1} [2^i(H_l(p) + H_r(p)) + 2j]\Bigg]$$

$$+ \sum_{p\in P_4\cup P_6}\Bigg[4\sum_{i=0}^{k-1} 16^{k-i-1} 2^i [f_i(L_l(p)) + g_i(L_r(p))] \tag{2.4}$$
$$+ 4\sum_{i=0}^{k-1} 16^{k-i-1} \sum_{j=0}^{2^i-1} [h_{i,j}(S_l(p)) + h_{i,j}(S_r(p))]$$
$$+ 2\sum_{i=0}^{k-1} 16^{k-i-1} 2^i [2^i(H_l(p) + H_r(p)) + 2^i - 1]\Bigg].$$

It remains show that this sum can be written as $A \cdot 2^{4k} + B \cdot 2^{3k} + C \cdot 2^{2k} + D \cdot 2^k$, where A, B, C, and D depend on $L_l(p), L_r(p), S_l(p), S_r(p), H_l(p)$, and $H_r(p)$ for every point p in P. For that, we use the following observations:

Observation 2.2.

$$\begin{aligned}\sum_{i=0}^{k-1} 16^{k-i-1} 2^i f_i(x) &= \frac{2^{4k}}{2^4} \sum_{i=0}^{k-1} \frac{1}{2^{3i}} 2^{i-1} x(2^i x - 1) \\ &= \frac{2^{4k}}{2^4} \sum_{i=0}^{k-1} \frac{x^2}{2 \cdot 2^i} - \frac{x}{2 \cdot 2^{2i}} \\ &= \frac{2^{4k}}{2^4} \left[\frac{3x^2 - 2x}{3} - \frac{x^2}{2^k} + \frac{2x}{3 \cdot 2^{2k}} \right] \\ &= \frac{3x^2 - 2x}{48} \cdot 2^{4k} - \frac{x^2}{16} \cdot 2^{3k} + \frac{x}{24} \cdot 2^{2k}.\end{aligned}$$

Observation 2.3.

$$\begin{aligned}\sum_{i=0}^{k-1} 16^{k-i-1} 2^i g_i(x) &= \frac{2^{4k}}{2^4} \sum_{i=0}^{k-1} \frac{1}{2^{3i}} (2^i x + 2^i - 1)(2^{i-1} x + 2^{i-1} - 1) \\ &= \frac{2^{4k}}{2^4} \sum_{i=0}^{k-1} \frac{(x+1)^2}{2 \cdot 2^i} - \frac{3(x+1)}{2 \cdot 2^{2i}} + \frac{1}{2^{3i}} \\ &= \frac{2^{4k}}{2^4} \left[\frac{7x^2+1}{7} - \frac{(x+1)^2}{2^k} + \frac{2(x+1)}{2^{2k}} - \frac{8}{7 \cdot 2^{3k}} \right] \\ &= \frac{7x^2+1}{112} \cdot 2^{4k} - \frac{(x+1)^2}{16} \cdot 2^{3k} + \frac{x+1}{8} \cdot 2^{2k} - \frac{1}{14} \cdot 2^k.\end{aligned}$$

Observation 2.4.

$$\begin{aligned}\sum_{j=0}^{2^i-1} h_{i,j}(x) &= \frac{1}{2} \sum_{j=0}^{2^i-1} 2^{2i} x^2 + 2^i x(2j-1) + j(j-1) \\ &= \frac{1}{2} \left[2^i \cdot 2^{2i} x^2 + 2^i x(2^{2i} - 2^{i+1}) + \frac{2^i(2^i-1)(2^i-2)}{3} \right] \\ &= \frac{3(x^2+x)+1}{6} \cdot 2^{3i} - \frac{2x+1}{2} \cdot 2^{2i} + \frac{1}{3} \cdot 2^i.\end{aligned}$$

Observation 2.5.

$$\begin{aligned}&\sum_{i=0}^{k-1} 16^{k-i-1} \sum_{j=0}^{2^i-1} h_{i,j}(x) \\ &\quad= \frac{2^{4k}}{2^4} \sum_{i=0}^{k-1} \frac{1}{2^{4i}} \left[\frac{3(x^2+x)+1}{6} \cdot 2^{3i} - \frac{2x+1}{2} \cdot 2^{2i} + \frac{1}{3} \cdot 2^i \right] \\ &\quad= \frac{2^{4k}}{2^4} \sum_{i=0}^{k-1} \frac{1}{2^{4i}} \left[\frac{3(x^2+x)+1}{6 \cdot 2^i} - \frac{2x+1}{2 \cdot 2^{2i}} + \frac{1}{3 \cdot 2^{3i}} \right] \\ &\quad= \frac{2^{4k}}{2^4} \left[\frac{3(x^2+x)+1}{48}(2^{4k} - 2^{3k}) - \frac{2x+1}{24}(2^{4k} - 2^{2k}) + \frac{1}{42}(2^{4k} - 2^k) \right]\end{aligned}$$

$$= \frac{21x^2 - 7x + 1}{336} \cdot 2^{4k} - \frac{3(x^2 + x) + 1}{48} \cdot 2^{3k} + \frac{2x + 1}{24} \cdot 2^{2k} - \frac{1}{42} \cdot 2^k.$$

We show that (2.1), (2.2), (2.3) and (2.4) can be written as

$$a \cdot 2^{4k} + b \cdot 2^{3k} + c \cdot 2^{2k} + d \cdot 2^k.$$

For (2.1), it follows from Observation 2.2. For (2.2), it follows from Observations 2.2 and 2.3. For (2.3) and (2.4), it follows from Observations 2.2, 2.3 and 2.5. Thus, $\overline{\text{cr}}_2(Q_k, \chi_k)$ can be written as $A \cdot 2^{4k} + B \cdot 2^{3k} + C \cdot 2^{2k} + D \cdot 2^k$. □

Applying Claim 2 to an initial drawing on m vertices and letting $n = 2^k m$, we get

$$\overline{\text{cr}}_2(K_n) \leq \overline{\text{cr}}_2(Q_k, \chi_k) = \frac{24A}{m^4}\binom{n}{4} + \Theta(n^3),$$

which completes the proof of Theorem 2.1 when n is of the form $2^k m$. The proof for $2^k m < n < 2^{k+1} m$ then follows from the fact that $\overline{\text{cr}}_2(K_n)$ is an increasing function. □

Observe that the proof of Theorem 2.1 gives us a closed formula to calculate the number of monochromatic crossings after k duplications of a given point set (together with a 2-edge-coloring χ and a χ-halving matching) following the above-described procedure.

We remark that the duplication process described in the proof of Theorem 2.1 can also be applied if the initial set P has odd cardinality. However, then it might happen that the resulting matching is not χ'-halving for the resulting set Q. Moreover, a similar process can even be applied with any matching between the points of P and edges induced by P, where in that situation one needs to specify how the colors for the edges between duplicates of points (and possibly a matching for the resulting set) is chosen.

In the uncolored duplication process for obtaining bounds on $\overline{\text{cr}}(K_n)$, halving matchings always yield the best asymptotic behavior, which only depends on $|P|$ and $\overline{\text{cr}}(P)$. This is not the case for the 2-colored setting, where we ideally would like to achieve simultaneously for every point $p \in P$ that (i) both color classes are of similar size, (ii) both color classes are evenly split by the matching edge, and (iii) $\overline{\text{cr}}_2(P)$ is small. Yet, this is in general not possible. Starting with a χ-halving matching for P we obtain (ii) at least for the larger color class at every point of P. Moreover, this is hereditary by the design of our duplication process.

2.2.1. Small configurations

The previous section implies that for large cardinality we can obtain straight-line drawings of the complete graph with a reasonably small 2-colored crossing number by starting from *good* sets of constant size. Thus, in this section we describe how to obtain those small good sets.

Similar as in [18] we combine three different methods to obtain straight-line drawings of the complete graph with few monochromatic crossings. Our heuristic iterates three steps of (1) locally improving a set, (2) generating larger good sets, and (3) extracting good subsets, where also after steps (2) and (3) a local optimization is done.

For reducing the crossing number of a given drawing we use a combined approach. Similar to [18] we use a simple version of simulated annealing to optimize the placement of the points in the given drawing. Whenever the position of a point is changed, we in addition use a heuristic to recolor the edges of the drawing in order to reduce the number of monochromatic crossings. By iterating these steps and also allowing changes which (temporarily) increase the rectilinear number of monochromatic crossings (as typical for approaches based on simulated annealing), we can obtain better drawings for fixed cardinality. As discussed in the introduction, this approach combines the optimization of two hard problems and will thus in general not lead to a global optimum.

The tool for step (2) is actually implicitly given by the proof of Theorem 2.1. This proof not only provides the bounds for the resulting crossing number, but also gives the details how to construct a duplicated point set with twice the number of points. We used an implementation developed for [18] and extended it by also determining the color of the edges of the complete graph on the point set as described in the proof.

The third tool is to compute subsets of a given drawing. That is, we remove one vertex from the drawing and consider the induced subdrawing. Again, a heuristic is used to find a good coloring of the remaining edges. The process of removing a vertex is repeated for all vertices of the given drawing and we select the best obtained subdrawing of cardinality $n-1$. Note that again this requires more computations and care than in the uncolored case, where an optimal point to be removed can be identified rather easily: For each point one can add up the number of crossings in which its incident edges are involved and take any of the points for which this sum is maximized. Alternatively, we can also remove a (small) constant number of vertices at the same time and recolor the remaining edges. The process of removing one (or more) vertex (vertices) can of course be iterated to obtain smaller sets with few monochromatic crossings.

We use the three methods in alternating order by starting with a known set, apply the duplication process from Theorem 2.1 to obtain a larger set, locally optimize it to get a better set, find good subsets, locally optimize them, duplicate the resulting sets and so on. In this way it can actually happen that after a few iterations we end up with a set of the same cardinality as the starting set, but with less monochromatic crossings. Our experiments show clearly that this combined heuristic leads to significantly better sets than just local optimization.

The currently best (with respect to the crossing constant, see below) straight-line drawing D with 2-edge coloring χ we found[1] has $n = 135$ vertices, the number of monochromatic crossings is $\overline{\text{cr}}_2(D,\chi) = 1470756$, and it contains a χ-halving matching.

[1]The interested reader can get a file with the coordinates of the points, the colors of the edges, and a χ-halving matching from `http://www.crossingnumbers.org/projects/monochromatic/sets/n135.php`.

2.2.2. Rectilinear 2-colored crossing constant

Let $\overline{\text{cr}}_2$ be the *rectilinear* 2*-colored crossing constant*, that is, the constant such that the best straight-line drawing of K_n for large values of n has at most $\overline{\text{cr}}_2\binom{n}{4}$ monochromatic crossings. Its existence follows from the fact that the limit $\lim_{n\to\infty} \overline{\text{cr}}_2(K_n)/\binom{n}{4}$ exits and is a positive number. The proof is essentially the same as for the (rectilinear) crossing constant [172]. Moreover, it shows that the sequence $\overline{\text{cr}}_2(K_n)/\binom{n}{4}$ is an increasing sequence. For completeness, we include the proof here.

Theorem 2.6. *The limit* $\lim_{n\to\infty} \frac{\overline{\text{cr}}_2(K_n)}{\binom{n}{4}}$ *exists and is a positive number.*

Proof. First, in any straight-line drawing of a graph, the edges incident to a subset of four vertices produce at most one (monochromatic) crossing. This shows that $\overline{\text{cr}}_2(K_n)/\binom{n}{4}$ is at most one. Moreover, as we will show in Section 2.3, $\overline{\text{cr}}_2(9) = 2$. Thus, to show that the limit exists and is positive it suffices now to show that $\overline{\text{cr}}_2(K_n)/\binom{n}{4} \geq \overline{\text{cr}}_2(K_{n-1})/\binom{n-1}{4}$. We can rewrite this inequality as $(n-4)\overline{\text{cr}}_2(K_n) \geq n\overline{\text{cr}}_2(K_{n-1})$. Consider a straight-line drawing D of K_n that has exactly $\overline{\text{cr}}_2(K_n)$ monochromatic crossings. If we remove one vertex (and its incident edges) we get a drawing with at least $\overline{\text{cr}}_2(K_{n-1})$ monochromatic crossings. Doing that for every vertex we count every monochromatic crossing exactly $n-4$ times and we get that the number of monochromatic crossings in D is at least $\overline{\text{cr}}_2(K_n) \geq \frac{n}{n-4}\overline{\text{cr}}_2(K_{n-1})$, as we wanted to show. □

The results in this section imply that we can derive an upper bound for the rectilinear 2-colored crossing constant from a given set of constant size. Taking the set of 135 points obtained in Section 2.2.1, together with its 2-coloring and its matching, and duplicating it once, we obtain a point set of even cardinality with a coloring and matching, where Case 5 does not show up. Plugging the values of this set into the machinery developed in the proof of Theorem 2.1 we get the upper bound of $\overline{\text{cr}}_2 < 0.11798016$.

Theorem 2.7. *The rectilinear 2-colored crossing constant satisfies*

$$\overline{\text{cr}}_2 \leq \frac{182873519}{1550036250} < 0.11798016.$$

In [7] a lower bound of $\overline{\text{cr}} \geq \frac{277}{729} > 0.37997267$ has been shown for the rectilinear crossing constant. We can thus give an upper bound on the asymptotic ratio between the best rectilinear 2-colored drawing of K_n and the best rectilinear drawing of K_n of $\overline{\text{cr}}_2/\overline{\text{cr}} \leq 0.31049652$.

2.3. Lower bounds on $\overline{\text{cr}}_2(K_n)$

In this section we consider lower bounds for the rectilinear 2-colored crossing number and the biplanar crossing number of K_n.

In related work [159], the authors present lower and upper bounds on the $\sup \overline{\text{cr}}_k(G)/\overline{\text{cr}}(G)$ where the supremum is taken over all non-planar graphs. We remark that this lower bound

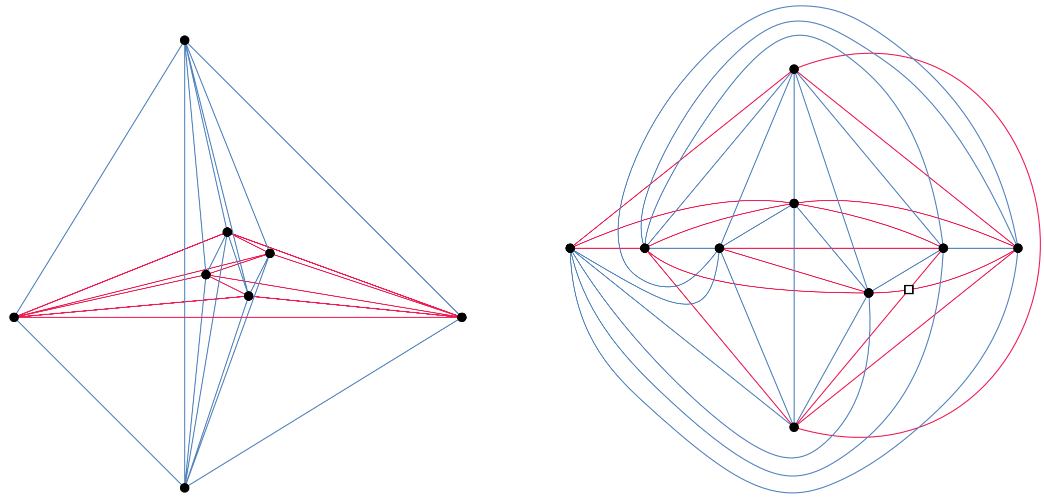

Figure 2.3.: Left: A 2-edge-colored rectilinear drawing of K_8 without monochromatic crossings. Right: A 2-edge-colored drawing of K_9 with only one monochromatic (red) crossing marked with a square.

does not yield a lower bound for $\overline{\mathrm{cr}}_2(K_n)$ as their bound is obtained for "midrange" graphs (graphs with a subquadratic but superlinear number of edges). Czabarka et al. mention a lower bound on the biplanar crossing number of general graphs depending on the number of edges [71, Equation 3]. For the complete graph, this yields a lower bound of $\overline{\mathrm{cr}}_2(K_n) \geq 1/1944n^4 - O(n^3)$.

A better bound can be obtained from the crossing lemma [145] in the following way. A version of the improved crossing lemma [9] states that for an undirected simple graph with n vertices and e edges with $e > 7n$, the crossing number of the graph is at least $\frac{e^3}{29n^2}$. Consider a general (not even necessarily simple) 2-edge-colored drawing of K_n with $e = \binom{n}{2}$ edges. If for $0 < \alpha < 1$ the two color classes have αe and $(1-\alpha)e$ edges, respectively, then a lower bound for the biplanar crossing number $\mathrm{cr}_2(K_n)$ provided by the crossing lemma is obtained for $\alpha = 1 - \alpha = 1/2$. In that case $\mathrm{cr}_2(K_n) \geq 2\frac{(n(n-1)/4)^3}{29n^2} = \frac{1}{32}\frac{(n(n-1))^3}{29n^2}$. This implies that $\overline{\mathrm{cr}}_2 \geq \frac{24}{29\cdot 32} = 3/116 > 1/39$.

Alternatively, the next result shows that from the 2-colored rectilinear crossing number of small sets we can obtain lower bounds for larger sets.

Lemma 2.8. *Let $\overline{\mathrm{cr}}_2(m) = \hat{c}$ for some $m \geq 4$. Then for $n > m$ we have*

$$\overline{\mathrm{cr}}_2(K_n) \geq \frac{24\hat{c}}{m(m-1)(m-2)(m-3)}\binom{n}{4}$$

which implies

$$\overline{\mathrm{cr}}_2 \geq \frac{24\hat{c}}{m(m-1)(m-2)(m-3)}.$$

Proof. Every subset of m points of K_n induces a drawing with at least $\hat{c}$ crossings, and thus we have $\hat{c}\binom{n}{m}$ crossings in total. In this way every crossing is counted $\binom{n-4}{m-4}$ times. This results in a total of $\frac{24\hat{c}}{m(m-1)(m-2)(m-3)}\binom{n}{4}$ crossings. □

We next determine $\overline{\text{cr}}_2(K_9)$, as K_8 can be drawn such that $\overline{\text{cr}}_2(K_8) = 0$; see Figure 2.3 (left). We use the optimization heuristic mentioned from Section 2.2 to obtain good colorings for all 158 817 order types of K_9 (which are provided by the order type data base [11]). In this way, it is guaranteed that all (crossing-wise) different straight-line drawings of K_9 (uncolored) are considered.

To prove that the heuristics indeed found the best colorings we consider the intersection graph for each drawing D. In the intersection graph every edge in D is a vertex, and two vertices are connected if their edges in D cross. Note that each odd cycle in the intersection graph of D gives rise to a monochromatic crossing in D. On the other hand, several odd cycles might share a crossing and only one monochromatic crossing is forced by them.

We thus set up an integer linear program, where for every crossing of D we have a non-negative variable and for each odd cycle the sum of the variables corresponding to the crossings of the cycle has to be at least one. The objective function aims to minimize the sum of all variables, which by construction is a lower bound for the number of monochromatic crossings in D.

Notice that in the linear relaxation, arising from dropping the integrality constraint, no variable in an optimal solution will be larger than one. Therefore, as linear programs usually converge much faster than integer programs, we used a two stage approach. We first check the drawing D with the linear program, and if the result differs by less than one from the number of monochromatic crossings obtained by the coloring from the heuristic, the 2-colored crossing number of D is determined and we are done. Otherwise, we use the integer program. Another way to speed up the computations is to consider odd cycles only up to a predefined length. Again, this can be done in an iterative way. Start with a linear program which considers only 3- and 5-cycles, and only if this is not sufficient, it adds 7-cycles and so on.

In the described way we have been able to determine all the 2-colored crossing numbers of all drawings of K_9 within a few hours. The longest cycles we considered have been 7-cycles, and only for two cases we had to use integer programming (see [95] for details). The best drawings we found have 2 monochromatic crossings, and thus $\overline{\text{cr}}_2(K_9) = 2$. Using Lemma 2.8 for $m = 9$ and $\hat{c} = 2$ we get a bound of $\overline{\text{cr}}_2 \geq 1/63$, which is worse than the bound that we obtained from the crossing lemma.

Repeating the process of computing lower bounds for sets of small cardinality we checked all order types of size up to 11 [12, 26] (the largest cardinality for which a complete data base exists). We obtained $\overline{\text{cr}}_2(K_{10}) = 5$ and $\overline{\text{cr}}_2(K_{11}) = 10$. By Lemma 2.8, the latter gives the improved lower bound of $\overline{\text{cr}}_2 \geq 1/33$.

2.3.1. Straight-line versus general drawings

The best straight-line drawings of K_n with $n \leq 8$ have no monochromatic crossing; see again Figure 2.3 (left). In [159, Section 3] the authors state that no graph is known were the k-planar crossing number is strictly smaller than the rectilinear k-planar crossing number for any $k \geq 2$. Moreover, according to personal communication [188], the similar question whether a graph exists where the k-planar crossing number is strictly smaller than the rectilinear k-colored crossing number was open. We next argue that K_9 is such an example. From the previous section we know that $\overline{\text{cr}}_2(K_9) = 2$. Inspecting rotation systems for $n = 9$ [1] which have the minimum number of 36 crossings, we have been able to construct a drawing of K_9 which has only one monochromatic crossing; see Figure 2.3 (right). As the graph thickness of K_9 is 3 [50, 192], we cannot draw K_9 with just two colors without monochromatic crossings. Thus, we get the following result.

Observation 2.9. *The biplanar crossing number for K_9 is one and is thus strictly smaller than the rectilinear 2-colored crossing number $\overline{\text{cr}}_2(K_9) = 2$.*

2.4. Upper bounds on the ratio $\overline{\text{cr}}_2(D)/\overline{\text{cr}}(D)$

In this section we study the extreme values that $\overline{\text{cr}}_2(D)/\overline{\text{cr}}(D)$ can attain for straight-line drawings D of K_n.

2.4.1. General straight-line drawings of K_n

Using a simple probabilistic argument as in [159], 2-coloring the edges uniformly at random, it can be shown that $\overline{\text{cr}}_2(D)/\overline{\text{cr}}(D) < 1/2$ for every straight-line drawing D, even if the underlying graph is not K_n. For completeness we spell out this argument here.

Lemma 2.10. *For any straight-line drawing D of any graph G, the ratio $\overline{\text{cr}}_2(D)/\overline{\text{cr}}(D)$ is at most 1/2.*

Proof. Color each edge of D independently at random with equal probability red or blue. For every pair of edges of D that cross, the probability of them having the same color is exactly 1/2. Therefore, the expected value of the ratio between monochromatic crossings and crossings is equal to 1/2. Since there is a coloring (only using one color) for which the ratio is equal to 1, there exists a 2-edge-coloring χ of D such that $\overline{\text{cr}}_2(D, \chi)/\overline{\text{cr}}(D) < 1/2$, and the result follows. □

In the following, we show that for K_n this upper bound on $\overline{\text{cr}}_2(D)/\overline{\text{cr}}(D)$ can be improved. To obtain our improved bound, we find subdrawings of D and colorings such that many of the crossings in these drawings are between edges of different colors. To this end, we need to find large subsets of vertices of D with identical geometric properties. We use the following definition and theorem. Let $(Y_1, ..., Y_k)$ be a tuple of finite subsets of points in the plane. A *transversal* of $(Y_1, ..., Y_k)$ is a tuple of points $(y_1, \ldots, y_k)$ such that $y_i \in Y_i$ for all i.

Theorem 2.11 (Positive fraction Erdős-Szekeres theorem). *For every integer $k \geq 4$ there is a constant $c_k > 0$ such that every sufficiently large finite point set $X \subset \mathbb{R}^2$ in general position contains k disjoint subsets $Y_1, \ldots, Y_k$, of at least $c_k|X|$ points each, such that each transversal of $(Y_1, \ldots, Y_k)$ is in convex position.*

The Positive Fraction Erdős-Szekeres theorem was proved by Bárány and Valtr [48] (see also Matoušek's book [145]). Although it is not stated in the theorem, every transversal of the $(Y_1, \ldots, Y_k)$ has the same (labelled) order type. Making use of that result we obtain the following theorem.

Theorem 2.12. *There exists an integer $n_0 > 0$ and a constant $c > 0$ such that for any straight-line drawing D of K_n on $n \geq n_0$ vertices, $\overline{\mathrm{cr}}_2(D)/\overline{\mathrm{cr}}(D) < \frac{1}{2} - c$.*

Proof. Let c_4 be as in Theorem 2.11 and let n_0 be such that Theorem 2.11 holds for $k = 4$ and for point sets with at least n_0 points. Let D be a straight-line drawing of K_n, where $n \geq n_0$.

Our general strategy is as follows. We first find subsets of edges of D that can be 2-colored such that many of the crossings between these edges are between pairs of edges of different colors. We remove these edges and search for a subset of edges with the same property. We repeat this process as long as possible. We 2-color the remaining edges so that at most half of the crossings are monochromatic. Afterwards, we put back the edges we removed while 2-coloring them in a convenient way.

We define a sequence of subsets $V = X_0 \supset X_1 \supset \cdots \supset X_m$ of vertices of D, where $V = X_0$ is the set of vertices of D, and tuples $(F_1, F_1'), \ldots, (F_m, F_m')$ of sets of edges of D as follows. Suppose that X_i has been defined. If $|X_i| < n_0$, we stop the process. Otherwise we apply Theorem 2.11 to X_i, to obtain a tuple (Y_1, Y_2, Y_3, Y_4) of disjoint subsets of points X_i, each with exactly $\lfloor c_4|X_i| \rfloor$ vertices, such that every transversal (y_1, y_2, y_3, y_4) of (Y_1, Y_2, Y_3, Y_4) is a convex quadrilateral. Without loss of generality we assume that (y_1, y_2, y_3, y_4) appear in clockwise order around this quadrilateral. This implies that the edge y_1y_3 crosses the edge y_2y_4. Let F_i be the set of edges with an endpoint in Y_1 and an endpoint in Y_3; let F_i' be the set of edges with an endpoint in Y_2 and an endpoint in Y_4; and finally, let $X_{i+1} = X_i \setminus (Y_1 \cup Y_2)$. Note that every edge in F_i crosses every edge in F_i'.

We now consider the remaining edges. Let $\overline{F}$ be the set of edges of D that are not contained in any F_i nor in any F_i' for $1 \leq i \leq m$. Let H be the straight-line drawing with the same vertices as D and with edge set equal to $\overline{F}$. By a probabilistic argument 2-coloring the edges uniformly at random, there is a coloring χ' of the edges of H so that $\overline{\mathrm{cr}}(H)/\overline{\mathrm{cr}}_2(H, \chi') \geq 2$.

We now 2-color the edges in F_i and F_i'. We define a sequence of straight-line drawings $H = D_{m+1}, \subset D_m \subset \cdots \subset D_0 = D$ and a corresponding sequence of 2-edge-colorings $\chi' = \chi_{m+1}, \chi_m, \ldots, \chi_0 = \chi$ that satisfies the following. Each χ_i is a 2-edge-coloring of D_i. Also χ_{i-1} when restricted to D_i equals χ_i. Suppose that D_i and χ_i have been defined and that $0 < i \leq m+1$. Let D_{i-1} be the straight-line drawing with the same vertices as D and with edge set E_{i-1} equal to $E_i \cup F_{i-1} \cup F_{i'-1}$ (where E_i is the edge set of D_i). Since χ_{i-1} coincides with χ_i in the edges of E_i, we only need to specify the colors of F_{i-1} and F_{i-1}'. We color the edges of F_i with the same color and the edges of F_{i-1}' with the other color. There are two options for doing this, and one of them guarantees that at most half of the crossings between an edge of $F_{i-1} \cup F_{i-1}'$ and an edge of D_i are monochromatic. We choose this option to define χ_{i-1}.

In what follows we assume that D has been colored by χ. Let C be the set of pairs of edges of D that cross. Of these, let C_1 be the subset of pairs of edges such that both of them are contained in $F_i \cup F_i'$ for some $1 \le i \le m$. Let $C_2 := C \setminus C_1$. Note that, by construction of χ, at most half of the pairs of edges in C_2 are of edges of the same color. For a given i, let E_i' be the subset of pairs of edges in C_1 such that both edges are in $F_i \cup F_i'$. Let (Y_1, Y_2, Y_3, Y_4) be the tuple of disjoint subsets of points X_i used to define F_i and F_i'. Recall that each Y_i consists of $\lfloor c_4 |X_i| \rfloor$ points. Every pair of crossing edges defines a convex quadrilateral and, conversely, every convex quadrilateral defines a unique pair of crossing edges. Therefore, by construction there at most ${c_4}^4 \lfloor |X_i| \rfloor^4/2$ pairs of edges in E_i' such that both edges are of the same color; and there are exactly $\lfloor c_4 |X_i| \rfloor^4$ pairs of edges in E_i' such that the edges are of different color. Thus, at most $1/3$ of the pairs of edges in E_i' are edges of the same color.

Therefore,

$$\frac{\overline{\text{cr}}_2(D,\chi)}{\overline{\text{cr}}(D)} \le \frac{\frac{1}{2}|C_1| + \frac{1}{3}|C_2|}{|C_1| + |C_2|}.$$

This is maximized when C_1 is as large as possible. Since there in total at most $\binom{n}{4}$ pairs of edges that cross, we have $|C_1| \le \binom{n}{4} - |C_2|$. Thus,

$$\frac{\overline{\text{cr}}_2(D,\chi)}{\overline{\text{cr}}(D)} \le \frac{\frac{1}{2}\binom{n}{4} - \frac{1}{6}|C_2|}{\binom{n}{4}}.$$

We now obtain a lower bound for the size of C_2. Note that $|X_0| = n$ and $|X_i| \ge (1-4c_4)|X_{i-1}|$. This implies that $|X_i| \ge (1-4c_4)^i n$ and that $|E_i| \ge {c_4}^4 (1-4c_4)^{4i} n^4$. Therefore,

$$|C_2| = \sum_{i=1}^{m} |E_i| \ge \sum_{i=1}^{m} {c_4}^4 (1-4c_4)^{4i} n^4 = 24{c_4}^4 \left(\frac{1}{1-(1-4c_4)^4} - 1 - o(1) \right) \binom{n}{4},$$

which completes the proof. □

2.4.2. Special cases: convex position and the double chain

We now consider the ratio $\overline{\text{cr}}_2(D)/\overline{\text{cr}}(D)$ for particular families of drawings D of K_n.

If the vertices of a straight-line drawing D are in convex position then the drawing D is said to be *convex*. For a convex straight-line drawing D of K_n the problem of finding a 2-edge-coloring that minimizes $\overline{\text{cr}}_2(D)$ is equivalent to the problem of finding the 2-page crossing number of the complete graph K_n; see Figure 2.4. In [5], Ábrego et al. proved that the 2-page crossing number of K_n is equal to

$$\frac{1}{4}\left\lfloor \frac{n}{2} \right\rfloor \left\lfloor \frac{n-1}{2} \right\rfloor \left\lfloor \frac{n-2}{2} \right\rfloor \left\lfloor \frac{n-3}{2} \right\rfloor.$$

Note that for general graphs testing whether the 2-page crossing number is zero is NP-hard [67, Corollary 4.4 on p. 47], though computing the 2-page crossing number is fixed-parameter tractable with respect to the sum of the number of crossings and the treewidth [45].

Since the number of crossings in a convex straight-line drawing of K_n is $\binom{n}{4}$, we obtain the following theorem.

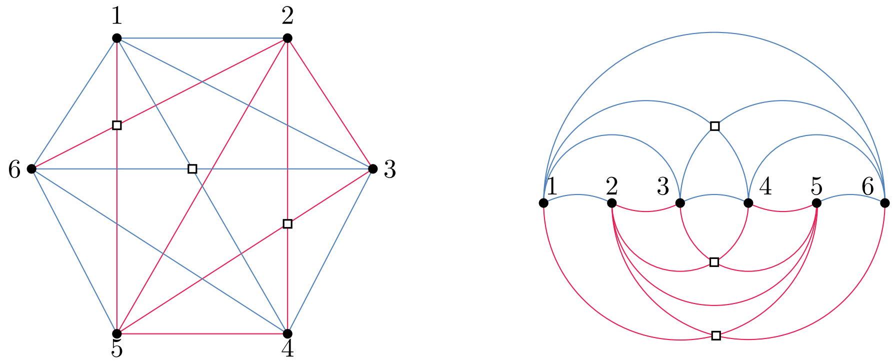

Figure 2.4.: Drawings of K_6. Left: Optimal 2-edge-coloring of a convex straight-line drawing. Right: Optimal 2-page book drawing.

Theorem 2.13. *If D is a convex straight-line drawing of K_n, then* $\overline{\mathrm{cr}}_2(D)/\overline{\mathrm{cr}}(D) = 3/8 - o(1)$.

Note that this theorem implies that the maximum value of the constant c such that Theorem 2.12 holds is at most $1/8$.

The other special case we consider consists of non-complete straight-line drawing whose vertices form a *double-chain*. This configuration is defined as follows. For $n \geq 3$, an (n,n)-*double-chain* consists of two (upper and lower) convex chains of n points each, linearly separable, and facing each other so that (i) two successive points of one chain and two successive points of the other are always in convex position, and (ii) three successive points of one chain and one point of the other are never in convex position.

Theorem 2.14. *Let D be a straight-line drawing of a graph whose vertex set is an (n,n)-double-chain, and in which there exists an edge between two vertices if and only if they belong to different chains. Then* $\overline{\mathrm{cr}}_2(D)/\overline{\mathrm{cr}}(D) \leq 1/3 + o(1)$.

Proof. We label the vertices of the upper chain from left to right as $1, \ldots, n$ and we label the vertices on lower chain from left to right also as $1, \ldots, n$. Let ij be an edge of D, with i in the upper chain and j in the lower chain. If $i < j$ then we color ij blue; if $i > j$ then we color ij red; and if $i = j$ then we color ij red or blue.

Let $I = (i, j, k, l)$ be a tuple of indices with $1 \leq i \leq j \leq k \leq l \leq n$, and at most two of them equal. Let S be a set of four vertices of D, whose labels are in $\{i, j, k, l\}$, and such that two vertices are in the upper chain and the other two are in the lower chain. Note that S defines a unique pair of edges of D that cross; and conversely, every pair of edges that cross has two vertices in the upper chain and the other two in the lower chain. There are six possible choices for S (for a given I) and each defines a different pair of crossing edges (except when at least two indices are the same). Of these six pairs of crossing edges, only two are between edges of the same color. Since the number of possible tuples (i, j, k, l) in which at most two indices are equal is $\binom{n}{4} + O(n^3)$, the result follows. $\square$

2.5. Chapter summary

In this chapter we have proved both lower and upper bounds on the geometric 2-colored crossing number for the complete graph K_n. To obtain this result, we showed that asymptotic bounds can be derived from optimal and near-optimal instances with few vertices. We obtained such instances using a combination of heuristics, integer programming, and the order type database [12, 26]. Moreover, for any fixed drawing of K_n, we improved the bound on the ratio between its geometric 2-colored crossing number and its rectilinear crossing number.

3. Minimal representations of order types by straight-line drawings

Part of the results in this chapter have been accepted for publication [17].

3.1. Introduction

Let $S, T \subset \mathbb{R}^2$ be two sets of n labeled points in general position (no three collinear). We say that S and T have *the same order type* if there is a bijection $\varphi : S \to T$ such that any triple $(p, q, r) \in S^3$ of three distinct points has the same orientation (clockwise or counterclockwise) as the image $(\varphi(p), \varphi(q), \varphi(r)) \in T^3$. The resulting equivalence relation on planar n-point sets has a finite number of equivalence classes, the *order types* [105]. Representatives of several distinct order types of five or six points are illustrated in Figure 3.2. Among other things, the order type of the set of vertices determines which edges cross in a straight-line drawing. Thus, order types appear ubiquitously in the study of extremal problems on straight-line drawings.

Now, suppose we have discovered an interesting order type, and we would like to illustrate it in a publication. One solution is to give explicit coordinates of a representative point set S; see Figure 3.1 (left). This is unlikely to satisfy most readers. We could also present S as a set of dots in a figure. For some point sets (particularly those with extremal properties), the reader may find it difficult to discern the orientation of an almost collinear point triple. To mend this, we could draw all lines spanned by two points in S. In fact, it suffices to present only the *segments* between the point pairs (the straight-line drawing of the complete graph on S). The orientation of a triple can then be obtained by inspecting the corresponding triangle; see Figure 3.1 (middle). However, such a drawing is rather dense, and we may have trouble following an edge from one endpoint to the other. Therefore, we want to reduce the number of edges in the drawing as much as possible, but so that the order type remains uniquely identifiable; see Figure 3.1 (right).

```
(-1,1)
(1,1)
(-1,-1)
(1,-1)
(-0.6,0.4)
(-0.6,-0.4)
```

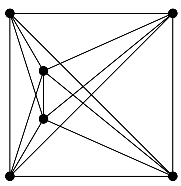

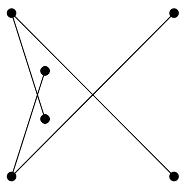

Figure 3.1.: Three different representations of an order type of six points.

We remark that there are order types requiring doubly exponential grid size in any drawing [106]. This implies that the coordinate representation of an order type might require exponential storage. However, order types of random point sets with high probability have representatives that can be drawn on an integer grid of small size [87].

We introduce the concept of *exit edges* to capture which edges are sufficient to uniquely describe a given order type in a robust way under continuous motion of vertices. More precisely, in a straight-line drawing on a representative point set with all exit edges, at least one vertex needs to move across an (exit) edge in order to change the order type. We give an alternative characterization of exit edges in terms of the dual line arrangement, where an exit edge corresponds to one or two empty triangular cells. This allows us to efficiently compute the set of exit edges for a given set of n points in $O(n^2)$ time and space.

Using the more general framework of abstract order types and their dual pseudoline arrangements, we prove that every set of $n \geq 4$ points has at least $(3n-7)/5$ exit edges. We also describe a family of n points with $n-3$ exit edges, showing that the best possible lower bound is of order $\Omega(n)$. An upper bound of $n(n-1)/3$ follows from known results on the number of triangular cells in line arrangements [108]. Thus, compared to the drawing all $n(n-1)/2$ edges, using only exit edges we save at least one third of the edges.

Two straight-line drawings (of the same graph) are *isomorphic* if there is an orientation-preserving homeomorphism of the plane transforming one into the other. Each class of this equivalence relation may be described combinatorially by the cyclic orders of the edge segments around vertices and crossings, and by the incidences of vertices, crossings, edge segments, and cells. In the following, we will consider topology-preserving deformations. An *ambient isotopy* of the Euclidean plane is a continuous map $f : \mathbb{R}^2 \times [0,1] \to \mathbb{R}^2$ such that $f(\cdot, t)$ is a homeomorphism for every $t \in [0,1]$ and $f(\cdot, 0) = \mathrm{Id}$. Note that if there is an ambient isotopy transforming a straight-line drawing $D(G)$ of a graph G into another straight-line drawing $D'(H)$ of a graph H, then $D(G)$ and $D'(H)$ are isomorphic.

Definition 3.1. *Let $D(G)$ be a straight-line drawing of a graph G on a point set S. We say that $D(G)$ is* supporting *for S if every ambient isotopy f of $\mathbb{R}^2$ that keeps the images of the edges of G straight (thus, transforming $D(G)$ into another straight-line drawing) and that allows at most one triple of collinear points of $f(S,t)$ for every $t \in [0,1]$, also preserves the order type of the vertex set.*

The connection between order types and straight-line drawings has been studied intensively, both for planar drawings and for drawings minimizing the number of crossings. For example, it is NP-complete to decide whether a planar graph can be embedded on a given point set [58]. Continuous movements of the vertices of plane straight-line drawings of graphs have also been considered [31]. The continuous movement of points maintaining the order type was considered by Mnëv [89, 148]. He showed that there are point sets with the same order type such that there is no ambient isotopy between them preserving the order type, settling a conjecture by Ringel [175]. The orientations of triples that have to be fixed to determine the order type are strongly related to the concept of *minimal reduced systems* [55].

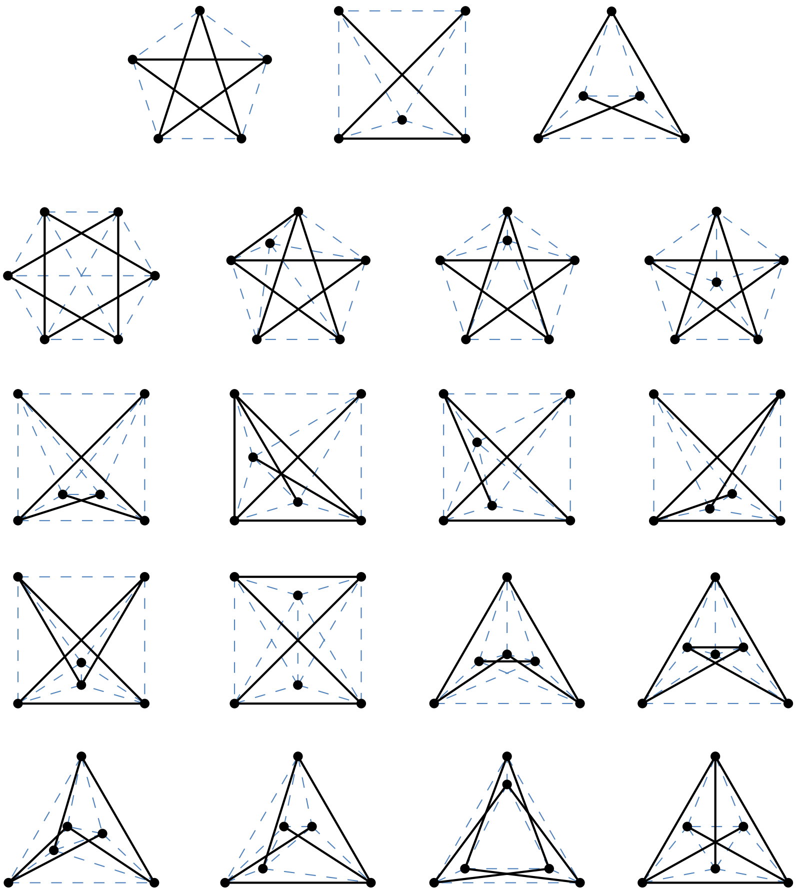

Figure 3.2.: Representatives of all the order types of five and six points in general position. Exit edges are drawn in black.

Outline. We introduce the concept of *exit edges* for a given point set. The resulting *exit drawings* are always supporting, though they are not necessarily minimal. In Section 3.2 we show that some exit edges are rendered unnecessary by non-stretchability of certain pseudoline arrangements. Despite being non-minimal in general, we argue that exit drawings are good candidates for supporting drawings by discussing their dual representation in pseudoline arrangements (Section 3.3). This connection allows us to both compute exit edges efficiently and give bounds on their number (Section 3.4). In particular we show that using exit edges to represent order types saves at least one third of the edges with respect to drawing all the edges. The undelying abstract graph of a supporting drawing does not need to be connected in general, and two minimal straight-line drawings that are supporting for point sets with different order types can be drawings of the same abstract graph; see Figure 3.2. Thus, the structure of the drawing is crucial. In Section 3.5 we provide some further properties of the exit drawings. We conjecture that graphs based on exit edges are not only supporting but also they encode the order type, as discussed in Section 3.6. In Section 3.7 we present a summary of the chapter.

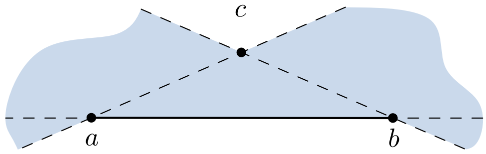

Figure 3.3.: If the blue region is empty of points, then the edge ab is an exit edge.

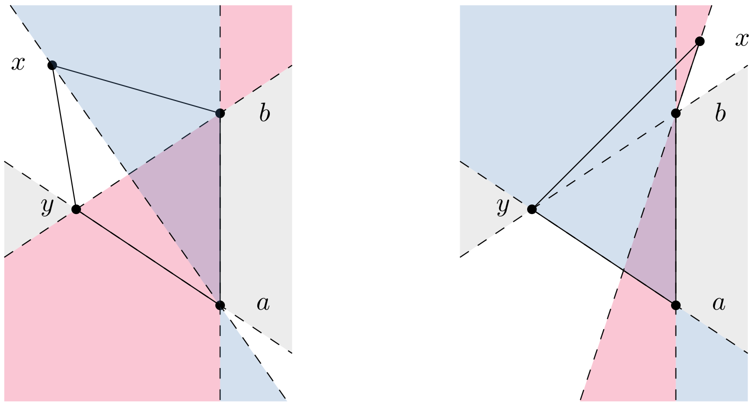

Figure 3.4.: An illustration of the proof of Proposition 3.3.

3.2. Exit edges

Clearly, a straight-line drawing of the complete graph K_n is supporting for every point set of n points in general position. To obtain a supporting drawing with fewer edges, we select edges so that no vertex of the resulting straight-line drawing can be moved to change the order type while preserving isomorphism.

Let a, b, c be three distinct points. We denote by $\overrightarrow{ab}$ the ray starting at a and going through b, and by $\overline{ab}$ the line through a and b directed from a to b.

Definition 3.2. *Let $S \subset \mathbb{R}^2$ be a finite set of points in general position. Let $a, b, c \in S$ be three distinct points. Then, ab is an* exit edge *with* witness *c if there is no $p \in S$ such that the line $\overline{ap}$ separates b from c or the line $\overline{bp}$ separates a from c. The straight line drawing with vertex set S and whose edges are the exit edges is called the* exit drawing *of S.*

Equivalently, ab is an *exit edge* with *witness* c if and only if the double-wedge through a between b and c and the double-wedge through b between a and c contain no point of S in their interior; see Figure 3.3 (left).

An exit edge has at most two witnesses. If $|S| \geq 4$ and ab is an exit edge in S with witness c, neither ac nor bc can be an exit edge with witness b or a, respectively. We illustrate the set of exit edges for sets of five and six points in Figure 3.2.

Exit edges can be characterized via 4-holes. For an integer $k \geq 3$, a *(general) k-hole* in S is a simple polygon $\mathcal{P}$ spanned by k points of S whose interior contains no point of S. If $\mathcal{P}$ is convex, we call $\mathcal{P}$ a *convex k-hole*. A point $a \in S$ or an edge ab with $a, b \in S$ is *extremal* for S if it lies on the boundary of the convex hull of S. A point or an edge in S that is not extremal in S is *internal* to S.

Proposition 3.3. *Let $S \subset \mathbb{R}^2$ be a set of points in general position and let $a, b \in S$. Then, ab is not an exit edge of S if and only if the following conditions hold:*

1. *If ab is extremal in S, then ab is an edge of at least one convex 4-hole in S.*
2. *If ab is internal in S, then there are two 4-holes $abxy$ and $bauv$, in counterclockwise order, such that their reflex angles (if any) are incident to ab.*

We remark that an internal exit edge either has a witness on both sides or is incident to at least one general 4-hole on one side.

Proof. Let ab be an exit edge with a witness c that lies, without loss of generality, to the left of $\overrightarrow{ab}$. Suppose there is a general 4-hole $abxy$, traced counterclockwise, such that the reflex angle of $abxy$ (if it exists) is incident to ab. We can assume that y lies to the left of $\overrightarrow{ab}$, as in Figure 3.4. First, suppose that $abxy$ is convex (this must hold if ab is extremal). Since ab is an exit edge with witness c, the line $\overline{ax}$ does not separate c from b and the line $\overline{by}$ does not separate c from a. Thus, c must be inside the 4-hole $abxy$, which is impossible. Second, suppose that $abxy$ is not convex (then, ab is internal), and x is to the right of $\overline{ab}$. Since ab is an exit edge with witness c, the line $\overline{bx}$ does not separate a from c and the line $\overline{ay}$ does not separate b from c, so c lies inside the 4-hole $abxy$, again a contradiction.

Conversely, assume that ab is not an exit edge. First, let ab be extremal, and let p be the closest point in $S \setminus \{a, b\}$ to the line $\overline{ab}$. The triangle abp is a 3-hole in S. Since p is not a witness for ab, there is a point $q \in S \setminus \{a, b, p\}$ such that, without loss of generality, the line $\overline{bq}$ separates a from p. Since ab is extremal, q lies on the same side of $\overline{ab}$ as p and, in particular, the polygon $abpq$ is convex. If we choose q so that it is the closest such point to the line $\overline{ap}$, the triangles bpq and abq are 3-holes in S. Altogether, we obtain a convex 4-hole $abpq$ in S.

Second, let ab be internal. Let p be closest in $S \setminus \{a, b\}$ to the line $\overline{ab}$ such that p lies to the left of $\overline{ab}$. The triangle abp is a 3-hole in S. Since p is not a witness for ab, there is a point $q \in S \setminus \{a, b, p\}$ such that either the line $\overline{bq}$ separates a from p or the line $\overline{aq}$ separates b from p. If q lies to the left of $\overline{ab}$, we obtain a convex 4-hole as in the previous case. Thus, we can assume that all such points q lie to the right of $\overline{ab}$. We choose the point q so that it is (one of the) closest to the line $\overline{ab}$ among all points that prevent ab from being an exit edge with witness p. Without loss of generality, we assume that the line $\overline{bq}$ separates a from p. The choice of q guarantees that bpq is a 3-hole in S. Thus, $abqp$ is a 4-hole in S incident to ab from the left. An analogous argument with a point p' from $S \setminus \{a, b\}$ that is closest to $\overline{ab}$ such that p' lies to the right of $\overline{ab}$ shows that there is an appropriate 4-hole in S incident to ab from the right. □

Proposition 3.4. *Let $S \subset \mathbb{R}^2$ be finite and in general position and, for every $t \in [0, 1]$, let $S(t)$ be a continuous deformation of S at time t. More formally, let $f : \mathbb{R}^2 \times [0, 1] \to \mathbb{R}^2$ be an ambient isotopy and $S(t) = \{f(s, t) \mid s \in S\}$, for $t \in [0, 1]$. Suppose that, for every $t \in [0, 1]$, there is at most one collinear triple of points in $S(t)$. Let (a, b, c) be the first triple to become*

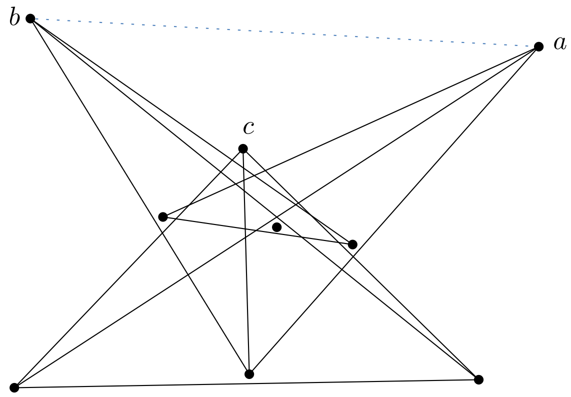

Figure 3.5.: Moving c overs ab to make (a, b, c) oriented clockwise without changing the orientation of other triples, would contradict Pappus's theorem [175].

collinear, at time $t_0 > 0$. If c lies on the segment ab in $S(t_0)$, then ab is an exit edge of $S(0)$ with witness c.

Proof. For $t \in [0, t_0)$, the triple orientations in $S(t)$ remain unchanged, and in $S(t_0)$ the point c lies on ab and the orientations of all triples except (a, b, c) are still unchanged. Thus, for $t \in [0, t_0)$, there is no line through two points of $S(t)$ that strictly separates the relative interior of ab from c. In particular, there is no such separating line through a or b in $S(0)$. Hence, ab is an exit edge with witness c. □

Corollary 3.5. *The exit drawing of every point set is supporting.*

The proof of Proposition 3.4 also shows that if a line separates c from the relative interior of ab, then there is such a line through a or b. This may suggest that the exit edges are necessary for a supporting drawing. However, this is not true in general. For example, in Figure 3.5, we see a construction by Ringel [175]: ab is an exit edge with witness c, but c cannot move over ab without violating Pappus' theorem. We note that in this situation, we might consider the *abstract order type* for the triple orientations we would obtain after moving c over ab. Since there is no planar point set with this set of triple orientations, this abstract order type is not *realizable*. Deciding realizability is (polynomial-time-)equivalent to the existential theory of the reals [148]. We will revisit these concepts in Section 3.4.

We note that there are point sets where two or more other exit edges prevent a witness c from crossing its corresponding exit edge ab; see, for example, Figure 3.6 (left). Since the two straight-line drawings in Figure 3.6 are not isomorphic, they cannot be transformed into each other by a continuous deformation as the one used in Definition 3.1. However, in this example, while c cannot move to ab without changing the order type in Figure 3.6 (left), if ab were not present, we could first change the point set to the one in Figure 3.6 (right) and then move c over ab. Thus, ab indeed has to be in a supporting drawing.

Figure 3.6.: The segment ab is an exit edge with witness c. In the left figure we cannot move c continuously to ab without first changing the order type, unless we also move other points.

3.3. Exit edges and empty triangular cells

The (real) *projective plane* $\mathbb{P}^2$ is a non-orientable surface obtained by augmenting the Euclidean plane $\mathbb{R}^2$ by a *line at infinity*. This line has one *point at infinity* for each direction, where all parallel lines with this direction intersect. Thus, in $\mathbb{P}^2$, each pair of parallel lines intersects in a unique point.

For a point set S in the Euclidean plane, add a line ℓ_∞ to obtain the projective plane. We use a duality transformation that maps a point s of $\mathbb{P}^2$ to a line s^* in $\mathbb{P}^2$. In this way, we get a set of lines S^* dual to S, giving a projective line arrangement $\mathcal{A}$. The removal of a line from $\mathcal{A}$ does not disconnect $\mathbb{P}^2$. Since $\mathbb{P}^2$ has non-orientable genus 1, removing any two lines ℓ_1 and ℓ_2 from $\mathbb{P}^2$ disconnects it into two components. We call the closure of each of the two components a *halfplane* determined by ℓ_1 and ℓ_2. The *marked cell* c_∞ is the cell of $\mathcal{A}$ that contains the point ℓ_∞^* dual to the line ℓ_∞. By appropriately choosing the duality transformation, we can assume that ℓ_∞^* lies at vertical infinity.

The combinatorial structure of $\mathcal{A}$, together with the marked cell, determines the order type of S. We show how to identify exit edges and their witnesses in dual line arrangements.

We use the marked cell c_∞ to orient the lines from S^*: First, we orient the lines on the boundary of c_∞ in one direction. Then, we iteratively remove lines that have already been oriented, and we define the orientation for the remaining lines from S^* by considering the new lines on the boundary of c_∞. Then, c_∞ is the only cell whose boundary is *oriented consistently*, that is, it can be traversed completely along the resulting orientation. In particular, for an unmarked triangular cell $\triangle$ in $\mathcal{A}$, the directed edges of $\triangle$ form a transitive order on its vertices, with a unique vertex of $\triangle$ in the middle (such that locally around that vertex the boundary is oriented consistently). We call this vertex the *exit vertex* of $\triangle$ and the line through the other two vertices of $\triangle$ the *witness line* of $\triangle$.

Note that if we consider the duality mapping a point $p = (p_x, p_y)$ from the real plane to the (non-vertical) line $p^* : y = p_x x - p_y$, then the described orientation procedure corresponds to orienting these dual lines from left to right.

Theorem 3.6. *Let $S \subset \mathbb{R}^2$ be in general position, and let $a, b, c \in S$. Then, ab is an exit edge with witness c if and only if the lines a^*, b^*, and c^* bound an unmarked triangular cell $\triangle$ in the arrangement $\mathcal{A}$ of lines from S^* so that c^* is the witness line of $\triangle$ and the point $\overline{ab}^* = a^* \cap b^*$ is the exit vertex of $\triangle$.*

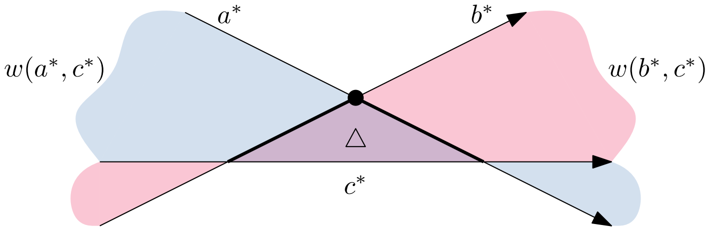

Figure 3.7.: An illustration of the proof of Theorem 3.6. If ab is an exit edge with witness c in S, then the two bold drawn segments of the corresponding triangular cell are unintersected, and thus, bound an unmarked triangular cell in S^*. The exit vertex is represented with a black disk.

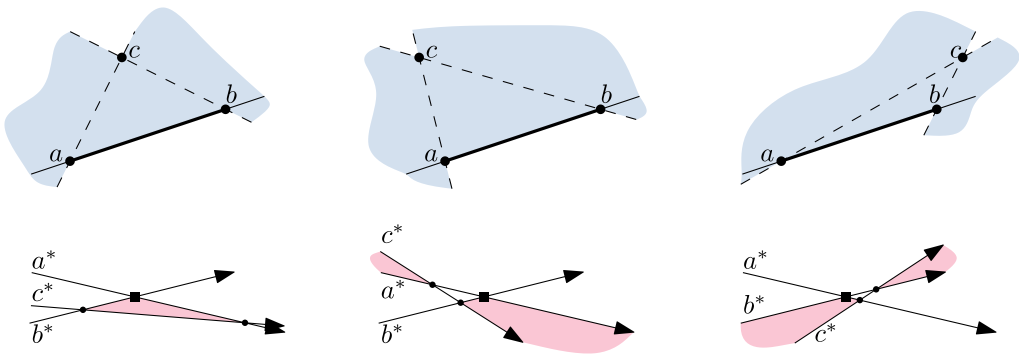

Figure 3.8.: When using the standard duality, the exit vertices might not be the ones with intermediate x-coordinate. Top: ab is an exit edge. The shaded area is empty of points. Bottom: Dual line arrangements. The red region is a triangular cell and the exit vertex is represented with a black disk.

Proof. For two points $p, q \in S$ and their dual lines $p^*, q^* \in S^*$, we denote by $w(p^*, q^*)$ the halfplane determined by p^* and q^* that does not contain the marked cell. Thus, the boundary of $w(p^*, q^*)$ is not oriented consistently. Since projective duality preserves incidences, the condition that no line spanned by two points of S intersects the edge pq is equivalent in S^* to $w(p^*, q^*)$ not containing any vertex of $\mathcal{A}$.

Let $\triangle$ be the triangular region determined by the intersection of the two halfplanes $w(a^*, c^*)$ and $w(b^*, c^*)$. By the projective duality, ab is an exit edge with witness c in S if and only if no line of S^* intersects a^* inside $w(b^*, c^*)$ or b^* inside $w(a^*, c^*)$. In other words, if and only if two sides of $\triangle$, lying on a^* and b^*, contain no intersection with lines from S^*. This is equivalent to $\triangle$ being a cell of the arrangement $\mathcal{A}$. Moreover, a^* and b^* share the exit vertex of $\triangle$; see Figure 3.7. Consequently, the exit vertex $a^* \cap b^*$ is the dual of the line containing the exit edge ab. $\square$

We remark that when considering the duality mapping a point $p = (p_x, p_y)$ from the real plane to the (non-vertical) line $p^* : y = p_x x - p_y$, the exit vertex of a triangular cell might not be the one with the intermediate x-coordinate; see Figure 3.8.

Figure 3.9.: Left: The two triangular cells $\triangle_1$ and $\triangle_2$ do not form an hourglass, because they share a vertex that is not an exit vertex. Right: The two triangular cells $\triangle_1$ and $\triangle_2$ form an hourglass because they share an exit vertex.

Corollary 3.7. *Let S be a set of n points in general position. Then the exit edges of S can be enumerated in $O(n^2)$ time by constructing the dual line arrangement of S and checking which cells are unmarked triangular cells.*

3.4. On the number of exit edges

Line arrangements can be generalized to so-called pseudoline arrangements. A *pseudoline* is a closed curve in the projective plane $\mathbb{P}^2$ whose removal does not disconnect $\mathbb{P}^2$. A set of pseudolines in $\mathbb{P}^2$, where any two pseudolines cross exactly once, determines a (projective) *pseudoline arrangement*. If no three pseudolines intersect in a common point, the pseudoline arrangement is *simple*. All notions that we have introduced for line arrangements, such as consistent orientations, exit vertices, or witness lines, naturally extend to pseudolines. For an introduction to pseudoline arrangements (and oriented matroids) we refer the reader to [89, 173]. (By the Topological Representation Theorem of Folkman and Lawrence [93], pseudoline arrangements are closely connected with oriented matrioids of rank 3.)

A pseudoline arrangement is *stretchable* if it is isomorphic to a line arrangement, that is, the corresponding cell complexes into which the two arrangements partition $\mathbb{P}^2$ are isomorphic. The combinatorial dual analogues of line arrangements and pseudoline arrangements are order types and abstract order types, respectively. Thus, deciding if a pseudoline arrangement is stretchable is (polynomial-time-)equivalent to the existential theory of the reals [89, 148].

As discussed in Section 3.3, the maximum number of *triangular cells* in a simple projective pseudoline arrangement gives an upper bound on the number of exit edges of a point set. However, one triangular cell could be c_∞, and there could be pairs of triangular cells with the same exit vertex. We call a configuration of the latter type an *hourglass*; see Figure 3.9. We say that the two pseudolines p and q that define the exit vertex of the two triangular cells of an hourglass H *slice* H and that H is *sliced* by p and by q.

Observation 3.8. *A triangular cell can be a part of at most one hourglass.*

Observation 3.9. *An exit edge ab with two witness points is dual to an hourglass with exit vertex $\overline{ab}^*$.*

Any projective arrangement of $n \geq 4$ lines has at least n triangular cells, as each line is incident to at least three triangular cells [141]. This is known to be tight. Therefore, taking into account the marked cell c_∞ and possible hourglasses, any set of $n \geq 4$ points has at least $\lceil \frac{n-1}{2} \rceil$ exit edges. We improve this lower bound by bounding from below the difference between the number of triangular cells and the number of hourglasses.

Proposition 3.10. *Any set of $n \geq 4$ points in the plane has at least $(3n-7)/5$ exit edges.*

For the proof of Proposition 3.10 we use the following two lemmas. The first is a theorem by Grünbaum [108, Theorem 3.7 on p. 50], and the second can be derived from the proof of that theorem.

Lemma 3.11 (Grünbaum [108])**.** *In a simple pseudoline arrangement L every pseudoline from L is incident to at least three triangular cells.*

Lemma 3.12 (Grünbaum [108])**.** *Let L be a simple arrangement of pseudolines, and let H be a closed halfplane determined by two pseudolines $\ell_1, \ell_2 \in L$. If two other pseudolines of L cross in the interior of H, then there is a triangular cell in H that is incident to ℓ_1 but not to ℓ_2.*

Proof of Proposition 3.10. Let L be a simple projective line arrangement of $n \geq 4$ pseudolines $\ell_1, \ell_2, \ldots, \ell_n$. For each pseudoline $\ell_i \in L$, let t_i be the number of triangular cells incident to ℓ_i and h_i the number of hourglasses sliced by ℓ_i. Set $x_i = t_i - h_i/2$. For each pseudoline $\ell_i \in L$, there are three possible cases.

Case 1: There is no hourglass sliced by ℓ_i. By Lemma 3.11, every pseudoline is incident to at least three triangular cells. Thus, we have $x_i = t_i \geq 3$.

Case 2: The pseudoline ℓ_i slices an hourglass together with some pseudoline ℓ_j and the interior of each of the two halfplanes determined by ℓ_i and ℓ_j contains at least one crossing of some other pair of pseudolines. By Lemma 3.12, ℓ_i is incident to the two triangular cells of the hourglass plus at least two other triangular cells, one in each closed halfplane. (We ignore here that a cell might be the marked one.) Thus, $t_i \geq 4$. Observation 3.8 implies $h_i \leq t_i/2$. Overall we get $x_i = t_i - h_i/2 \geq t_i - t_i/4 \geq (3/4) \cdot 4 = 3$.

Case 3: The pseudoline ℓ_i slices an hourglass together with some pseudoline ℓ_j, and one of the two closed halfplanes H_1 and H_2 determined by ℓ_i and ℓ_j contains no crossing of any other pair of pseudolines in its interior. Suppose the closed halfplane that contains no further crossing is H_1. Then, the hourglass sliced by ℓ_i and ℓ_j is in H_1, as the other two lines defining the hourglass do not cross in that halfplane; see Figure 3.10 (left). Since H_1 contains no crossing in its interior, it is divided by the other pseudolines into 4-gons and the two triangular cells of the hourglass. In particular, the marked cell is bounded by only four pseudolines, two of them being ℓ_i and ℓ_j; see Figure 3.10 (right). Thus, there can be at most four pseudolines for which Case 3 applies. Notice that in this case $h_i = 1$, since any other hourglass sliced by ℓ_i would have one triangular cell in each of the two halfplanes H_1 and H_2 and the two triangular cells in H_1 form the already-counted hourglass (and by Observation 3.8 they cannot be part of another hourglass). Thus, we can only guarantee that $x_i \geq 3 - 1/2 = 5/2$. However, as we showed, this case can happen at most for two pairs of pseudolines.

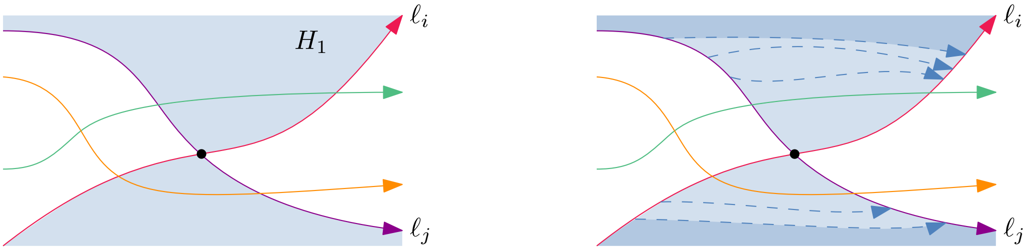

Figure 3.10.: In Case 3, both ℓ_i and ℓ_j must bound the marked cell, shown striped on the right picture. Moreover, that cell is bounded by four pseudolines. Thus, this case can happen for at most two pairs of pseudolines.

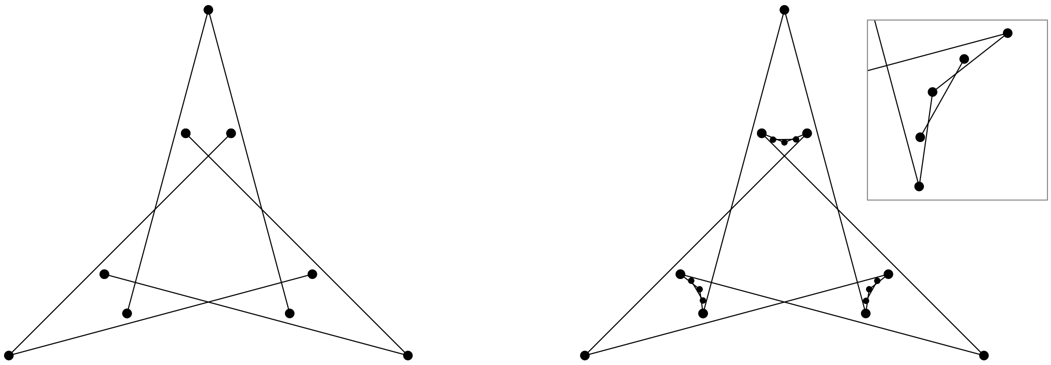

Figure 3.11.: Construction with $n-3$ exit edges.

Let T be the total number of triangular cells in L and let H be the total number of hourglasses. Summing the contributions of Cases 1–3, we have

$$3T - H = \sum_{i=1}^{n} t_i - \frac{1}{2}\sum_{i=1}^{n} h_i = \sum_{i=1}^{n} x_i \geq 3 \cdot (n-4) + 4 \cdot \binom{5}{2} = 3n - 2.$$

By Observation 3.8, we have $T \geq 2H$. Combining these inequalities, we get

$$T - H = \frac{3T - H + 2(T - 2H)}{5} \geq \frac{3T - H}{5} \geq \frac{3n-2}{5}.$$

By Theorem 3.6, the number of exit edges in a point set is equal to the number of exit vertices in its dual line arrangement. In general, the number of exit vertices in a pseudoline arrangement is bounded from below by $T - H - 1$. Therefore, there are at least $\frac{3}{5}n - \frac{7}{5}$ exit edges. $\square$

We do not know if the lower bound in Proposition 3.10 is tight. The smallest number of exit edges we could achieve is $n-3$ for $n \geq 9$; see Figure 3.11.

The number of triangular cells in a simple arrangement of n lines in the projective plane $\mathbb{P}^2$ is at most $n(n-1)/3$ [108, Theorem 2.21 on p. 26], so there are at most $n^2/3 + O(n)$ exit edges. This means that representing an order type with the exit drawing instead of with all the edges between pairs of points saves at least one third of the edges. Palásti and Füredi [98]

Number of points n	3	4	5	6	7	8	9	10
Min. EE	3	2	3	4	5	6	6	7
OTs with min. EE	1	1	1	3	13	80	1	3
Max. EE	3	3	5	6	9	14	16	21
OTs with max. EE	1	1	1	5	4	1	11	11
OTs	1	2	3	16	135	3 315	158 817	14 309 547

Table 3.1.: Minimum and maximum number of exit edges (EE) in a point set with at most 10 points and number of order types (OTs) achieving these minimum and maximum values.

Figure 3.12.: An illustration of Lemma 3.13.

showed that for every value of n there are simple arrangement of n lines in $\mathbb{P}^2$ with $n(n-3)/3$ triangular cells. Moreover, Roudneff [177] and Harborth [116] showed that the upper bound $n(n-1)/3$ is tight for infinitely many values of n (see also [53]). The point sets that are dual to the currently known triangle-maximizing arrangements have $n^2/6 + O(n)$ exit edges, since most of their exit edges have two witnesses. This gives a quadratic lower bound for the number of exit edges in the worst case, but the leading coefficient remains unknown.

It is worth noting that there are line arrangements with no pair of adjacent triangular cells [142], which implies the existence of point sets where every exit edge has precisely one witness.

Regarding precise values for small point sets, Table 3.1 shows the minimum and maximum number of exit edges in point sets with up to ten points. The values were obtained by exhaustive computer search in the database of order types [12, 26]. In particular, it shows that the example in Figure 3.11 displays the only order type of nine vertices with six exit edges.

The next lemma shows that the number of exit edges is non-decreasing when adding new points to a point set with at least four points.

Lemma 3.13. *Let S be a set of at least five points in general position. For any point $p \in S$, S has at least the same amount of exit edges as $S \backslash \{p\}$.*

Proof. Let ab with $a, b \in S\backslash\{p\}$ be an exit edge of $S\backslash\{p\}$ with witness $c \in S\backslash\{p\}$. For the proof we consider the dual line arrangement S^* of S. By Theorem 3.6, the lines a^*, b^*, and c^* define an empty triangular cell in $S^*\backslash\{p^*\}$ with $a^* \cap b^*$ as exit vertex. If p^* intersects this triangular cell, it can either intersect both a^* and b^* or c^*. In the former case, a smaller triangular cell exists inside the original one and $a^* \cap b^*$ is an exit vertex; see Figure 3.12(a). In the latter case, we can assume without loss of generality that p^* intersects both a^* and c^*. Again in this case a smaller triangular cell exists inside the original one; see Figure 3.12(b). Since $S^* \backslash \{p^*\}$ is a line arrangement with at least four lines, no two triangular cells in it are separated by a common boundary segment. Thus, the exit vertex of the smaller triangular cell cannot be the exit vertex for any other triangular cell in S^*. Therefore, when inserting p^* in $S^* \backslash \{p^*\}$ for each exit vertex that gets destroyed at least new one gets created. $\square$

We now show that the expected number of exit edges is quadratic for uniformly distributed point sets in convex shapes.

Theorem 3.14. *Let $S = \{p_1, \ldots, p_n\}$ be a set of n points in the plane with $p_i = (i, y_i)$ for every $i = 1, \ldots, n$, where each y_i is chosen uniformly at random from the real interval $[1, n]$. Then the expected number of exit edges in S is $\Theta(n^2)$.*

The main idea of the proof of Theorem 3.14 is inspired by the proof of Theorem 2.3 from [46].

Proof. The upper bound $O(n^2)$ on the number of exit edges in S follows from the fact that the number of pairs of points from S is $\binom{n}{2}$. In the rest of the proof we establish the lower bound $\Omega(n^2)$.

First, note that all points of S lie in the rectangle $R = [1, n] \times [1, n]$. Assume for convenience that n is divisible by 5. In the following, we identify each point p_i with the number i, which is the x-coordinate of p_i. Let $A = \{1, \ldots, \frac{n}{5}\}$, $B = \{\frac{2n}{5} + 1, \ldots, \frac{3n}{5}\}$, and $C = \{\frac{4n}{5} + 1, \ldots, n\}$. (We can assume that n is multiple of 5.) Let a, b, and c be fixed integers with $a \in A$, $b \in B$, and $c \in C$. We now find a lower bound on the probability that $p_a p_c$ is an exit edge of S with witness p_b.

The probability that the point p_b has vertical distance at most 1 from the line segment $p_a p_c$ is at least $\frac{1}{n}$, because the points from $\{b\} \times \mathbb{R}$ lying at distance at most 1 from $p_a p_c$ form a vertical line segment of length 2, and at least one half of this line segment is contained in R.

In the following, we assume that p_b has distance at most 1 from $p_a p_c$. Consider a point p_d with $d \in \{a+1, \ldots, n\} \setminus \{b, c\}$. Since $a \in A$ and $b \in B$, we have $b - a \geq n/5$ and $d - a \leq n$. Since p_b has distance at most 1 from $p_a p_c$, the vertical side of the triangle T bounded by the vertical line $\{b\} \times \mathbb{R}$ and by the rays $\overrightarrow{p_a p_b}$ and $\overrightarrow{p_a p_c}$ has length at most 1; see Figure 3.13. Since the triangle T' bounded by these two rays and by the vertical line $\{d\} \times \mathbb{R}$ is similar to T, and since $d - a \leq 5(b - a)$, the vertical side of T' has length at most 5. Thus, the probability that p_d lies in the convex wedge spanned by the rays $\overrightarrow{p_a p_b}$ and $\overrightarrow{p_a p_c}$ is at most $5/n$. An analogous argument shows that the probability that a point p_d with $d \in \{1, \ldots, c-1\} \setminus \{a, b\}$ lies in the convex wedge spanned by the rays $\overrightarrow{p_c p_a}$ and $\overrightarrow{p_c p_b}$ is at most $5/n$. In total, the probability that $p_a p_c$ is an exit edge of the point set $\{p_a, p_b, p_c, p_d\}$ with witness p_b is at least $1 - 10/n$.

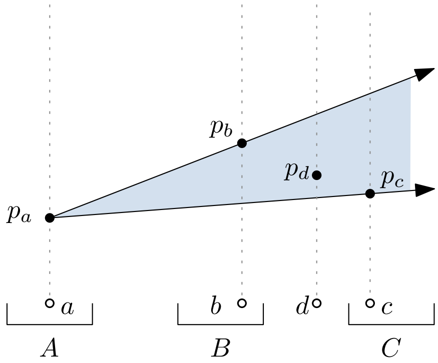

Figure 3.13.: An illustration of the proof of Theorem 3.14.

Altogether, the probability that $p_a p_c$ is an exit edge of S with witness p_b and that p_b is at vertical distance at most 1 from $p_a p_c$ is at least

$$\frac{1}{n} \cdot \prod_{d \in \{1,\dots,n\} \setminus \{a,b,c\}} \left(1 - \frac{10}{n}\right) = \frac{1}{n} \cdot \left(1 - \frac{10}{n}\right)^{n-3} \geq \frac{1}{n \cdot e^{20}},$$

where we use the inequality $1 - x \geq e^{-2x}$ for every real x with $0 \leq x \leq 1/2$.

Since every exit edge of S has at most two witnesses, the expected number of exit edges of S is at least

$$\frac{1}{2} \sum_{a \in A} \sum_{b \in B} \sum_{c \in C} \frac{1}{n \cdot e^{20}} \geq \Omega(n^2).$$

□

Combining the point-line duality that maps a point (a, b) to the line $\{(x, y) \in \mathbb{R}^2 \colon y = ax - b\}$ with Theorem 3.14, we obtain the following result.

Corollary 3.15. *Let $L = \{\ell_1, \dots, \ell_n\}$ be a set of lines, where $\ell_i = \{(x, y) \in \mathbb{R}^2 \colon y = i \cdot x - b_i\}$ and where b_i is chosen uniformly at random from the real interval $[1, n]$. Then the expected number of triangular cells in the line arrangement induced by L is $\Theta(n^2)$.*

3.5. Properties of exit drawings

We present some further results on supporting drawings and exit drawings.

Theorem 3.16. *Any straight-line drawing supporting a point set S, with $|S| \geq 9$, contains a crossing.*

Proof. Let $D(G)$ be a straight-line drawing with vertex set S without crossings. There is a point set S' with a different order type that also admits $D(G)$: Dujmović [81] showed that every plane graph admits a plane straight-line embedding with at least $\sqrt{n/2}$ points on a line.

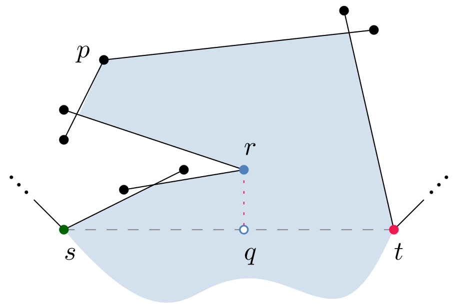

Figure 3.14.: An illustration of the setting for the proof of Proposition 3.17.

As we have a point set with a collinear triple that admits $D(G)$, there are at least two point sets S and S' in general position with a different order type that admit $D(G)$. Moreover, one can continuously morph S to S' while keeping the corresponding straight-line drawing plane and isomorphic to $D(G)$ (see, for example, [31]). Therefore, $D(G)$ does not support S. $\square$

Proposition 3.17. *Let S be a point set in general position in $\mathbb{R}^2$ and let $D(G)$ be its exit drawing. Every vertex on the boundary of the unbounded cell of $D(G)$ is extremal, that is, it lies on the boundary of the convex hull of S.*

Note that, as shown in Figure 3.5 (left), an analogous statement does not hold for general supporting drawings.

Proof. Suppose for contradiction that there is a point $p \in S$ on the boundary of the unbounded cell of the exit drawing of S and that is *non-extremal*, that is, lies in the interior of the convex hull $\mathrm{conv}(S)$ of S. This means that there is a polygonal path inside $\mathrm{conv}(S)$ from p to the boundary of $\mathrm{conv}(S)$ such that the interior of this path intersects no exit edge of S. Let $\delta(p)$ be the infimum of the lengths of such paths. Since $\mathrm{conv}(S)$ and S are both compact sets, there is a polygonal path P_p of length $\delta(p) > 0$ from p to the boundary of $\mathrm{conv}(S)$ that has no crossing with exit edges but may pass through other points of S. Among all such points p, let $r \in S$ be the point for which $\delta(r)$ is the minimum possible. Then P_r is a single segment. Let q be the endpoint of P_r on the boundary of $\mathrm{conv}(S)$.

If q coincides with an extremal point in S, we slightly perturb the point q so that q lies in the interior of an edge of $\mathrm{conv}(S)$ and the line segment rq does not intersect any exit edge of S. Let s and t be the endpoints of the edge of $\mathrm{conv}(S)$ containing q; see Figure 3.14 for an illustration.

Since exit edges are invariant to shearing, rotation, and mirroring, we assume without loss of generality that the following three conditions are satisfied.

(i) The points r and q lie on the y-axis, s has negative x-coordinate and t has positive x-coordinate,

(ii) the point r lies above the line $\overline{st}$, and

(iii) all points of S have distinct x-coordinates and S does not contain the origin.

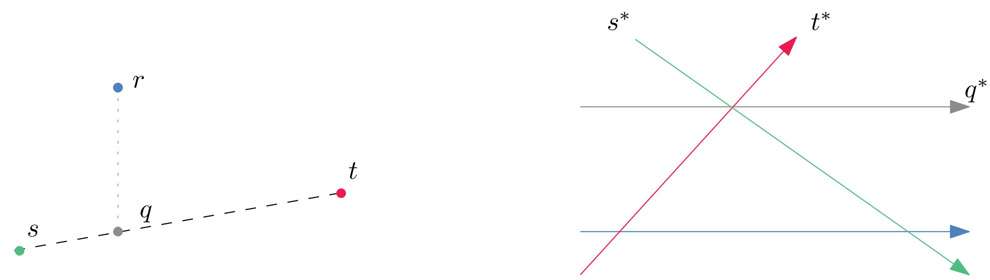

Figure 3.15.: Applying the dual transformation to the point set S (left) and obtaining the line arrangement S^* (right).

To obtain a contradiction, we will show that the segment rq intersects the interior of an exit edge of S. We will prove this in a dual setting.

By applying the duality transformation that maps each point $p = (a, b)$ to the line $p^* = \{(x, y); y = ax - b\}$, we map the point set S to the dual line arrangement S^*. Due to the three conditions above, the lines r^* and q^* are horizontal and the lines s^* and t^* have a negative and a positive slope, respectively; see Figure 3.15. A triple of points of S representing the endpoints of an exit edge together with its witness, such that the x-coordinate of the witness is between the x-coordinates of the endpoints of the exit edge, corresponds to a triangular cell Δ in S^* where the dual of the witness is the line bounding Δ with median slope.

Let $\triangle$ be the triangular region bounded by the lines r^*, s^*, and t^*. Since the line segment st is not an exit edge in S, the triangular region $\triangle$ is not a cell in S^*. Thus, the interior of $\triangle$ is intersected by some line from S^*. Since s and t are vertices of $\mathrm{conv}(S)$, their duals s^* and t^* are incident to the upper envelope of S^*.

Moving a point p vertically down from r to q corresponds to sweeping the dual S^* by a horizontal line p^* from r^* to q^*. Thus, meeting an exit edge of S with p corresponds to the situation in the dual in which the sweeping line p^* meets a vertex of a triangular cell of S^* such that the vertex is an intersection of a line with a positive slope and a line with a negative slope. Therefore, the line segment rq crosses an exit edge of S if and only if there is a triangular cell $\triangle'$ of S^* between r^* and q^* such that $\triangle'$ is bounded by two lines with positive and negative slope. To obtain a contradiction, we will show that $\triangle$ contains such a triangular cell $\triangle'$.

We start with the line arrangement containing the lines r^*, s^*, and t^* and we iteratively add the lines of S^* that intersect $\triangle$. We first add such lines with positive slopes, sorted by increasing slopes, and prove that in every step we have a triangular region inside $\triangle$ with one edge on s^*. In the beginning, $\triangle$ itself is such a triangular region. In each iterative step we have to show that if the existing triangular region is intersected, a new such triangular region is created. Note that s^* is the only line in our arrangement with a negative slope. Since we are iteratively inserting only lines with positive slopes sorted by increasing slopes, the slope of the newly inserted line is larger than the slopes of all remaining lines in our arrangement, possibly except t^*. The four possible cases are depicted in Figure 3.16.

Since we are inserting lines with positive slope, sorted by increasing slope, the newly inserted line cannot have minimum slope among the ones with positive slope. In fact, either t^* or the

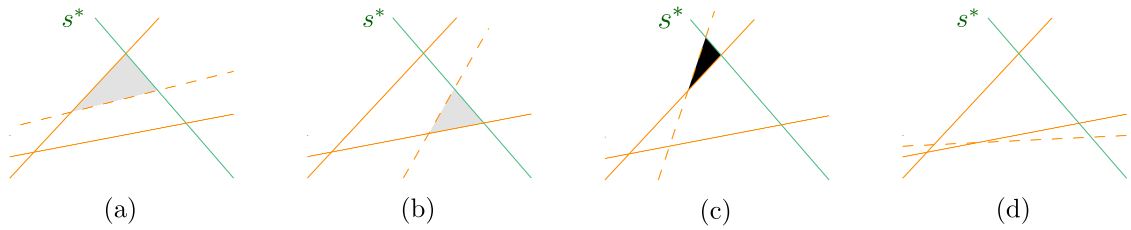

Figure 3.16.: All possible ways of intersecting the previous triangular region when inserting the lines from S^* with positive slopes. The line s^* is the only one with negative slope and is colored green. The newly inserted line with positive slope is depicted dashed. The leftmost case cannot occur in our setting, since the slope of r^* is zero and the lines are inserted by increasing slope. A newly created triangular cell is shaded and a new triangular region which contains a triangular cell is colored black.

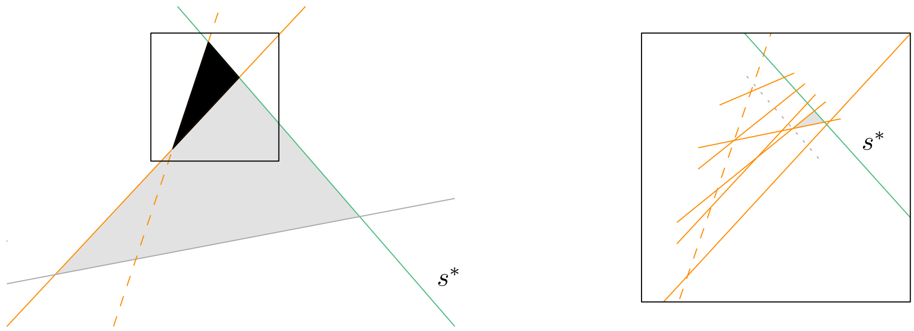

Figure 3.17.: Inside the black triangular region there is a triangular cell bounded by s^*.

newly inserted line has maximum slope. Therefore, the case shown in Figure 3.16(d) cannot occur in any step. For the cases shown in Figure 3.16(a) and 3.16(b), a new triangular cell (shaded gray) with an edge on s^* is created.

For the case shown in Figure 3.16(c), we show that inside the black triangular region there is a triangular cell inside $\triangle$ with an edge on s^*. Since the lines s^* and t^* must bound the upper envelope, the newly inserted line is below the intersection of s^* and t^*. Thus, the black triangular region is completely contained in $\triangle$. In particular, this means that the black triangular region cannot be bounded by t^*. If the black triangular region is a cell, we are done. Otherwise, to show that there is a triangular cell bounded by s^* inside the black triangular region, consider all the intersections of lines with positive slope added previously that lie in the closure of the black triangular region. The closest intersection to s^* is the leftmost vertex of a triangular cell bounded by s^*, since all the inserted lines have positive slope; see Figure 3.17 for an illustration. This shows that after inserting all lines with positive slope we have a triangular cell inside $\triangle$ with s^* on its boundary.

After all lines with positive slope have been inserted, in the resulting arrangement there are two lines $\tilde{r}^*$ and $\tilde{t}^*$ with non-negative slope (they possibly coincide with r^* or t^*), which, together with s^*, bound a triangular cell $\triangle'$ that lies inside the triangle $\triangle$.

We now iteratively add the lines with negative slope sorted by decreasing slope. In an analogous

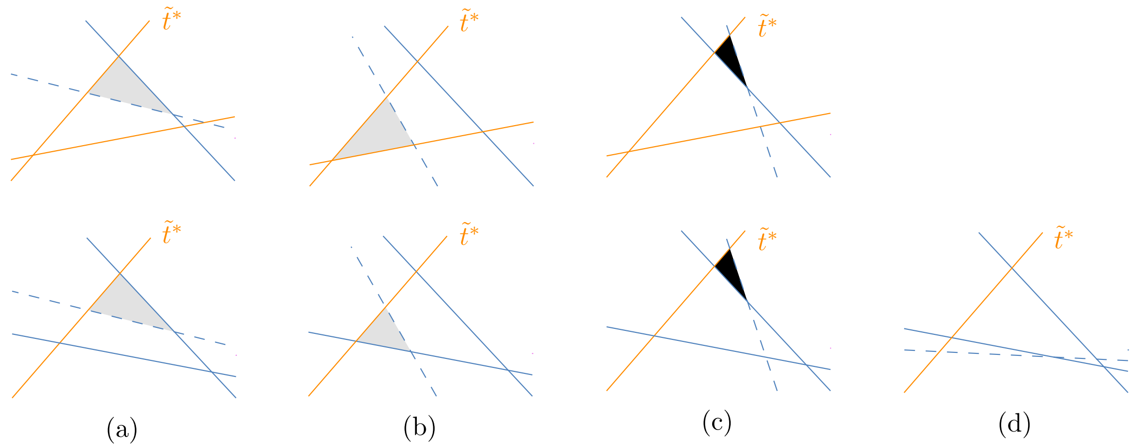

Figure 3.18.: All possible ways of intersecting the previous triangular region when inserting a new line with negative slope, whereas the rightmost case cannot occur in our setting. Lines with positive (negative) slope are colored orange (blue). The newly inserted line with negative slope is depicted dashed. A newly created triangular cell is shaded in gray and a new triangular region which contains a triangular cell is colored black.

manner as before, we show that in every step there is a triangular cell inside $\triangle'$ that is bounded by $\tilde{t}^*$ and at least one line with negative slope. All possible cases after adding a new line are depicted in Figure 3.18. Notice that the different cases are the mirrored versions of the ones in Figure 3.16. The symmetry allows us to use arguments analogous to the ones used before.

In the cases shown in Figures 3.18(a) and 3.18(b), a new triangular region (depicted shaded) bounded by $\tilde{t}^*$ and at least one line with negative slope is created. Since we are inserting lines with negative slope, sorted by decreasing slope, the newly inserted line cannot have maximum slope among the ones with negative slope. Thus, the case shown in Figure 3.18(d) cannot occur at any step.

For the case shown in Figure 3.18(c), we show that inside the black triangular region there is a triangular cell inside $\triangle$ with an edge on $\tilde{t}^*$. Since the lines s^* and t^* must bound the upper envelope, the newly inserted line is below the intersection of s^* and t^*. Thus, the black triangular region is contained in $\triangle$ and in $\triangle'$ (since this last one was bounded by s^*). In particular, this means that the black triangular region cannot be bounded by s^*. If the black triangular region is a cell, we are done. Otherwise, to show that there is a triangular cell bounded by $\tilde{t}^*$ inside the black triangular region, consider all the intersections of lines with negative slope previously added that lie in the closure of the black triangular region. The intersection closest to $\tilde{t}^*$ is the rightmost vertex of a triangular cell bounded by $\tilde{t}^*$, since the black triangular region is in $\triangle'$, and therefore all lines intersecting it have negative slope.

Altogether, by duality, we have that the segment rq crosses an exit edge of S, which is a contradiction. $\square$

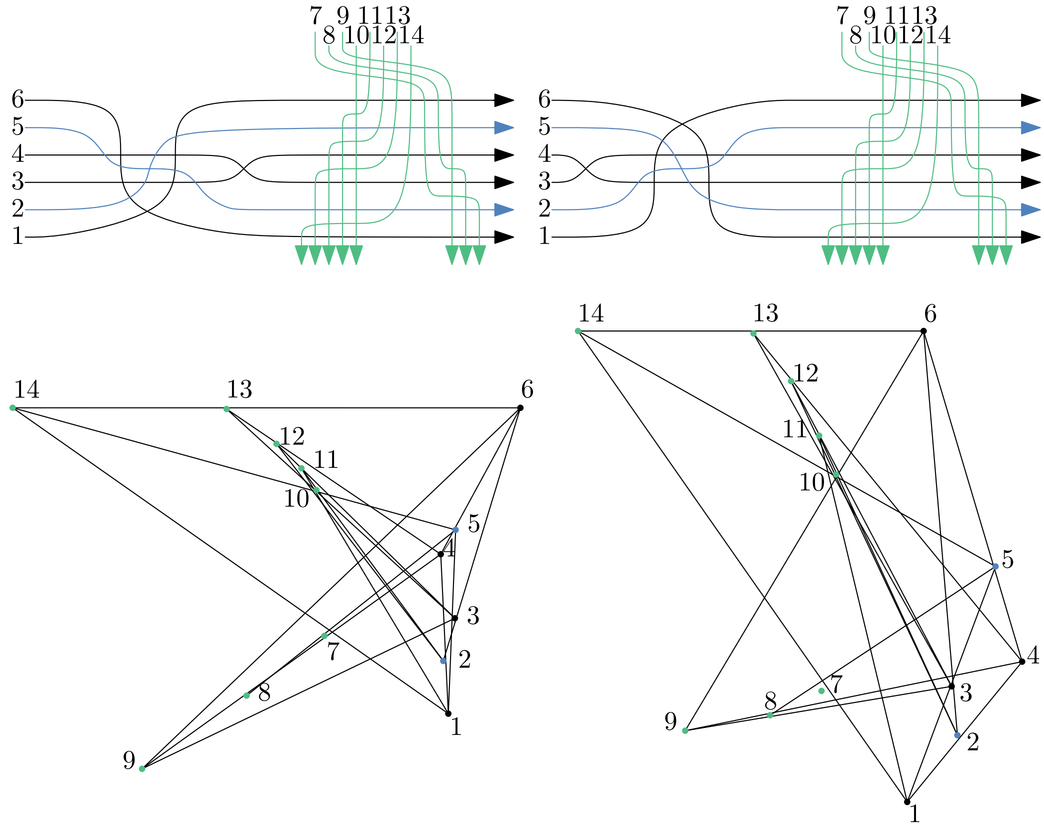

Figure 3.19.: Top: Two arrangements of 14 pseudolines with the same set of triangular cells (extending [90, Figure 3]). Bottom: Corresponding dual point sets. The order types are not the same (see for example the number of extremal points).

3.6. Reconstructing the order type

We conjecture that the straight-line drawing of the exit edges not only is supporting for S, but also that any point set S' that is the vertex set of an isomorphic straight-line drawing has the same order type as S. One might conjecture that already knowing all exit edges and their witnesses (in the dual line arrangement, all triangular cells and their boundary orientations) is sufficient to determine the order type. Surprisingly, this turns out to be not true. A counterexample is presented in Figure 3.19 as a dual (stretchable) pseudoline arrangement of 14 lines in the projective plane, based on an example by Felsner and Weil [90]. It consists of two arrangements of six lines in the Euclidean plane that are combinatorially different, but share the set of triangular cells and their boundary orientations. While the exit edges are the same for the two different order types, the corresponding exit drawings are not isomorphic.

In the dual of that example the order of the triangular cells along each pseudoline differs, but that extra information is not enough to distinguish the two order types: We can modify the pseudoline arrangements in Figure 3.19 by, essentially, duplicating pseudolines 1–6 and making a pseudoline and its duplication cross between the crossings with two green pseudolines (7–14). In

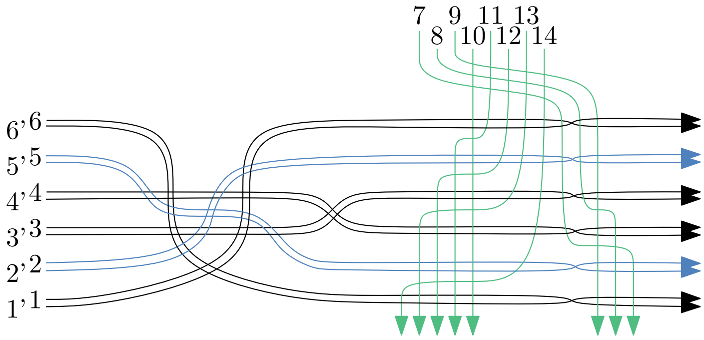

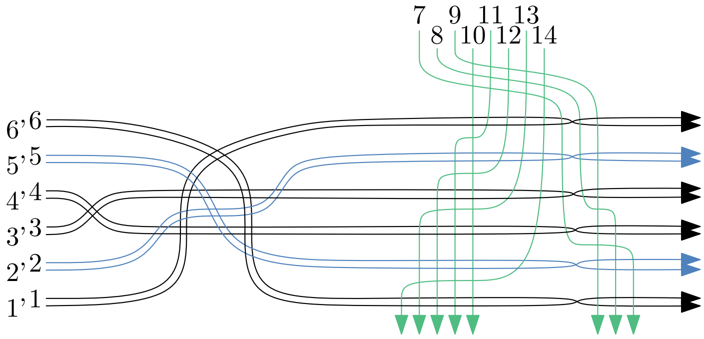

Figure 3.20.: Two arrangements of 20 pseudolines with the same set of triangular cells (extending [90, Figure 3]) and with the same order of the triangular cells along the pseudolines, but corresponding to different order types.

Figure 3.20 we show and illustration with two pseudoline arrangements with the same triangular cells (including their boundary orientations) and the same order of triangular cells along each pseudoline. However, the corresponding order types are not the same (see for example the number of extremal points). Note that the dual point sets of the pseudoline arrangements in Figure 3.20 can be obtained from the ones in Figure 3.19 by adding a copy of points 1–6 close to the original respective points. Thus, we cannot reconstruct the order type from that information.

3.7. Chapter summary

In this chapter we have introduced the concept of exit edges, that allows to have a compact visualization of the order type of a given point set S. The set of exit edges prevents the order type from changing under continuous motion of vertices. That is, in a drawing of a point set S with all its exit edges, in order to change the order type of S, at least one vertex needs to move across an exit edge. Moreover, exit edges have a natural dual characterization, which allows us to efficiently compute them and to bound their number. Due to the relevance of order types in a wide range of geometric problems, instances in which an order type must be displayed arise naturally. An example of the applicability of exit edges is presented in Figure 4.14 in the next chapter.

4. A superlinear lower bound on the number of 5-holes

The results presented in this chapter have been accepted for publication [15] and a preliminary version of them appeared in [14]. It is planned that these results also appear in the thesis of the coauthor Manfred Scheucher.

4.1. Introduction

The Erdős–Szekeres Theorem motivated a lot of further research, including numerous modifications and extensions of the theorem. In this chapter we study the number of empty convex pentagons in any set of n points in general position in the plane.

Let P be a finite set of points in general position in the plane. We say that a set H of k points from P is a *k-hole in P* if H is the vertex set of a convex k-gon containing no other points of P. For positive integers n and k, let $h_k(n)$ be the minimum number of k-holes in a set of n points in general position in the plane. For every n, Horton sets [122, 193, 195] are sets of n points in general position in the plane with no 7-hole. Thus, $h_k(n) = 0$ for every n and every $k \geq 7$. Asymptotically tight estimates for the functions $h_3(n)$ and $h_4(n)$ are known. The best known lower bounds are due to Aichholzer et al. [22] who showed that $h_3(n) \geq n^2 - \frac{32n}{7} + \frac{22}{7}$ and $h_4(n) \geq \frac{n^2}{2} - \frac{9n}{4} - o(n)$. The best known upper bounds $h_3(n) \leq 1.6196n^2 + o(n^2)$ and $h_4(n) \leq 1.9397n^2 + o(n^2)$ are due to Bárány and Valtr [49].

For $h_5(n)$ and $h_6(n)$, no matching bounds are known. So far, the best known asymptotic upper bounds on $h_5(n)$ and $h_6(n)$ were obtained by Bárány and Valtr [49] and give $h_5(n) \leq 1.0207n^2 + o(n^2)$ and $h_6(n) \leq 0.2006n^2 + o(n^2)$. For the lower bound on $h_6(n)$, Valtr [194] showed $h_6(n) \geq n/229 - 4$.

In this chapter we present a new lower bound on $h_5(n)$. It is widely conjectured that $h_5(n)$ grows quadratically in n, but to this date only lower bounds on $h_5(n)$ that are linear in n have been known. As noted by Bárány and Füredi [46], a linear lower bound of $\lfloor n/10 \rfloor$ follows directly from Harborth's result [115]. Bárány and Károlyi [47] improved this bound to $h_5(n) \geq n/6 - O(1)$. In 1987, Dehnhardt [75] showed $h_5(11) = 2$ and $h_5(12) = 3$, obtaining $h_5(n) \geq 3\lfloor n/12 \rfloor$. However, his result remained unknown to the scientific community until recently. García [99] then presented a proof of the lower bound $h_5(n) \geq 3\lfloor \frac{n-4}{8} \rfloor$ and a slightly better estimate $h_5(n) \geq \lceil 3/7(n-11) \rceil$ was shown by Aichholzer, Hackl, and Vogtenhuber [25]. Quite recently, Valtr [194] obtained $h_5(n) \geq n/2 - O(1)$. This was strengthened by Aichholzer et al. [22] to $h_5(n) \geq 3n/4 - o(n)$. All improvements on the multiplicative constant were achieved

n	≤ 9	10	11	12	13	14	15	16	17	18	19	20
$h_5(n)$	0	1	2	3	3	6	9	11	≤ 16	≤ 21	≤ 26	≤ 33

Table 4.1.: The minimum number $h_5(n)$ of 5-holes determined by any set of $n \leq 20$ points.

by utilizing the values of $h_5(10)$, $h_5(11)$, and $h_5(12)$. In the bachelor's thesis of Scheucher [181] the exact values $h_5(13) = 3$, $h_5(14) = 6$, and $h_5(15) = 9$ were determined with computer aid. Later, the value $h_5(16) = 11$ was also determined [180]. Table 4.1 summarizes our knowledge on the values of $h_5(n)$ for $n \leq 20$. The values $h_5(n)$ for $n \leq 16$ can be used to obtain further improvements on the multiplicative constant. By revising the proofs of [22, Lemma 1] and [22, Theorem 3], one can obtain $h_5(n) \geq n - 10$ and $h_5(n) \geq 3n/2 - o(n)$, respectively. We also note that it was shown in [170] that if $h_3(n) \geq (1 + \epsilon)n^2 - o(n^2)$, then $h_5(n) = \Omega(n^2)$.

As our main result, we give the first superlinear lower bound on $h_5(n)$. This solves an open problem, which was explicitely stated, for example, in a book by Brass, Moser, and Pach [56, Chapter 8.4, Problem 5] and in the survey [10].

Theorem 4.1. *There is an absolute constant $c > 0$ such that for every integer $n \geq 10$ we have $h_5(n) \geq cn \log^{4/5} n$.*

Let P be a finite set of points in the plane in general position and let ℓ be a line that contains no point of P. We say that P is *ℓ-divided* if there is at least one point of P in each of the two halfplanes determined by ℓ. For an ℓ-divided set P, we use $P = A \cup B$ to denote the fact that ℓ partitions P into the subsets A and B. In the rest of the chapter, we assume without loss of generality that ℓ is vertical and directed upwards, A is to the left of ℓ, and B is to the right of ℓ.

The following result, which might be of independent interest, is a crucial step in the proof of Theorem 4.1.

Theorem 4.2. *Let $P = A \cup B$ be an ℓ-divided set with $|A|, |B| \geq 5$ and with neither A nor B in convex position. Then there is an ℓ-divided 5-hole in P.*

The proof of Theorem 4.2 is computer-assisted. We reduce the result to several statements about point sets of size at most 11 and then verify each of these statements by an exhaustive computer search. To verify the computer-aided proofs we have implemented two independent programs, which, in addition, are based on different abstractions of point sets; see Subsection 4.5.2.

Using a result of García [99], we adapt the proof of Theorem 4.1 to provide improved lower bounds on the minimum numbers of 3-holes and 4-holes.

Theorem 4.3. *The following two bounds are satisfied for every positive integer n:*

(i) $h_3(n) \geq n^2 + \Omega(n \log^{2/3} n)$ *and*

(ii) $h_4(n) \geq \frac{n^2}{2} + \Omega(n \log^{3/4} n)$.

In the rest of the chapter, we assume that every point set P is planar, finite, and in general position. We also assume, without loss of generality, that all points in P have distinct x-coordinates. We use $\mathrm{conv}(P)$ to denote the convex hull of P and $\partial\,\mathrm{conv}(P)$ to denote the boundary of the convex hull of P.

A subset Q of P that satisfies $P \cap \mathrm{conv}(Q) = Q$ is called an *island of* P. Note that every k-hole in an island Q of P is also a k-hole in P. For any subset R of the plane, if R contains no point of P, then we say that R is *empty of points of* P.

Outline. In Section 4.2 we derive quite easily Theorem 4.1 from Theorem 4.2. Theorem 4.3 is proved in Section 4.3. Then, in Section 4.4, we give some preliminaries for the proof of Theorem 4.2, which is presented in Section 4.5. In Section 4.6, we give some final remarks. In particular, we show that the assumptions in Theorem 4.2 are necessary. To provide a better general view, we present a flow summary of the proof of Theorem 4.1 in Section 4.7. Finally, in Section 4.8 we summarize the chapter.

4.2. Proof of Theorem 4.1

We now apply Theorem 4.2 to obtain a superlinear lower bound on the number of 5-holes in a given set of n points. It clearly suffices to prove the statement for the case in which $n = 2^t$ for some integer $t \geq 5^5$.

We prove by induction on $t \geq 5^5$ that the number of 5-holes in an arbitrary set P of $n = 2^t$ points is at least $f(t) := c \cdot 2^t t^{4/5} = c \cdot n \log_2^{4/5} n$ for some absolute constant $c > 0$. For $t = 5^5$, we have $n > 10$ and, by the result of Harborth [115], there is at least one 5-hole in P. If c is sufficiently small, then $f(t) = c \cdot n \log_2^{4/5} n \leq 1$ and we have at least $f(t)$ 5-holes in P, which constitutes our base case.

For the inductive step we assume that $t > 5^5$. We first partition P with a line ℓ into two sets A and B of size $n/2$ each. Then we further partition A and B into smaller sets using the following well-known lemma, which is, for example, implied by a result of Steiger and Zhao [185, Theorem 1].

Lemma 4.4 ([185])**.** *Let $P' = A' \cup B'$ be an ℓ-divided set and let r be a positive integer such that $r \leq |A'|, |B'|$. Then there is a line that is disjoint from P' and that determines an open halfplane h with $|A' \cap h| = r = |B' \cap h|$.*

We set $r := \lfloor \log_2^{1/5} n \rfloor$, $s := \lfloor n/(2r) \rfloor$, and apply Lemma 4.4 iteratively in the following way to partition P into islands $P_1, \ldots, P_{s+1}$ of P so that the sizes of $P_i \cap A$ and $P_i \cap B$ are exactly r for every $i \in \{1, \ldots, s\}$. Let $P'_0 := P$. For every $i = 1, \ldots, s$, we consider a line that is disjoint from P'_{i-1} and that determines an open halfplane h with $|P'_{i-1} \cap A \cap h| = r = |P'_{i-1} \cap B \cap h|$. Such a line exists by Lemma 4.4 applied to the ℓ-divided set P'_{i-1}. We then set $P_i := P'_{i-1} \cap h$, $P'_i := P'_{i-1} \setminus P_i$, and continue with $i + 1$. Finally, we set $P_{s+1} := P'_s$.

Let $i \in \{1, \ldots, s\}$. If one of the sets $P_i \cap A$ and $P_i \cap B$ is in convex position, then there are at least $\binom{r}{5}$ 5-holes in P_i and, since P_i is an island of P, we have at least $\binom{r}{5}$ 5-holes in P. If this is the case for at least $s/2$ islands P_i, then, given that $s = \lfloor n/(2r) \rfloor$ and thus $s/2 \geq \lfloor n/(4r) \rfloor$, we obtain at least $\lfloor n/(4r) \rfloor \binom{r}{5} \geq c \cdot n \log_2^{4/5} n$ 5-holes in P for a sufficiently small absolute constant $c > 0$.

We thus further assume that for more than $s/2$ islands P_i, neither of the sets $P_i \cap A$ nor $P_i \cap B$ is in convex position. Since $r = \lfloor \log_2^{1/5} n \rfloor \geq 5$, Theorem 4.2 implies that there is an ℓ-divided 5-hole in each such P_i. Thus there is an ℓ-divided 5-hole in P_i for more than $s/2$ islands P_i. Since each P_i is an island of P and since $s = \lfloor n/(2r) \rfloor$, we have more than $s/2 \geq \lfloor n/(4r) \rfloor$ ℓ-divided 5-holes in P. As $|A| = |B| = n/2 = 2^{t-1}$, there are at least $f(t-1)$ 5-holes in A and at least $f(t-1)$ 5-holes in B by the inductive assumption. Since A and B are separated by the line ℓ, we have at least

$$2f(t-1) + n/(4r) = 2c(n/2) \log_2^{4/5}(n/2) + n/(4r) \geq cn(t-1)^{4/5} + n/(4t^{1/5})$$

5-holes in P. The right side of the above expression is at least $f(t) = cnt^{4/5}$, because the inequality $cn(t-1)^{4/5} + n/(4t^{1/5}) \geq cnt^{4/5}$ is equivalent to the inequality $(t-1)^{4/5}t^{1/5} + 1/(4c) \geq t$, which is true if c is sufficiently small, as $(t-1)^{4/5}t^{1/5} \geq t-1$. This finishes the proof of Theorem 4.1.

4.3. Proof of Theorem 4.3

In this section we improve the lower bounds on the minimum number of 3-holes and 4-holes. To this end we use the notion of generated holes as introduced by García [99].

Given a 5-hole H in a point set P, a 3-hole in P is *generated by* H if it is spanned by the leftmost point p of H and the two vertices of H that are not adjacent to p on the boundary of $\mathrm{conv}(H)$. Similarly, a 4-hole in P is *generated by* H if it is spanned by the vertices of H with the exception of one of the points adjacent to the leftmost point of H on the boundary of $\mathrm{conv}(H)$. We call a 3-hole or a 4-hole in P *generated* if it is generated by some 5-hole in P. We denote the number of generated 3-holes and generated 4-holes in P by $h_{3|5}(P)$ and $h_{4|5}(P)$, respectively. We also denote by $h_{3|5}(n)$ and $h_{4|5}(n)$ the minimum of $h_{3|5}(P)$ and $h_{4|5}(P)$, respectively, among all sets P of n points.

For an integer $k \geq 3$ and a point set P, let $h_k(P)$ be the number of k-holes in P. García [99] proved the following relationships between $h_3(P)$ and $h_{3|5}(P)$ and between $h_4(P)$ and $h_{4|5}(P)$.

Theorem 4.5 ([99]). *Let P be a set of n points and let $\gamma(P)$ be the number of extremal points of P. Then the following two equalities are satisfied:*

(i) $h_3(P) = n^2 - 5n + \gamma(P) + 4 + h_{3|5}(P)$ *and*

(ii) $h_4(P) = \frac{n^2}{2} - \frac{7n}{2} + \gamma(P) + 3 + h_{4|5}(P)$.

The proofs of both parts of Theorem 4.3 are carried out by induction on n similarly to the proof of Theorem 4.1. The base cases follow from the fact that each set P of $n \geq 10$ points contains at least one 5-hole in P and thus a generated 3-hole in P and a generated 4-hole in P. For the inductive step, let $P = A \cup B$ be an ℓ-divided set of n points with $|A|, |B| \geq \lfloor \frac{n}{2} \rfloor$, where n is a sufficiently large positive integer.

To show part (i), it suffices to prove $h_{3|5}(P) \geq \Omega(n \log^{2/3} n)$ as the statement then follows from Theorem 4.5. We use the recursive approach from the proof of Theorem 4.1, where we choose $r = \lfloor \log_2^{1/3} n \rfloor$. In each step of the recursion we either obtain $\lfloor \frac{n}{4r} \rfloor$ pairwise disjoint r-holes in P or $\lfloor \frac{n}{4r} \rfloor$ pairwise disjoint ℓ-divided 5-holes in P.

In the first case, each r-hole in P admits $\binom{r}{3}$ 3-holes in P and, by Theorem 4.5, it contains $\binom{r}{3} - r^2 + 5r - r - 4$ generated 3-holes in P. Thus, in total, we count at least $\frac{n}{4r}\binom{r}{3} - O(nr) \geq \Omega(n \log^{2/3} n)$ generated 3-holes in P.

In the second case, we have at least $\lfloor \frac{n}{4r} \rfloor$ ℓ-divided 5-holes in P. Without loss of generality, we can assume that at least $\frac{1}{2} \lfloor \frac{n}{4r} \rfloor \geq \lfloor \frac{n}{8r} \rfloor$ of those ℓ-divided 5-holes in P contain at least two points to the right of ℓ, as we otherwise continue with the horizontal reflection of P, which has ℓ as the axis of reflection. Therefore we have at least $\lfloor \frac{n}{8r} \rfloor$ ℓ-divided generated 3-holes in P and, analogously as in the proof of Theorem 4.1, we obtain

$$h_{3|5}(P) \geq 2h_{3|5}\left(\left\lfloor \frac{n}{2} \right\rfloor\right) + \left\lfloor \frac{n}{4r} \right\rfloor \geq \Omega(n \log^{2/3} n).$$

This finishes the proof of part (i).

The proof of part (ii) is almost identical. We choose $r = \lfloor \log_2^{1/4} n \rfloor$ and use the facts that every r-hole in P contains $\binom{r}{4} - \frac{r^2}{2} + \frac{7r}{2} - r - 3$ generated 4-holes in P and that every ℓ-divided 5-hole in P generates two 4-holes in P, at least one of which is ℓ-divided. This finishes the proof of Theorem 4.3.

4.4. Preliminaries for the proof of Theorem 4.2

Before proceeding with the proof of Theorem 4.2, we first introduce some notation and definitions, and state some immediate observations.

Let a, b, c be three distinct points in the plane. We denote by $\overrightarrow{ab}$ the ray starting at a and going through b, and by $\overline{ab}$ the line through a and b directed from a to b. We say c is to the *left* (*right*) of $\overline{ab}$ if the triple (a, b, c) traced in this order is oriented counterclockwise (clockwise). Note that c is to the left of $\overline{ab}$ if and only if c is to the right of $\overline{ba}$, and that the triples (a, b, c), (b, c, a), and (c, a, b) have the same orientation. We say a point set S is to the *left* (*right*) of $\overline{ab}$ if every point of S is to the left (right) of $\overline{ab}$.

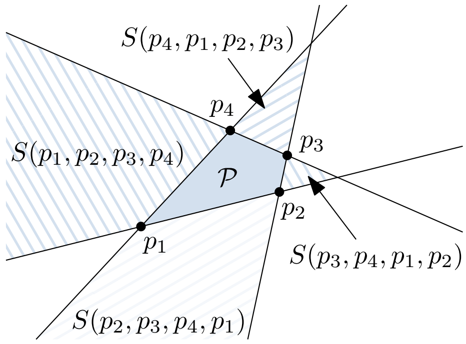

Figure 4.1.: An example of sectors.

Sectors of polygons. For an integer $k \geq 3$, let $\mathcal{P}$ be a convex polygon with vertices $p_1, p_2, \ldots, p_k$ traced counterclockwise in this order. We denote by $S(p_1, p_2, \ldots, p_k)$ the open convex region to the left of each of the three lines $\overline{p_1p_2}$, $\overline{p_1p_k}$, and $\overline{p_{k-1}p_k}$. We call the region $S(p_1, p_2, \ldots, p_k)$ a *sector* of $\mathcal{P}$. Note that every convex k-gon defines exactly k sectors. Figure 4.1 gives an illustration.

We use $\triangle(p_1, p_2, p_3)$ to denote the closed triangle with vertices p_1, p_2, p_3. We also use $\square(p_1, p_2, p_3, p_4)$ to denote the closed quadrilateral with vertices p_1, p_2, p_3, p_4 traced in the counterclockwise order along the boundary.

The following simple observation summarizes some properties of sectors of polygons.

Observation 4.6. *Let $P = A \cup B$ be an ℓ-divided set with no ℓ-divided 5-hole in P. Then the following conditions are satisfied.*

(i) Every sector of an ℓ-divided 4-hole in P is empty of points of P.

(ii) If S is a sector of a 4-hole in A and S is empty of points of A, then S is empty of points of B.

ℓ-critical sets and islands. An ℓ-divided set $C = A \cup B$ is called *ℓ-critical* if it fulfills the following two conditions.

(i) Neither A nor B is in convex position.

(ii) For every extremal point x of C, one of the sets $(C \setminus \{x\}) \cap A$ and $(C \setminus \{x\}) \cap B$ is in convex position.

Note that every ℓ-critical set $C = A \cup B$ contains at least four points in each of A and B. Figure 4.2 shows some examples of ℓ-critical sets. If $P = A \cup B$ is an ℓ-divided set with neither A nor B in convex position, then there exists an ℓ-critical island of P. This can be seen by iteratively removing extremal points so that none of the parts is in convex position after the removal.

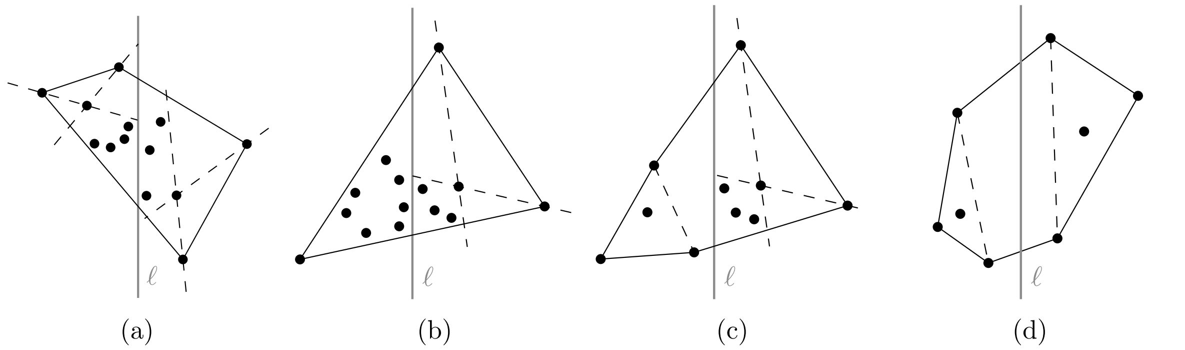

Figure 4.2.: Examples of ℓ-critical sets.

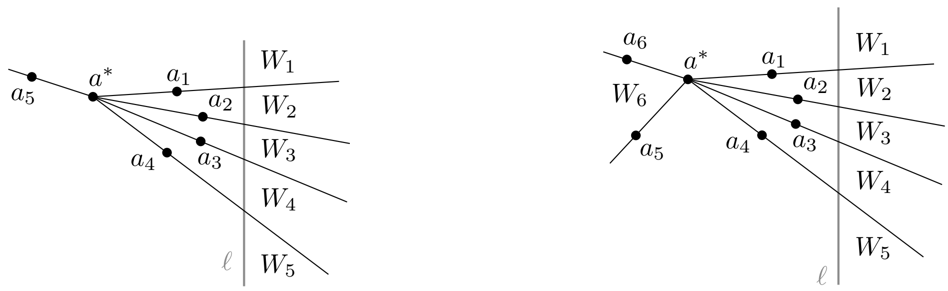

Figure 4.3.: Examples of a^*-wedges. Left: An example with $t = |A| - 1$. Right: An example with $t < |A| - 1$.

***a*-wedges and *a**-wedges.** Let $P = A \cup B$ be an ℓ-divided set. For a point a in A, the rays $\overrightarrow{aa'}$ for all $a' \in A \setminus \{a\}$ partition the plane into $|A| - 1$ regions. We call the closures of those regions *a-wedges* and label them as $W_1^{(a)}, \ldots, W_{|A|-1}^{(a)}$ in the clockwise order around a, where $W_1^{(a)}$ is the topmost a-wedge that intersects ℓ. Let $t^{(a)}$ be the number of a-wedges that intersect ℓ. Note that $W_1^{(a)}, \ldots, W_{t^{(a)}}^{(a)}$ are the a-wedges that intersect ℓ sorted in top-to-bottom order on ℓ. Also note that all a-wedges are convex if a is an inner point of A, and that there exists exactly one non-convex a-wedge otherwise. The indices of the a-wedges are considered modulo $|A| - 1$. In particular, $W_0^{(a)} = W_{|A|-1}^{(a)}$ and $W_{|A|}^{(a)} = W_1^{(a)}$.

If A is not in convex position, we denote by a^* the rightmost inner point of A and write $t := t^{(a^*)}$ and $W_k := W_k^{(a^*)}$ for $k = 1, \ldots, |A| - 1$. Recall that a^* is unique, since all points have distinct x-coordinates. Figure 4.3 gives an illustration.

We set $w_k := |B \cap W_k|$ and label the points of A so that W_k is bounded by the rays $\overrightarrow{a^* a_{k-1}}$ and $\overrightarrow{a^* a_k}$ for $k = 1, \ldots, |A| - 1$. Again, the indices are considered modulo $|A| - 1$. In particular, $a_0 = a_{|A|-1}$ and $a_{|A|} = a_1$.

Observation 4.7. *Let $P = A \cup B$ be an ℓ-divided set with A not in convex position. Then the points $a_1, \ldots, a_{t-1}$ lie to the right of a^* and the points $a_t, \ldots, a_{|A|-1}$ lie to the left of a^*.*

4.5. Proof of Theorem 4.2

First, we give a high-level overview of the main ideas of the proof of Theorem 4.2. We proceed by contradiction and we suppose that there is no ℓ-divided 5-hole in a given ℓ-divided set $P = A \cup B$ with $|A|, |B| \geq 5$ and with neither A nor B in convex position. If $|A|, |B| = 5$, then the statement follows from the result of Harborth [115]. Thus we assume that $|A| \geq 6$ or $|B| \geq 6$. We reduce P to an island Q of P by iteratively removing points from the convex hull until one of the two parts $Q \cap A$ and $Q \cap B$ contains exactly five points or Q is ℓ-critical with $|Q \cap A|, |Q \cap B| \geq 6$. If $|Q \cap A| = 5$ and $|Q \cap B| \geq 6$ or vice versa, then we reduce Q to an island of Q with eleven points and, using a computer-aided result (Lemma 4.14), we show that there is an ℓ-divided 5-hole in that island and hence in P. If Q is ℓ-critical with $|Q \cap A|, |Q \cap B| \geq 6$, then we show that $|A \cap \partial \operatorname{conv}(Q)|, |B \cap \partial \operatorname{conv}(Q)| \leq 2$ and that, if $|A \cap \partial \operatorname{conv}(Q)| = 2$, then a^* is the only interior point of $Q \cap A$ and similarly for B (Lemma 4.19). Without loss of generality, we assume that $|A \cap \partial \operatorname{conv}(Q)| = 2$ and thus a^* is the only interior point of $Q \cap A$. Using this assumption, we prove that $|Q \cap B| < |Q \cap A|$ (Proposition 4.21). By exchanging the roles of $Q \cap A$ and $Q \cap B$, we obtain $|Q \cap A| \leq |Q \cap B|$ (Proposition 4.22), which gives a contradiction.

To prove that $|Q \cap B| < |Q \cap A|$, we use three results about the sizes of the parameters $w_1, \dots, w_t$ for the ℓ-divided set Q, that is, about the numbers of points of $Q \cap B$ in the a^*-wedges $W_1, \dots, W_t$ of Q. We show that if we have $w_i = 2 = w_j$ for some $1 \leq i < j \leq t$, then $w_k = 0$ for some k with $i < k < j$ (Lemma 4.12). Further, for any three or four consecutive a^*-wedges whose union is convex and contains at least four points of $Q \cap B$, each of those a^*-wedges contains at most two such points (Lemma 4.18). Finally, we show that $w_1, \dots, w_t \leq 3$ (Lemma 4.20). The proofs of Lemmas 4.18 and 4.20 rely on some results about small ℓ-divided sets with computer-aided proofs (Lemmas 4.15, 4.16, and 4.17). Altogether, this is sufficient to show that $|Q \cap B| < |Q \cap A|$.

We now start the proof of Theorem 4.2 by showing that if there is an ℓ-divided 5-hole in the intersection of P with a union of consecutive a^*-wedges, then there is an ℓ-divided 5-hole in P.

Lemma 4.8. *Let $P = A \cup B$ be an ℓ-divided set with A not in convex position. For integers i, j with $1 \leq i \leq j \leq t$, let $W := \bigcup_{k=i}^{j} W_k$ and $Q := P \cap W$. If there is an ℓ-divided 5-hole in Q, then there is an ℓ-divided 5-hole in P.*

Proof. If W is convex then Q is an island of P and the statement immediately follows. Hence we assume that W is not convex. The region W is bounded by the rays $\overrightarrow{a^*a_{i-1}}$ and $\overrightarrow{a^*a_j}$ and all points of $P \setminus Q$ lie in the convex region $\mathbb{R}^2 \setminus W$; see Figure 4.4.

Since W is non-convex and every a^*-wedge contained in W intersects ℓ, at least one of the points a_{i-1} and a_j lies to the left of a^*. Moreover, the points $a_i, \dots, a_{j-1}$ are to the right of a^* by Observation 4.7. Without loss of generality, we assume that a_{i-1} is to the left of a^*.

If a_j is to the left of a^*, then we let h be the closed halfplane determined by the vertical line through a^* such that a_{i-1} and a_j lie in h. Otherwise, if a_j is to the right of a^*, then we let

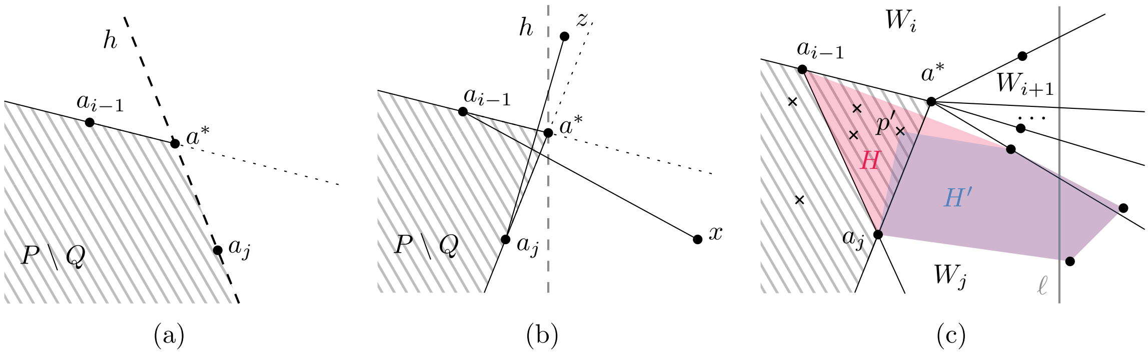

Figure 4.4.: Illustration of the proof of Lemma 4.8. (a) The point a_j is to the right of a^*. (b) The point a_j is to the left of a^*. (c) The hole H properly intersects the ray $\overrightarrow{a^*a_j}$. The convex hull of H is shaded in red and the convex hull of H' is shaded in blue.

h be the closed halfplane determined by the line $\overline{a^*a_j}$ such that a_{i-1} lies in h. In either case, $h \cap A \cap Q = \{a^*, a_{i-1}, a_j\}$.

Let H be an ℓ-divided 5-hole in Q. We say that H *properly intersects* a ray r if the interior of conv(H) intersects r. Now we show that if H properly intersects the ray $\overrightarrow{a^*a_j}$, then H contains a_{i-1}. Assume there are points $p, q \in H$ such that the interior of pq intersects $r := \overrightarrow{a^*a_j}$. Since r lies in h and neither of p and q lies in r, at least one of the points p and q lies in $h \setminus r$. Without loss of generality, we assume $p \in h \setminus r$. From $h \cap A \cap Q = \{a^*, a_{i-1}, a_j\}$ we have $p = a_{i-1}$. By symmetry, if H properly intersects the ray $\overrightarrow{a^*a_{i-1}}$, then H contains a_j.

Suppose for contradiction that H properly intersects both rays $\overrightarrow{a^*a_{i-1}}$ and $\overrightarrow{a^*a_j}$. Then H contains the points a_{i-1}, a_j, x, y, z for some points $x, y, z \in Q$, where $a_{i-1}x$ intersects $\overrightarrow{a^*a_j}$, and a_jz intersects $\overrightarrow{a^*a_{i-1}}$. Observe that z is to the left of $\overline{a_{i-1}a^*}$ and that x is to the right of $\overline{a_ja^*}$. If a_j lies to the right of a^*, then z is to the left of a^*, and thus z is in A; see Figure 4.4(a). However, this is impossible as z also lies in h. Hence, a_j lies to the left of a^*; see Figure 4.4(b). As x and z are both to the right of a^*, the point a^* is inside the convex quadrilateral $\square(a_{i-1}, a_j, x, z)$. This contradicts the assumption that H is a 5-hole in Q.

So assume that H properly intersects exactly one of the rays $\overrightarrow{a^*a_{i-1}}$ and $\overrightarrow{a^*a_j}$, say $\overrightarrow{a^*a_j}$; see Figure 4.4(c). In this case, H contains a_{i-1}. The interior of the triangle $\triangle(a^*, a_{i-1}, a_j)$ is empty of points of Q, since the triangle is contained in h. Moreover, conv(H) cannot intersect the line that determines h both strictly above and strictly below a^*. Thus, all remaining points of $H \setminus \{a_{i-1}\}$ lie to the right of $\overline{a_{i-1}a^*}$ and to the right of $\overline{a_ja^*}$. If H is empty of points of $P \setminus Q$, we are done. Otherwise, we let $H' := (H \setminus \{a_{i-1}\}) \cup \{p'\}$ where $p' \in P \setminus Q$ is a point inside $\triangle(a^*, a_{i-1}, a_j)$ closest to $\overline{a_ja^*}$. Note that the point p' might not be unique. By construction, H' is an ℓ-divided 5-hole in P. An analogous argument shows that there is an ℓ-divided 5-hole in P if H properly intersects $\overrightarrow{a^*a_{i-1}}$.

Finally, if H does not properly intersect any of the rays $\overrightarrow{a^*a_{i-1}}$ and $\overrightarrow{a^*a_j}$, then conv(H) contains no point of $P \setminus Q$ in its interior, and hence H is an ℓ-divided 5-hole in P. $\square$

4.5.1. Sequences of a^*-wedges with at most two points of B

In this subsection we consider an ℓ-divided set $P = A \cup B$ with A not in convex position. We consider the union W of consecutive a^*-wedges, each containing at most two points of B, and derive an upper bound on the number of points of B that lie in W if there is no ℓ-divided 5-hole in $P \cap W$; see Corollary 4.13.

Observation 4.9. *Let $P = A \cup B$ be an ℓ-divided set with A not in convex position. Let W_k be an a^*-wedge with $w_k \geq 1$ and $1 \leq k \leq t$ and let b be the leftmost point in $W_k \cap B$. Then the points a^*, a_{k-1}, b, and a_k form an ℓ-divided 4-hole in P.*

From Observation 4.6(i) and Observation 4.9 we obtain the following result.

Observation 4.10. *Let $P = A \cup B$ be an ℓ-divided set with A not in convex position and with no ℓ-divided 5-hole in P. Let W_k be an a^*-wedge with $w_k \geq 2$ and $1 \leq k \leq t$ and let b be the leftmost point in $W_k \cap B$. For every point b' in $(W_k \cap B) \setminus \{b\}$, the line $\overline{bb'}$ intersects the segment $a_{k-1}a_k$. Consequently, b is inside $\triangle(a_{k-1}, a_k, b')$, to the left of $\overline{a_kb'}$, and to the right of $\overline{a_{k-1}b'}$.*

The following lemma states that there is an ℓ-divided 5-hole in P if two consecutive a^*-wedges both contain exactly two points of B.

Lemma 4.11. *Let $P = A \cup B$ be an ℓ-divided set with A not in convex position and with $|A|, |B| \geq 5$. Let W_i and W_{i+1} be consecutive a^*-wedges with $w_i = 2 = w_{i+1}$ and $1 \leq i < t$. Then there is an ℓ-divided 5-hole in P.*

Proof. Suppose for contradiction that there is no ℓ-divided 5-hole in P. Let $W := W_i \cup W_{i+1}$ and let $Q := P \cap W$. By Lemma 4.8, there is also no ℓ-divided 5-hole in Q. We label the points in $B \cap W_i$ as b_{i-1} and b_i so that b_{i-1} is to the right of b_i. Similarly, we label the points in $B \cap W_{i+1}$ as b_{i+1} and b_{i+2} so that b_{i+2} is to the right of b_{i+1}. By Observation 4.10, the point a_i is to the right of $\overline{b_ib_{i-1}}$ and to the left of $\overline{b_{i+1}b_{i+2}}$. If the points $b_{i-1}, b_i, b_{i+1}, b_{i+2}$ are in convex position, then $a_i, b_{i+1}, b_{i+2}, b_{i-1}, b_i$ form an ℓ-divided 5-hole in P; see Figure 4.5 (left). Thus, we assume the points $b_{i-1}, b_i, b_{i+1}, b_{i+2}$ are not in convex position. Without loss of generality, we assume that $\overline{b_ib_{i-1}}$ intersects $\overline{b_{i+1}b_{i+2}}$.

We show that the segments a_ib_{i-1} and b_ib_{i+1} intersect. As $\overline{b_ib_{i-1}}$ intersects a_ia_{i-1} and $b_{i+1}b_{i+2}$, the point b_{i-1} lies in the triangle $\triangle(b_i, b_{i+1}, b_{i+2})$. Moreover, b_{i-1} is to the right of $\overline{b_{i+1}b_i}$, a_i is to the left of $\overline{b_{i+1}b_i}$, b_i is to the left of $\overline{a_ib_{i-1}}$, and b_{i+1} is to the right of $\overline{a_ib_{i-1}}$. Consequently, the points $a_i, b_{i+1}, b_{i-1}, b_i$ form an ℓ-divided 4-hole in P; see Figure 4.5 (right).

The points $a_{i-1}, b_i, b_{i-1}, b_{i+2}$ are in convex position because a_{i-1} is the leftmost and b_{i+2} is the rightmost of those four points and because both a_{i-1} and b_{i+2} lie to the left of $\overline{b_ib_{i-1}}$. Moreover, the points $a_{i-1}, b_i, b_{i-1}, b_{i+2}$ form an ℓ-divided 4-hole in P as $\square(a_{i-1}, b_i, b_{i-1}, b_{i+2})$ lies in W and $w_i = w_{i+1} = 2$.

We consider the four points $b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1}$. The point b_{i+2} is the rightmost of those four points. By Observation 4.10, b_{i+1} lies to the right of $\overline{a_ib_{i+2}}$ and a_{i+1} lies to the right of $\overline{b_{i+1}b_{i+2}}$.

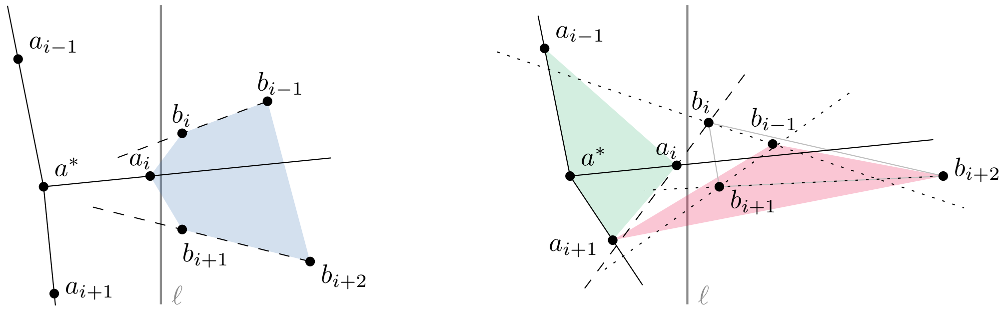

Figure 4.5.: Left: If $b_{i-1}, b_i, b_{i+1}, b_{i+2}$ are in convex position, then there is an ℓ-divided 5-hole in P. Right: The points $a^*, a_{i+1}, a_i, a_{i-1}$ form a 4-hole in P.

Since $b_{i-1} \in W_i$ and $b_{i+2} \in W_{i+1}$, the point b_{i-1} lies to the left of $\overline{a_i b_{i+2}}$. Thus, the clockwise order around b_{i+2} is $a_{i+1}, b_{i+1}, b_{i-1}$.

Suppose for contradiction that the points $b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1}$ form a convex quadrilateral. Due to the clockwise order around b_{i+2}, the convex quadrilateral is $\square(b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1})$. The only points of P that can lie in the interior of this quadrilateral are a^*, a_{i-1}, a_i, and b_i. Since the triangle $\triangle(b_{i+2}, b_{i+1}, a_{i+1})$ is contained in W_{i+1}, it contains neither of the points a^*, a_{i-1}, a_i, and b_i. Since the triangle $\triangle(b_{i+2}, b_{i-1}, b_{i+1})$ is contained in the convex hull of B, it does not contain a^*, a_{i-1}, nor a_i. Moreover, as b_{i-1} lies in the triangle $\triangle(b_i, b_{i+1}, b_{i+2})$, the triangle $\triangle(b_{i+2}, b_{i-1}, b_{i+1})$ also does not contain b_i. Thus the quadrilateral $\square(b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1})$ is empty of points of P. By Observation 4.6(i), the two sectors $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$ and $S(b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1})$ contain no point of P. Since every point of $B \setminus \{b_{i-1}, b_i, b_{i+1}, b_{i+2}\}$ is either in $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$ or in $S(b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1})$, we have $B = \{b_{i-1}, b_i, b_{i+1}, b_{i+2}\}$. This contradicts the assumption that $|B| \geq 5$.

Therefore the points $b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1}$ are not in convex position. In particular, the point b_{i+1} lies in the triangle $\triangle(b_{i-1}, a_{i+1}, b_{i+2})$, since a_{i+1} is the leftmost and b_{i+2} is the rightmost of the points $b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1}$ and since b_{i-1} lies in W_i; see the red area in Figure 4.5 (right) for an illustration.

Consequently, the point a_{i+1} lies to the left of $\overline{b_{i+1} b_{i-1}}$. By Observation 4.6(i), the point a_{i+1} is not in the sector $S(b_{i+1}, b_{i-1}, b_i, a_i)$, as otherwise the points $b_{i+1}, b_{i-1}, b_i, a_i, a_{i+1}$ form an ℓ-divided 5-hole in P. Thus the point a_{i+1} lies to the left of $\overline{a_i b_i}$; see Figure 4.5 (right).

The points $a^*, a_{i+1}, a_i, a_{i-1}$ do not form a 4-hole in P because otherwise b_i lies in the sector $S(a_{i-1}, a^*, a_{i+1}, a_i)$, which is impossible by Observation 4.6(ii).

Therefore the points $a^*, a_{i+1}, a_i, a_{i-1}$ are not in convex position. Now we show that a^* is inside the triangle $\triangle(a_{i-1}, a_{i+1}, a_i)$. The point a_i is not inside $\triangle(a_{i-1}, a_{i+1}, a^*)$, since, by Observation 4.7, a_i is to the right of a^* and since a^* is the rightmost inner point of A. Since a_{i-1} is to the left of $\overline{a^* a_i}$ and a_{i+1} is to the right of $\overline{a^* a_i}$, a^* is the inner point of $a^*, a_{i+1}, a_i, a_{i-1}$. Figure 4.6 gives an illustration.

Since $|B| \geq 5$, there is another a^*-wedge besides W_i and W_{i+1} that intersects ℓ. Now we show that all points of $B \setminus Q$ lie in a^*-wedges below W_{i+1}. The rays $\overrightarrow{b_i a_{i-1}}$ and $\overrightarrow{b_{i-1} b_{i+2}}$ both start

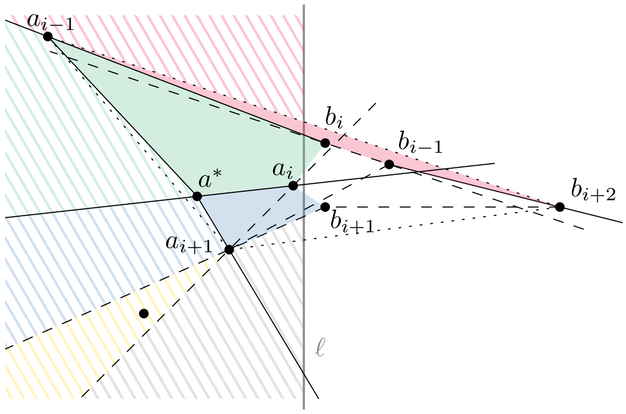

Figure 4.6.: Location of the points of $A \setminus Q$ in the proof of Lemma 4.11.

in W_i and then leave W_i. Moreover, the segment $b_i a_{i-1}$ intersects ℓ and $b_{i-1}b_{i+2}$ intersects $\overrightarrow{a^*a_i}$. As both b_i and b_{i-1} lie to the right of $\overline{a_{i-1}b_{i+2}}$, all points of $B \setminus Q$ that lie in an a^*-wedge above W_i also lie in the sector $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$. We recall that, by Observation 4.6(i), the sector $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$ is empty of points of P. Hence all points of $B \setminus Q$ lie in a^*-wedges below W_{i+1}.

We show that $i = 1$. That is, W_i is the topmost a^*-wedge that intersects ℓ. By Observation 4.7, a_{i+1} lies to the right of a^*. Since a_i and a_{i+1} are both to the right of a^* and since a^* is inside the triangle $\triangle(a_{i-1}, a_{i+1}, a_i)$, the point a_{i-1} is to the left of a^*. By Observation 4.7, we have $i = 1$.

Now we show that all points of $A \setminus Q$ lie to the left of $\overline{a_{i+1}a_i}$, to the right of $\overline{a_{i+1}b_{i+1}}$, and to the right of $\overline{a^*a_{i+1}}$. The yellow area in Figure 4.6 gives an illustration where the remaining points of $A \setminus Q$ lie. We recall that the sector $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$ (red shaded area in Figure 4.6) is empty of points of P. By Observation 4.9, both sets $\{a^*, a_i, b_i, a_{i-1}\}$ and $\{a^*, a_{i+1}, b_{i+1}, a_i\}$ form ℓ-divided 4-holes in P. By Observation 4.6(i), the two sectors $S(a^*, a_i, b_i, a_{i-1})$ (green shaded area in Figure 4.6) and $S(a^*, a_{i+1}, b_{i+1}, a_i)$ (blue shaded area in Figure 4.6) are thus empty of points of P. Therefore, no point of $A \setminus Q$ lies to the left of $\overline{a_{i+1}b_{i+1}}$. Since W is non-convex, every point of P that is to the left of $\overline{a^*a_{i+1}}$ lies in Q. Thus every point of $A \setminus Q$ lies to the right of $\overline{a^*a_{i+1}}$. Moreover, no point a of $A \setminus Q$ lies to the right of $\overline{a_{i+1}a_i}$ (gray area in Figure 4.6) because otherwise, a_{i+1} is an inner point of $\triangle(a_i, a^*, a)$, which is impossible since a^* is the rightmost inner point of A and a_{i+1} is to the right of a^*.

Now we have restricted where the points of $A \setminus Q$ lie. In the rest of the proof we show that the points $b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2}$ form an ℓ-divided 4-hole in P. We will then use the sectors $S(b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2})$ and $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$ to argue that $|B| = |B \cap Q| = 4$, which then contradicts the assumption $|B| \geq 5$.

We consider a_{i+2} and show that the points $a_{i+1}, a^*, a_{i-1}, a_{i+2}$ are in convex position. It suffices to show that a_{i+2} does not lie in the triangle $\triangle(a^*, a_{i-1}, a_{i+1})$ because of the cyclic order of $A \setminus \{a^*\}$ around a^*. Recall that a^* lies inside the triangle $\triangle(a_{i-1}, a_{i+1}, a_i)$, that b_{i+1} lies inside the triangle $\triangle(a_i, a_{i+1}, b_{i+2})$, and that b_{i-1} lies inside the triangle $\triangle(a_{i-1}, a_i, b_{i+2})$. Since the

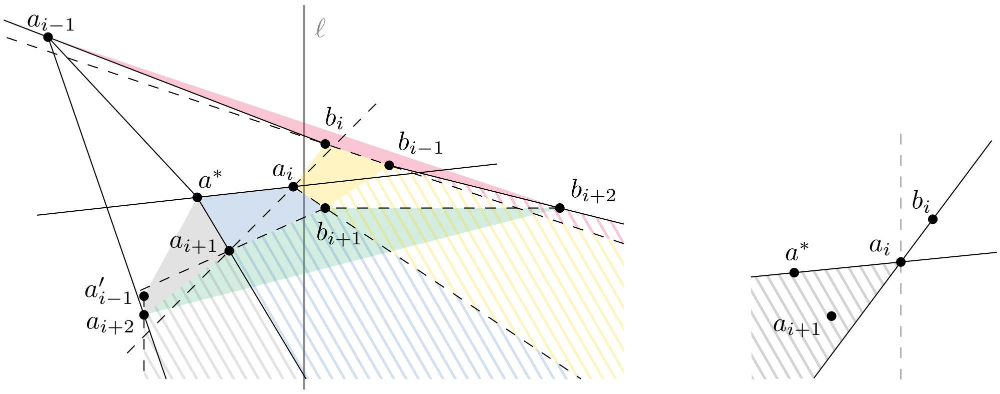

Figure 4.7.: Left: Location of the points of $B \setminus Q$. Right: The point a_{i+1} lies to the left of a_i.

triangles $\triangle(a_{i-1}, a_{i+1}, a_i)$, $\triangle(a_i, a_{i+1}, b_{i+2})$, and $\triangle(a_{i-1}, a_i, b_{i+2})$ are oriented counterclockwise along the boundary, the point a_i lies inside $\triangle(a_{i-1}, a_{i+1}, b_{i+2})$. Thus also the points a^*, b_i, b_{i+1} lie in the triangle $\triangle(a_{i-1}, a_{i+1}, b_{i+2})$. Consequently, the triangle $\triangle(a^*, a_{i-1}, a_{i+1})$ is contained in the union of the sectors $S(a_{i+1}, b_{i+1}, a_i, a^*)$ (blue shaded area in Figure 4.6) and $S(a^*, a_i, b_i, a_{i-1})$ (green shaded area in Figure 4.6). Thus a_{i+2} does not lie in the triangle $\triangle(a^*, a_{i-1}, a_{i+1})$ and the points $a_{i+1}, a^*, a_{i-1}, a_{i+2}$ are in convex position.

We now show that the sector $S(a_{i+1}, a^*, a_{i-1}, a_{i+2})$ is empty of points of P. If the quadrilateral $\square(a_{i+1}, a^*, a_{i-1}, a_{i+2})$ is not empty of points of P, then there is a point a'_{i-1} of A in $\triangle(a^*, a_{i-1}, a_{i+2})$. This is because $\triangle(a^*, a_{i+2}, a_{i+1})$ is empty of points of A due to the cyclic order of $A \setminus \{a^*\}$ around a^*. We can choose a'_{i-1} to be a point that is closest to the line $\overline{a^* a_{i+2}}$ among the points of A inside $\triangle(a^*, a_{i+2}, a_{i+1})$. If the quadrilateral $\square(a_{i+1}, a^*, a_{i-1}, a_{i+2})$ is empty of points of P, then we set $a'_{i-1} := a_{i-1}$.

By the choice of a'_{i-1}, the quadrilateral $\square(a_{i+1}, a^*, a'_{i-1}, a_{i+2})$ is empty of points of P. Since a_{i+1} and a_{i+2} are consecutive in the order around a^*, no point of A lies in the sector $S(a_{i+1}, a^*, a'_{i-1}, a_{i+2})$. By Observation 4.6(ii), the sector $S(a_{i+1}, a^*, a'_{i-1}, a_{i+2})$, the gray shaded area in Figure 4.7 (left), is empty of points of P. Since the sector $S(a_{i+1}, a^*, a_{i-1}, a_{i+2})$ is a subset of $S(a_{i+1}, a^*, a'_{i-1}, a_{i+2})$, the sector $S(a_{i+1}, a^*, a_{i-1}, a_{i+2})$ is empty of points of P.

We show that a_{i+1} is to the left of a_i and to the right of a_{i+2}. Recall that a_i lies to the right of a^* and to the left of b_i. The point b_i lies to the left of $\overline{a^* a_i}$ and the point a_{i+1} lies to the right of this line; see Figure 4.7 (right). The point a_{i+1} then lies to the left of a_i, since we know already that a_{i+1} lies to the left of $\overline{a_i b_i}$. Recall that a_{i+1} is to the right of a^*. Consequently, the point a_{i+2} lies to the left of a_{i+1}, as a_{i+2} lies to the right of $\overline{a^* a_{i+1}}$ and to the left of $\overline{a_{i+1} a_i}$.

Now we are ready to prove that the points $b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2}$ form an ℓ-divided 4-hole in P, the green area in Figure 4.7 (left). Recall that b_{i+2} and a_{i+2} both lie to the right of $\overline{a_{i+1} b_{i+1}}$, and that a_{i+2} is the leftmost and b_{i+2} is the rightmost of those four points. Altogether, we see that the points $b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2}$ are in convex position. The four sectors $S(b_{i+2}, a_{i-1}, b_i, b_{i-1})$, the red shaded area in Figure 4.7 (left), $S(b_{i-1}, b_i, a_i, b_{i+1})$, the yellow shaded area in Figure 4.7 (left), $S(b_{i+1}, a_i, a^*, a_{i+1})$, the blue shaded area in Figure 4.7 (left), and $S(a_{i+1}, a^*, a'_{i-1}, a_{i+2})$, the gray shaded area in Figure 4.7 (left) contain the quadrilateral $\square(b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2})$, the

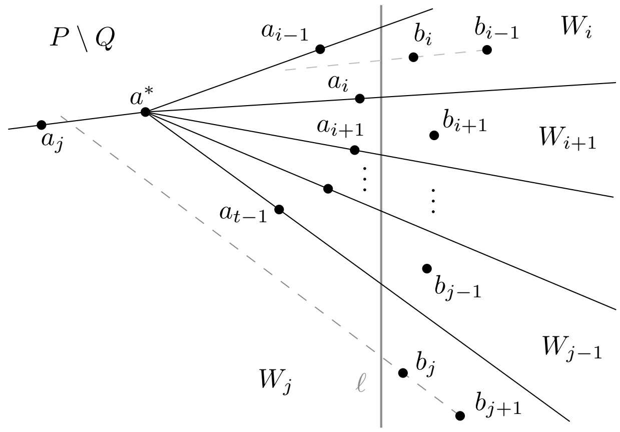

Figure 4.8.: An illustration of a^*-wedges $W_i, \dots, W_j$ in the proof of Lemma 4.12.

green area in Figure 4.7 (left). The sectors are empty of points of P by Observation 4.6(i). Consequently, the convex quadrilateral $\square(b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2})$ is an ℓ-divided 4-hole in P.

To finish the proof, recall that all points of $B \setminus Q$ lie in a^*-wedges below W_{i+1} as $i = 1$. Since a_{i+2} is to the left of a_{i+1}, the line $\overline{a_{i+2}a_{i+1}}$ intersects ℓ above $\ell \cap W_{i+2}$. The line $\overline{a_{i+1}b_{i+1}}$ also intersects ℓ above $\ell \cap W_{i+2}$, since a_{i+1} and b_{i+1} both lie in W_{i+1}. From $i = 1$, every point of $B \setminus Q$ is to the right of $\overline{a_{i+2}a_{i+1}}$ and to the right of $\overline{a_{i+1}b_{i+1}}$. Since the points $b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2}$ form an ℓ-divided 4-hole in P, Observation 4.6(i) implies that the sector $\overline{S(b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2})}$ is empty of points of P. Thus every point of $B \setminus Q$ lies to the left of $\overline{b_{i+1}b_{i+2}}$. Since $\overline{b_{i+1}b_{i+2}}$ intersects $\ell \cap W_{i+1}$ above $\ell \cap a_{i+1}b_{i+1}$ and since b_{i-1} lies to the left of b_{i+2} and to the left of $\overline{b_{i+1}b_{i+2}}$, every point of $B \setminus Q$ lies to the left of $\overline{b_{i-1}b_{i+2}}$ and to the right of b_{i+2}, and thus in the sector $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$. However, by Observation 4.6(i), this sector is empty of points of P. Thus we obtain $B = \{b_{i-1}, b_i, b_{i+1}, b_{i+2}\}$, which contradicts the assumption $|B| \geq 5$. $\square$

Next we show that if there is a sequence of consecutive a^*-wedges where the first and the last a^*-wedge both contain two points of B and every a^*-wedge in between them contains exactly one point of B, then there is an ℓ-divided 5-hole in P.

Lemma 4.12. *Let $P = A \cup B$ be an ℓ-divided set with A not in convex position and with $|A| \geq 5$ and $|B| \geq 6$. Let $W_i, \dots, W_j$ be consecutive a^*-wedges with $1 \leq i < j \leq t$, $w_i = 2 = w_j$, and $w_k = 1$ for every k with $i < k < j$. Then there is an ℓ-divided 5-hole in P.*

Proof. For $i = j - 1$, the statement follows by Lemma 4.11. Thus we assume $j \geq i + 2$. That is, we have at least three consecutive a^*-wedges. Suppose for contradiction that there is no ℓ-divided 5-hole in P. Let $W := \bigcup_{k=i}^{j} W_k$ and $Q := P \cap W$. By Lemma 4.8, there is also no ℓ-divided 5-hole in Q. Note that $|Q \cap B| = j - i + 3$. Also observe that $|Q \cap A| = j - i + 2$ if $a_{i-1} = a_j = a_t$ and $|Q \cap A| = j - i + 3$ otherwise. We label the points in $B \cap W_i$ as b_{i-1} and b_i so that b_{i-1} is to the right of b_i. Further, we label the only point in $B \cap W_k$ as b_k for each $i < k < j$, and the two points in $B \cap W_j$ as b_j and b_{j+1} so that b_{j+1} is to the right of b_j; see Figure 4.8.

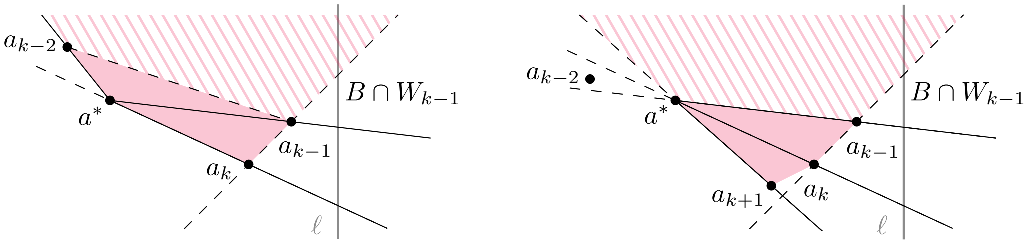

Figure 4.9.: An illustration of the proof of Claim 1.

Claim 1. *All points of* $B \cap (W_{k-1} \cup W_k \cup W_{k+1})$ *are to the right of* $\overline{a_k a_{k-1}}$ *for every* k *with* $i < k < j$.

The claim clearly holds for points from $B \cap W_k$. Thus it suffices to prove the claim only for points from $B \cup W_{k-1}$, as for points from $B \cup W_{k+1}$ it follows by symmetry. Since $i < k < j$, Observation 4.7 implies that the points a_{k-1} and a_k are both to the right of a^*.

We now distinguish the following two cases.

Case 1: The point a_{k-2} is to the left of $\overline{a^* a_k}$; see Figure 4.9 (left). Since a^* is the rightmost inner point of A, a_{k-1} does not lie inside the triangle $\triangle(a^*, a_k, a_{k-2})$ and thus $\square(a_{k-2}, a^*, a_k, a_{k-1})$ is a 4-hole in P. All points of $B \cap W_{k-1}$ lie to the right of $\overline{a^* a_{k-2}}$ and to the left of $\overline{a_{k-2} a_{k-1}}$. By Observation 4.6(ii), no point of $B \cap W_{k-1}$ lies in the sector $S(a_{k-2}, a^*, a_k, a_{k-1})$, the red shaded area in Figure 4.9 (left), and thus all points of $B \cap W_{k-1}$ are to the right of $\overline{a_k a_{k-1}}$.

Case 2: The point a_{k-2} is to the right of $\overline{a^* a_k}$; see Figure 4.9 (right). Since a_{k-1} and a_k are to the right of a^* and since a_{k-2} is to the left of $\overline{a^* a_{k-1}}$ and to the right of $\overline{a^* a_k}$, the point a_{k-2} is to the left of a^*. By Observation 4.7, we have $k = 2$. That is, W_{k-1} is the topmost a^*-wedge that intersects ℓ.

There is another a^*-wedge below W_{k+1}, since otherwise $|B| = |B \cap (W_{k-1} \cup W_k \cup W_{k+1})| \leq 2+1+2 = 5$, which is impossible according to the assumption $|B| \geq 6$. By Observation 4.7, the point a_{k+1} is to the right of a^*. Moreover, since a^* is the rightmost inner point of A, the point a_k does not lie inside the triangle $\triangle(a^*, a_{k+1}, a_{k-1})$. The points $a^*, a_{k+1}, a_k, a_{k-1}$ then form a 4-hole in P, which has a^* as the leftmost point.

By definition, all points of $B \cap W_{k-1}$ lie to the left of $\overline{a^* a_{k-1}}$. As the ray $\overrightarrow{a^* a_{k+1}}$ intersects ℓ, all points of $B \cap W_{k-1}$ lie also to the left of $\overline{a^* a_{k+1}}$. By Observation 4.6(ii), no point of $B \cap W_{k-1}$ lies in the sector $S(a^*, a_{k+1}, a_k, a_{k-1})$. Thus all points of $B \cap W_{k-1}$ lie to the right of $\overline{a_k a_{k-1}}$.

This finishes the proof of Claim 1.

We say that points p_1, p_2, p_3, p_4 form a *counterclockwise-oriented convex quadrilateral* if every triple (p_x, p_y, p_z) with $1 \leq x < y < z \leq 4$ is oriented counterclockwise.

Claim 2. *The points* $b_{i-1}, b_i, a_i, a_{i+1}$ *form a counterclockwise-oriented convex quadrilateral.*

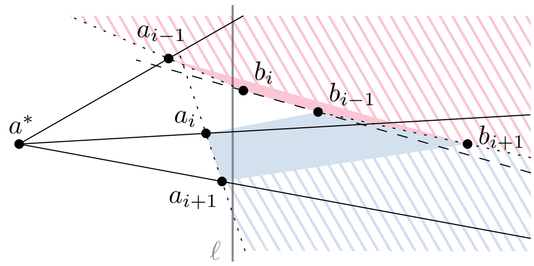

Figure 4.10.: The point b_{i+1} cannot lie to the left of $\overline{b_i b_{i-1}}$.

Due to Claim 1, the points b_{i-1} and b_i are both to the right of $\overline{a_{i+1}a_i}$. Thus the points a_i and a_{i+1} are both extremal points of those four points. Also the point b_{i-1} is extremal, since it is the rightmost of those four points. The point b_i does not lie inside the triangle $\triangle(a_{i+1}, a_i, b_{i-1})$, since, by Observation 4.10, b_i lies to the left of $\overline{a_i b_{i-1}}$. To finish the proof of Claim 2, it suffices to observe that the triples (b_{i-1}, b_i, a_i), (b_{i-1}, b_i, a_{i+1}), (b_{i-1}, a_i, a_{i+1}), and (b_i, a_i, a_{i+1}) are all oriented counterclockwise.

Claim 3. *The point b_{i+1} lies to the right of $\overline{b_i b_{i-1}}$.*

Suppose for contradiction that b_{i+1} lies to the left of $\overline{b_i b_{i-1}}$. We consider the five points $a_{i-1}, a_i, b_{i-1}, b_i, b_{i+1}$; see Figure 4.10. By Claim 1, the points b_{i-1}, b_i, and b_{i+1} lie to the right of $\overline{a_i a_{i-1}}$. Moreover, since b_{i-1} and b_i lie in W_i and since b_{i+1} lies in W_{i+1}, the points b_{i-1} and b_i both lie to the left of $\overline{a_i b_{i+1}}$. By Observation 4.10, the point a_{i-1} lies to the left of $\overline{b_i b_{i-1}}$ and b_{i+1} is to the right of b_{i-1}. Consequently, the points b_{i-1} and b_i lie in the triangle $\triangle(a_{i-1}, a_i, b_{i+1})$. Altogether, the points a_{i-1}, b_i, b_{i-1}, and b_{i+1} are in convex position.

By Claim 1, the points b_{i-1} and b_{i+1} lie to the right of $\overline{a_{i+1}a_i}$. Moreover, since b_{i-1} is to the left of b_{i+1} and to the left of $\overline{a_i b_{i+1}}$, the points b_{i+1}, b_{i-1}, a_i, and a_{i+1} are in convex position. Since there are no further points in W_i and W_{i+1}, the sets $\{a_{i-1}, b_i, b_{i-1}, b_{i+1}\}$ and $\{b_{i+1}, b_{i-1}, a_i, a_{i+1}\}$ are ℓ-divided 4-holes in P. By Observation 4.6(i), the point b_{i+2} lies neither in $S(a_{i-1}, b_i, b_{i-1}, b_{i+1})$ nor in $S(b_{i+1}, b_{i-1}, a_i, a_{i+1})$. Recall that the ray $\overrightarrow{b_{i-1}b_{i+1}}$ intersects $\overrightarrow{a^* a_i}$ and the ray $\overrightarrow{b_i a_{i-1}}$ does not intersect $\overrightarrow{a^* a_i}$. Therefore b_{i+2} is to the right of $\overline{a_i a_{i+1}}$. This contradicts Claim 1 and finishes the proof of Claim 3.

Claim 4. *For each k with $i < k < j$, the point b_k lies to the left of $\overline{a_k b_{i-1}}$ and to the left of b_{i-1}.*

We show by induction on k that

(i) the points b_{i-1}, b_{k-1}, a_{k-1}, and a_k form a counterclockwise-oriented convex quadrilateral, which has b_{i-1} as the rightmost point, and

(ii) the point b_k lies inside this convex quadrilateral and, in particular, to the left of $\overline{a_k b_{i-1}}$.

Claim 4 then clearly follows.

For the base case, we consider $k = i + 1$. By Claim 2, the points b_{i-1}, b_i, a_i, and a_{i+1} form a counterclockwise-oriented convex quadrilateral. By definition, b_{i-1} is the rightmost of those

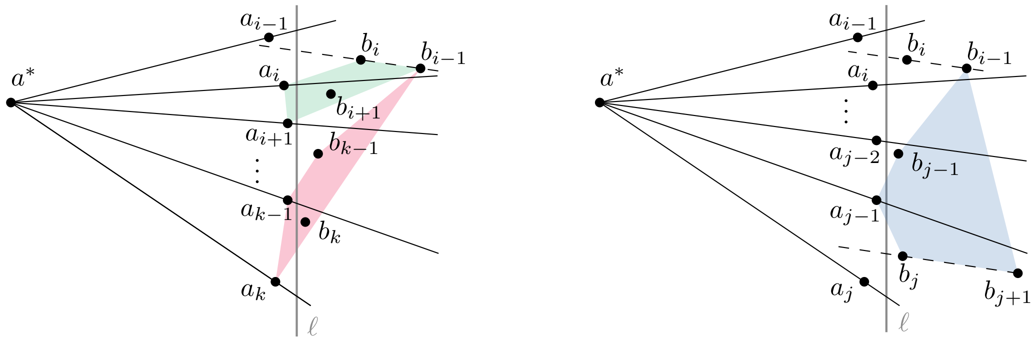

Figure 4.11.: An illustration of the proofs of Claim 4 and Lemma 4.12.

four points; see Figure 4.11 (left) for an illustration. The point b_{i+1} lies to the right of $\overline{a_{i+1}a_i}$ and, by Claim 3, to the right of $\overline{b_ib_{i-1}}$. Moreover, since b_{i+1} lies in W_{i+1}, it lies to the right of $\overline{a_ib_i}$. By Observation 4.6(i), b_{i+1} does not lie in the sector $S(b_{i-1}, b_i, a_i, a_{i+1})$. Consequently, b_{i+1} lies inside the quadrilateral $\square(b_{i-1}, b_i, a_i, a_{i+1})$.

For the inductive step, let $i+1 < k < j$. By the inductive assumption, the point b_{k-1} lies to the left of $\overline{a_{k-1}b_{i-1}}$ and to the left of b_{i-1}. By Claim 1, b_{k-1} lies to the right of $\overline{a_ka_{k-1}}$. Hence, the points a_k and b_{i-1} both lie to the right of $\overline{a_{k-1}b_{k-1}}$. Recall that the points $b_{i-1}, b_{k-1}, a_{k-1}, a_k$ lie to the right of a^*. Since b_{i-1} is the first and a_k is the last in the clockwise order around a^*, the points $b_{i-1}, b_{k-1}, a_{k-1}, a_k$ form a counterclockwise-oriented convex quadrilateral,

Recall that the points b_{k-1} and b_k both lie to the right of $\overline{a_ka_{k-1}}$ and that b_{k-1} is to the left of $\overline{a_{k-1}b_{i-1}}$. Since $b_k \in W_k$, the point b_k lies to the right of $\overline{a_{k-1}b_{i-1}}$. Therefore the clockwise order of $\{b_{k-1}, b_{i-1}, b_k\}$ around a_{k-1} is b_{k-1}, b_{i-1}, b_k. Since b_{i-1} is not contained in $W_{k-1} \cup W_k$, the point b_{i-1} is not contained in the triangle $\triangle(a_{k-1}, b_k, b_{k-1})$. Consequently, the points $a_{k-1}, b_k, b_{i-1}, b_{k-1}$ form a convex quadrilateral and, in particular, b_k lies to the right of $\overline{b_{k-1}b_{i-1}}$; see Figure 4.11 (left) for an illustration. Since b_k lies in W_k, it lies to the right of $\overline{a_{k-1}b_{k-1}}$. By Observation 4.6(i), the point b_k does not lie in the sector $S(b_{i-1}, b_{k-1}, a_{k-1}, a_k)$. Thus b_k lies inside the quadrilateral $\square(b_{i-1}, b_{k-1}, a_{k-1}, a_k)$. This finishes the proof of Claim 4.

Using Claim 4, we now finish the proof of Lemma 4.12, by finding an ℓ-divided 5-hole in Q and thus obtaining a contradiction with the assumption that there is no ℓ-divided 5-hole in P. In the following, we assume, without loss of generality, that b_{j+1} is to the right of b_{i-1}. Otherwise we can consider a vertical reflection of P.

We consider the polygon $\mathcal{P}$ through the points $b_{i-1}, b_{j-1}, a_{j-1}, b_j, b_{j+1}$ and we show that $\mathcal{P}$ is convex and empty of points of Q; see Figure 4.11 (right) for an illustration. This will give us an ℓ-divided 5-hole in Q.

We show that $\mathcal{P}$ is convex by proving that every point of $\{b_{i-1}, b_{j-1}, a_{j-1}, b_j, b_{j+1}\}$ is a convex vertex of $\mathcal{P}$. The point a_{j-1} is a convex vertex of $\mathcal{P}$ because it is the leftmost point in $\mathcal{P}$. The point b_{i-1} is a convex vertex of $\mathcal{P}$ because all points of $\mathcal{P}$ lie to the right of a^* and b_{i-1} is the topmost point in the clockwise order around a^*. The point b_{j+1} is a convex vertex of $\mathcal{P}$ because b_{j+1} is the rightmost point in $\mathcal{P}$ by Claim 4 and by the assumption that b_{j+1} is to the right of b_{i-1}. The point b_{j-1} is a convex vertex of $\mathcal{P}$ because b_{j-1} lies to the left of $\overline{a_{j-1}b_{i-1}}$ by Claim 4 while b_j and b_{j+1} both lie to the right of this line. The point b_j is a convex vertex of $\mathcal{P}$ because, by Observation 4.10, b_j lies to the right of $\overline{a_{j-1}b_{j+1}}$ while b_{j-1} and b_{i-1} both lie to the right of this line. Consequently, $\mathcal{P}$ is a convex pentagon with vertices from both A and B. Moreover, by Claim 4, all points b_k with $i < k < j$ lie to the left of $\overline{a_k b_{i-1}}$. Since b_i is to the left of $\overline{b_{j-1}b_{i-1}}$, $\mathcal{P}$ is thus empty of points of Q, which gives us a contradiction with the assumption that there is no ℓ-divided 5-hole in P. $\square$

We now use Lemma 4.12 to show the following upper bound on the total number of points of B in a sequence $W_i, \dots, W_j$ of consecutive a^*-wedges with $w_i, \dots, w_j \leq 2$.

Corollary 4.13. *Let $P = A \cup B$ be an ℓ-divided set with no ℓ-divided 5-hole, with A not in convex position, and with $|A| \geq 5$ and $|B| \geq 6$. For $1 \leq i \leq j \leq t$, let $W_i, \dots, W_j$ be consecutive a^*-wedges with $w_k \leq 2$ for every k with $i \leq k \leq j$. Then $\sum_{k=i}^{j} w_k \leq j - i + 2$.*

Proof. Let n_0, n_1, and n_2 be the number of a^*-wedges from $W_i, \dots, W_j$ with 0, 1, and 2 points of B, respectively. Due to Lemma 4.12, we can assume that between any two a^*-wedges from $W_i, \dots, W_j$ with two points of B each, there is an a^*-wedge with no point of B. Thus $n_2 \leq n_0 + 1$. Since $n_0 + n_1 + n_2 = j - i + 1$, we have $\sum_{k=i}^{j} w_k = 0n_0 + 1n_1 + 2n_2 = (j - i + 1) + (n_2 - n_0) \leq j - i + 2$. $\square$

4.5.2. Computer-assisted results

We now provide lemmas that are key ingredients in the proof of Theorem 4.2. All these lemmas have computer-aided proofs. Each result was verified by two independent implementations, which are also based on different abstractions of point sets; see below for details.

Lemma 4.14. *Let $P = A \cup B$ be an ℓ-divided set with $|A| = 5$, $|B| = 6$, and with A not in convex position. Then there is an ℓ-divided 5-hole in P.*

Lemma 4.15. *Let $P = A \cup B$ be an ℓ-divided set with no ℓ-divided 5-hole in P, $|A| = 5$, $4 \leq |B| \leq 6$, and with A in convex position. Then for every point a of A, every convex a-wedge contains at most two points of B.*

Lemma 4.16. *Let $P = A \cup B$ be an ℓ-divided set with no ℓ-divided 5-hole in P, $|A| = 6$, and $|B| = 5$. Then for each point a of A, every convex a-wedge contains at most two points of B.*

Lemma 4.17. *Let $P = A \cup B$ be an ℓ-divided set with no ℓ-divided 5-hole in P, $5 \leq |A| \leq 6$, $|B| = 4$, and with A in convex position. Then for every point a of A, if the non-convex a-wedge is empty of points of B, every a-wedge contains at most two points of B.*

To prove these lemmas, we employ an exhaustive computer search through all combinatorially different sets of $|P| \leq 11$ points in the plane. Notice that these statements only depend on the order type of the point set, and thus, we only need to check a finite number of equivalence classes of point sets.

Lemmas 4.14 to 4.17 were verified by two different programs available online. The first one [180] uses the order type database [12, 26] and runs in few hours on a standard computer. The second one [41] neither uses the order type database nor the program used to generate the database. Instead, it relies on the description of point sets by so-called *signature functions* [42, 91]. In this description, points are sorted according to their x-coordinates and every unordered triple of points is represented by a sign from $\{-,+\}$, where the sign is $-$ if the triple traced in the order by increasing x-coordinates is oriented clockwise and the sign is $+$ otherwise. Every 4-tuple of points is then represented by four signs of its triples, which are ordered lexicographically. There are only eight 4-tuples of signs that we can obtain (out of 16 possible ones); see [42, Theorem 3.2] or [91, Theorem 7] for details. All possible signature functions are generated using a simple depth-first search algorithm. Then the conditions of the lemmas are checked for every signature. The running time of each of the programs in this implementation takes up to a few hundreds of hours.

4.5.3. Applications of the computer-assisted results

Here we present some applications of the computer-assisted results from Section 4.5.2.

Lemma 4.18. *Let $P = A \cup B$ be an ℓ-divided set with no ℓ-divided 5-hole in P, with $|A| \geq 6$, and with A not in convex position. Then the following two conditions are satisfied.*

(i) Let W_i, W_{i+1}, W_{i+2} be three consecutive a^*-wedges whose union is convex and contains at least four points of B. Then $w_i, w_{i+1}, w_{i+2} \leq 2$.

(ii) Let $W_i, W_{i+1}, W_{i+2}, W_{i+3}$ be four consecutive a^*-wedges whose union is convex and contains at least four points of B. Then $w_i, w_{i+1}, w_{i+2}, w_{i+3} \leq 2$.

Proof. To show part (i), let $W := W_i \cup W_{i+1} \cup W_{i+2}$, $A' := A \cap W$, $B' := B \cap W$, and $P' := A' \cup B'$. Since W is convex, P' is an island of P and thus there is no ℓ-divided 5-hole in P'. Note that $|A'| = 5$ and A' is in convex position. If $|B'| \leq 5$, then every convex a^*-wedge in P' contains at most two points of B' by Lemma 4.15 applied to P'. So assume that $|B'| \geq 6$. We remove points from P' from the right, if necessary, to obtain $P'' = A' \cup B''$, where B'' contains exactly six points of B'. Note that there is no ℓ-divided 5-hole in P'', since P'' is an island of P'. By Lemma 4.15, each a^*-wedge in P'' contains exactly two points of B''. Let $\tilde{B}$ be the set of points of B that are to the left of the rightmost point of B'', including this point, and let $\tilde{P} := A \cup \tilde{B}$. Note that $B'' \subseteq \tilde{B}$. Since $|B''| = 6$ and since $W \cap \tilde{B} = B''$, each of the a^*-wedges W_i, W_{i+1}, W_{i+2} contains exactly two points of $\tilde{B}$. The a^*-wedges W_i, W_{i+1}, and W_{i+2} are also a^*-wedges in $\tilde{P}$. Thus, Lemma 4.11 applied to $\tilde{P}$ and W_i, W_{i+1} then gives us an ℓ-divided 5-hole in $\tilde{P}$. From the choice of $\tilde{P}$, we then have an ℓ-divided 5-hole in P, a contradiction.

To show part (ii), let $W := W_i \cup W_{i+1} \cup W_{i+2} \cup W_{i+3}$, $A' := A \cap W$, $B' := B \cap W$, and $P' := A' \cup B'$. Since W is convex, P' is an island of P and thus there is no ℓ-divided 5-hole in P'. Note that $|A'| = 6$ and A' is in convex position. If $|B'| = 4$, then the statement follows from Lemma 4.17 applied to P' since a^* is an extremal point of P'. If $|B'| = 5$, then the statement follows from Lemma 4.16 applied to P' and thus we can assume $|B'| \geq 6$. Suppose for contradiction that $w_j \geq 3$ for some $i \leq j \leq i + 3$. If necessary, we remove points from

P from the right to obtain P'' so that $B'' := P'' \cap B$ contains exactly six points of $W \cap B$. By applying part (i) for P'' and $W_i \cup W_{i+1} \cup W_{i+2}$ and $W_{i+1} \cup W_{i+2} \cup W_{i+3}$, we obtain that $|B'' \cap W_i|, |B'' \cap W_{i+3}| = 3$ and $|B'' \cap W_{i+1}|, |B'' \cap W_{i+2}| = 0$. Let b be the rightmost point from $P'' \cap W$. By Lemma 4.16 applied to $W \cap (P'' \setminus \{b\})$, there are at most two points of $B'' \setminus \{b\}$ in every a^*-wedge in $W \cap (P'' \setminus \{b\})$. This contradicts the fact that either $|(B'' \cap W_i) \setminus \{b\}| = 3$ or $|(B'' \cap W_{i+3}) \setminus \{b\}| = 3$. □

4.5.4. Extremal points of ℓ-critical sets

Recall the definition of ℓ-critical sets: An ℓ-divided point set $C = A \cup B$ is called ℓ-critical if neither $C \cap A$ nor $C \cap B$ is in convex position and if for every extremal point x of C, one of the sets $(C \setminus \{x\}) \cap A$ and $(C \setminus \{x\}) \cap B$ is in convex position.

In this section, we consider an ℓ-critical set $C = A \cup B$ with $|A|, |B| \geq 5$. We first show that C has at most two extremal points in A and at most two extremal points in B. Later, under the assumption that there is no ℓ-divided 5-hole in C, we show that $|B| \leq |A| - 1$ if A contains two extremal points of C (Section 4.5.4) and that $|B| \leq |A|$ if B contains two extremal points of C (Section 4.5.4).

Lemma 4.19. *Let $C = A \cup B$ be an ℓ-critical set. Then the following statements are true.*

(i) If $|A| \geq 5$, then $|A \cap \partial\,\mathrm{conv}(C)| \leq 2$.

(ii) If $A \cap \partial\,\mathrm{conv}(C) = \{a, a'\}$, then a^* is the only interior point in A and every point of $A \setminus \{a, a'\}$ lies in the convex region spanned by the lines $\overline{a^*a}$ and $\overline{a^*a'}$ that does not have any of a and a' on its boundary.

(iii) If $A \cap \partial\,\mathrm{conv}(C) = \{a, a'\}$, then the a^*-wedge that contains a and a' contains no point of B.

By symmetry, analogous statements hold for B.

Proof. To show statement (i), suppose for contradiction that $|A \cap \partial\,\mathrm{conv}(C)| \geq 3$. Let a, a', and a'' be three such consecutive points. If there is no point of A in the triangle $\triangle(a, a', a'')$ spanned by the points a, a', and a'', then $A \setminus \{a'\}$ is not in convex position. This is impossible, since C is an ℓ-critical set. If there is at least one point $a^{(1)}$ in $\triangle(a, a', a'')$, then we consider an arbitrary point $a^{(2)}$ from $A \setminus \{a, a', a'', a^{(1)}\}$. Such a point $a^{(2)}$ exists, since $|A| \geq 5$. The point $a^{(1)}$ lies inside one of the triangles $\triangle(a, a', a^{(2)})$, $\triangle(a, a'', a^{(2)})$, or in $\triangle(a', a'', a^{(2)})$ and thus one of the sets $A \setminus \{a''\}$, $A \setminus \{a'\}$, or $A \setminus \{a\}$ is not in convex position, which is again impossible. In any case, C cannot be ℓ-critical and we obtain a contradiction.

To show statement (ii), assume that $A \cap \partial\,\mathrm{conv}(C) = \{a, a'\}$. Every triangle in A with a point of A in its interior has a and a' as vertices, as otherwise $A \setminus \{a\}$ or $A \setminus \{a'\}$ is not in convex position, which is impossible. Consider points $a^{(1)}$ and $a^{(2)}$ from A such that $\triangle(a, a', a^{(1)})$ contains $a^{(2)}$. Denote by R the region bounded by $\overline{aa^{(2)}}$ and $\overline{a'a^{(2)}}$ that contains $a^{(1)}$. If there is a point $a^{(3)}$ in $A \setminus (R \cup \{a, a'\})$ then $a^{(2)}$ lies in one of $\triangle(a, a^{(1)}, a^{(3)})$ and $\triangle(a', a^{(1)}, a^{(3)})$, implying that $A \setminus \{a\}$ or $A \setminus \{a'\}$ is not in convex position. Hence all points of $A \setminus \{a, a', a^{(2)}\}$ lie in R. Moreover, any further interior point $a^{(4)}$ from $A \cap R$ lies in some triangle $\triangle(a, a', a^{(5)})$ for some $a^{(5)} \in A \cap R$. Thus, $a^{(4)}$ also lies in one of the triangles $\triangle(a, a^{(2)}, a^{(5)})$ or $\triangle(a', a^{(2)}, a^{(5)})$. This implies that $A \setminus \{a\}$ or $A \setminus \{a'\}$ is not in convex position. Hence $a^{(2)}$ is the only interior point of A.

To show statement (iii), assume that $A \cap \partial \operatorname{conv}(C) = \{a, a'\}$. Let W_i be the wedge that contains a and a'. Since a and a' are the only extremal points of C contained in A, the segment aa' is an edge of $\operatorname{conv}(C)$. The points a, a', and a^* all lie in A and thus the triangle $\triangle(a, a', a^*)$ contains no points of B. Since all points of C lie in the closed halfplane that is determined by the line $\overline{aa'}$ and that contains a^*, the wedge W_i contains no points of B. $\square$

We remark that the assumption $|A| \geq 5$ in part (i) of Lemma 4.19 is necessary. In fact, arbitrarily large ℓ-critical sets with only four points in A and with three points of A on $\partial \operatorname{conv}(C)$ exist, and analogously for B; see Figure 4.2(c) for an illustration.

Lemma 4.20. *Let $C = A \cup B$ be an ℓ-critical set with no ℓ-divided 5-hole in C and with $|A| \geq 6$. Then $w_i \leq 3$ for every $1 < i < t$. Moreover, if $|A \cap \partial \operatorname{conv}(C)| = 2$, then $w_1, w_t \leq 3$.*

Proof. Recall that, since C is ℓ-critical, we have $|B| \geq 4$. Let i be an integer with $1 \leq i \leq t$. We assume that there is a point a in $A \cap \partial \operatorname{conv}(C)$, which lies outside of W_i, as otherwise there is nothing to prove for W_i (either $|A \cap \partial \operatorname{conv}(C)| = 1$ and $i \in \{1, t\}$ or $|A \cap \partial \operatorname{conv}(C)| = 2$ and, by Lemma 4.19(iii), $W_i \cap B = \emptyset$). We consider $C' := C \setminus \{a\}$. Since C is an ℓ-critical set, $A' := C' \cap A$ is in convex position. Thus, there is a non-convex a^*-wedge W' of C'. Since W' is non-convex, all other a^*-wedges of C' are convex. Moreover, since W' is the union of the two a^*-wedges of C that contain a, all other a^*-wedges of C' are also a^*-wedges of C. Let W be the union of all a^*-wedges of C that are not contained in W'. Note that W is convex and contains at least $|A| - 3 \geq 3$ a^*-wedges of C. Since $|A| \geq 6$, the statement follows from Lemma 4.18(i). $\square$

Two extremal points of C in A

Proposition 4.21. *Let $C = A \cup B$ be an ℓ-critical set with no ℓ-divided 5-hole in C, with $|A|, |B| \geq 6$, and with $|A \cap \partial \operatorname{conv}(C)| = 2$. Then $|B| \leq |A| - 1$.*

Proof. Since $|A \cap \partial \operatorname{conv}(C)| = 2$, Lemma 4.20 implies that $w_i \leq 3$ for every $1 \leq i \leq t$. Let a and a' be the two points in $A \cap \partial \operatorname{conv}(C)$. By Lemma 4.19(ii), all points of $A \setminus \{a, a'\}$ lie in the convex region R spanned by the lines $\overline{a^*a}$ and $\overline{a^*a'}$ that does not have any of a and a' on its boundary. That is, without loss of generality, $a = a_{h-1}$ and $a' = a_h$ for some $1 \leq h \leq |A| - 1$ and, by Lemma 4.19(iii), we have $w_h = 0$. Since all points of $A \setminus \{a, a'\}$ lie in the convex region R, the regions $W := \operatorname{cl}(\mathbb{R}^2 \setminus (W_{h-1} \cup W_h))$ and $W' := \operatorname{cl}(\mathbb{R}^2 \setminus (W_h \cup W_{h+1}))$ are convex. Here $\operatorname{cl}(X)$ denotes the closure of a set $X \subseteq \mathbb{R}^2$. Recall that the indices of the a^*-wedges are considered modulo $|A| - 1$ and that $\mathbb{R}^2$ is the union of all a^*-wedges.

First, suppose for contradiction that $|A| = 6$ and $|B| \geq 6$. There are exactly five a^*-wedges $W_1, \ldots, W_5$, and only four of them can contain points of B, since $w_h = 0$. We apply Lemma 4.18(i) to W and to W'. An easy case analysis shows that either $w_i \leq 2$ for every $1 \leq i \leq t$ or $w_{h-1}, w_{h+1} = 3$ and $w_i = 0$ for every $i \notin \{h-1, h+1\}$. In the first case, Corollary 4.13 implies that $|B| \leq 5$ and in the latter case Lemma 4.16 applied to $P \setminus \{b\}$, where b is the rightmost point of B, gives $|B| \leq 5$, a contradiction. Hence, we assume $|A| \geq 7$.

Claim 1. *For $1 \leq k \leq t-3$, if one of the four consecutive a^*-wedges W_k, W_{k+1}, W_{k+2}, or W_{k+3} contains 3 points of B, then $w_k + w_{k+1} + w_{k+2} + w_{k+3} = 3$.*

There are $|A| - 1 \geq 6$ a^*-wedges and, in particular, W and W' are both unions of at least four a^*-wedges. For every W_i with $w_i = 3$ and $1 \leq i \leq t$, the a^*-wedge W_i is either contained in W or in W'. Thus we can find four consecutive a^*-wedges $W_k, W_{k+1}, W_{k+2}, W_{k+3}$ whose union is convex and contains W_i. Lemma 4.18(ii) implies that each of $W_k, W_{k+1}, W_{k+2}, W_{k+3}$ except of W_i is empty of points of B. This finishes the proof of Claim 1.

Claim 2. *For all integers i and j with $1 \leq i < j \leq t$, we have $\sum_{k=i}^{j} w_k \leq j - i + 2$.*

Let $S := (w_i, \ldots, w_j)$ and let S' be the subsequence of S obtained by removing every 1-entry from S. If S contains only 1-entries, the statement clearly follows. Thus we can assume that S' is non-empty. Recall that S' contains only 0-, 2-, and 3-entries, since $w_i \leq 3$ for all $1 \leq i \leq t$. Due to Claim 1, there are at least three consecutive 0-entries between every pair of nonzero entries of S' that contains a 3-entry. Together with Lemma 4.12, this implies that there is at least one 0-entry between every pair of 2-entries in S'.

By applying the following iterative procedure, we show that $\sum_{s \in S'} s \leq |S'| + 1$. While there are at least two nonzero entries in S', we remove the first nonzero entry s from S'. If $s = 2$, then we also remove the 0-entry from S' that succeeds s in S. If $s = 3$, then we also remove the two consecutive 0-entries from S' that succeed s in S'. The procedure stops when there is at most one nonzero element s' in the remaining subsequence S'' of S'. If $s' = 3$, then S'' contains at least one 0-entry and thus S'' contains at least $s' - 1$ elements. Since the number of removed elements equals the sum of the removed elements in every step of the procedure, we have $\sum_{s \in S'} s \leq |S'| + 1$. This implies

$$\sum_{k=i}^{j} w_k = \sum_{s \in S} s = |S| - |S'| + \sum_{s \in S'} s \leq |S| - |S'| + |S'| + 1 = j - i + 2$$

and finishes the proof of Claim 2.

If W_h does not intersect ℓ, that is, $t < h \leq |A| - 1$, then the statement follows from Claim 2 applied with $i = 1$ and $j = t$. Otherwise, we have $h = 1$ or $h = t$ and we apply Claim 2 with $(i,j) = (2,t)$ or $(i,j) = (1, t-1)$, respectively. Since $t \leq |A| - 1$ and $w_h = 0$, this gives us $|B| \leq |A| - 1$. □

Two extremal points of C in B

Proposition 4.22. *Let $C = A \cup B$ be an ℓ-critical set with no ℓ-divided 5-hole in C, with $|A|, |B| \geq 6$, and with $|B \cap \partial\, \mathrm{conv}(C)| = 2$. Then $|B| \leq |A|$.*

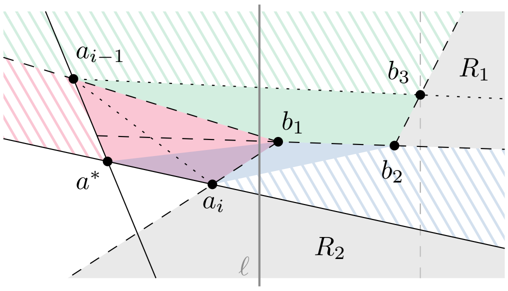

Figure 4.12.: An illustration of the proof of Proposition 4.22.

Proof. If $w_k \leq 2$ for all $1 \leq k \leq t$, then the statement follows from Corollary 4.13, since $|B| = \sum_{k=1}^{t} w_k \leq t+1 \leq |A|$. Therefore we assume that there is an a^*-wedge W_i that contains at least three points of B. Let b_1, b_2, and b_3 be the three leftmost points in $W_i \cap B$ from left to right. Without loss of generality, we assume that b_3 is to the left of $\overline{b_1 b_2}$. Otherwise we can consider a vertical reflection of P. Figure 4.12 gives an illustration.

Let R_1 be the region that lies to the left of $\overline{b_1 b_2}$ and to the right of $\overline{b_2 b_3}$ and let R_2 be the region that lies to the right of $\overline{a_i b_1}$ and to the right of $\overline{a^* a_i}$. Let $B' := B \setminus \{b_1, b_2, b_3\}$.

Claim 1. *Every point of B' lies in $R_1 \cup R_2$.*

We first show that every point of B' that lies to the left of $\overline{b_1 b_2}$ lies in R_1. Then we show that every point of B' that lies to the right of $\overline{b_1 b_2}$ lies in R_2.

By Observation 4.10, both lines $\overline{b_1 b_2}$ and $\overline{b_1 b_3}$ intersect the segment $a_{i-1}a_i$. Since the segment $a_{i-1}b_1$ intersects ℓ and since b_1 is the leftmost point of $W_i \cap B$, all points of B' that lie to the left of $\overline{b_1 b_2}$ lie to the left of $\overline{a_{i-1} b_1}$. The four points a_{i-1}, b_1, b_2, b_3 form an ℓ-divided 4-hole in P, since a_{i-1} is the leftmost and b_3 is the rightmost point of a_{i-1}, b_1, b_2, b_3 and both a_{i-1} and b_3 lie to the left of $\overline{b_1 b_2}$. By Observation 4.6(i), the sector $S(a_{i-1}, b_1, b_2, b_3)$ is empty of points of P (green shaded area in Figure 4.12). Altogether, all points of B' that lie to the left of $\overline{b_1 b_2}$ are to the right of $\overline{b_2 b_3}$ and thus lie in R_1.

Since the segment $a_i b_1$ intersects ℓ and since b_1 is the leftmost point of $W_i \cap B$, all points of B' that lie to the right of $\overline{b_1 b_2}$ lie to the right of $\overline{a_i b_1}$. By Observation 4.6(i), the sector $S(b_1, b_2, b_3, a_{i-1})$ is empty of points of P. Combining this with the fact that a^* is to the right of $\overline{a_{i-1} b_3}$, we see that a^* lies to the right of $\overline{b_1 b_2}$. Since b_1 and b_2 both lie to the left of $\overline{a^* a_i}$ and since a^* and a_i both lie to the right of $\overline{b_1 b_2}$, the points b_2, b_1, a^*, a_i form an ℓ-divided 4-hole in P. By Observation 4.6(i), the sector $S(b_2, b_1, a^*, a_i)$ (blue shaded area in Figure 4.12) is empty of points of P. Altogether, all points of B' that lie to the right of $\overline{b_1 b_2}$ are to the right of $\overline{a^* a_i}$ and to the right of $\overline{a_i b_1}$ and thus lie in R_2. This finishes the proof of Claim 1.

Claim 2. *If b_4 is a point from $B' \setminus R_1$, then b_2 lies inside the triangle $\triangle(b_3, b_1, b_4)$.*

By Claim 1, b_4 lies in R_2 and thus to the right of $\overline{a_i b_1}$ and to the right of $\overline{a^* a_i}$. We recall that b_4 lies to the right of $\overline{b_1 b_2}$.

We distinguish two cases. First, we assume that the points b_2, b_3, b_1, a_i are in convex position. Then b_2, b_3, b_1, a_i form an ℓ-divided 4-hole in P and, by Observation 4.6(i), the sector $S(b_2, b_3, b_1, a_i)$ is empty of points from P. Thus b_4 lies to the right of $\overline{b_2 b_3}$ and the statement follows.

Second, we assume that the points b_2, b_3, b_1, a_i are not in convex position. Due to Observation 4.10, b_2 and b_3 both lie to the right of $\overline{a_i b_1}$. Moreover, since b_3 is the rightmost of those four points, b_2 lies inside the triangle $\triangle(b_3, b_1, a_i)$. In particular, a_i lies to the right of $\overline{b_2 b_3}$. Therefore, since b_2 and b_3 are to the left of $\overline{a^* a_i}$, the line $\overline{b_2 b_3}$ intersects ℓ in a point p above $\ell \cap \overline{a^* a_i}$. Let q be the point $\ell \cap \overline{b_1 b_2}$. Note that q is to the left of $\overline{a^* a_i}$. The point b_4 is to the right of $\overline{b_2 b_3}$, as otherwise b_4 lies in $\triangle(p, q, b_2)$, which is impossible because the points p, q, b_2 are in W_i while b_4 is not. Altogether, b_2 is inside $\triangle(b_3, b_1, b_4)$ and this finishes the proof of Claim 2.

Claim 3. *Either every point of B' is to the right of b_3 or b_3 is the rightmost point of B.*

By Observation 4.6(i), the sector $S(b_3, a_{i-1}, b_1, b_2)$ is empty of points of P and thus all points of $B' \cap R_1$ lie to the left of $\overline{a_{i-1} b_3}$ and, in particular, to the right of b_3.

Suppose for contradiction that the claim is not true. That is, there is a point $b_4 \in B'$ that is the rightmost point in B and there is a point $b_5 \in B'$ that is to the left of b_3. Note that b_4 is an extremal point of C. By Claim 1 and by the fact that all points of $B' \cap R_1$ lie to the right of b_3, b_5 lies in $R_2 \setminus R_1$. By Claim 2, b_2 lies in the triangle $\triangle(b_1, b_5, b_3)$, and thus $B \setminus \{b_4\}$ is not in convex position. This contradicts the assumption that C is an ℓ-critical island. This finishes the proof of Claim 3.

Claim 4. *The point b_3 is the third leftmost point of B. In particular, W_i is the only a^*-wedge with at least three points of B.*

Suppose for contradiction that b_3 is not the third leftmost point of B. Then by Claim 3, b_3 is the rightmost point of B and therefore an extremal point of B. This implies that $B' \subseteq R_2 \setminus R_1$, since all points of $B' \cap R_1$ lie to the right of b_3. By Claim 2, each point of B' then forms a non-convex quadrilateral together with b_1, b_2, and b_3. Since neither b_1 nor b_2 are extremal points of C and since $|B \cap \partial \operatorname{conv}(C)| = 2$, there is a point $b_4 \in B$ that is an extremal point of C. Since $|B| \geq 5$, the set $C \setminus \{b_4\}$ has none of its parts separated by ℓ in convex position, which contradicts the assumption that C is an ℓ-critical set. Since W_i is an arbitrary a^*-wedge with $w_i \geq 3$, Claim 4 follows.

Claim 5. *Let W be a union of four consecutive a^*-wedges that contains W_i. Then $|W \cap B| \leq 4$.*

Suppose for contradiction that $|W \cap B| \geq 5$. Let $C' := C \cap W$. Note that $|C' \cap A| = 6$ and that a^*, a_{i-1}, a_i lie in C'. By Lemma 4.8, there is no ℓ-divided 5-hole in C'. We obtain C'' by removing points from C' from the right until $|C'' \cap B| = 5$. Since C'' is an island of C', there is no ℓ-divided 5-hole in C''. From Claim 4 we know that b_1, b_2, b_3 are the three leftmost points in C and thus lie in C''. We apply Lemma 4.16 to C'' and, since b_1, b_2, b_3 lie in a convex a^*-wedge of C'', we obtain a contradiction. This finishes the proof of Claim 5.

We now complete the proof of Proposition 4.22. First, we assume that $1 \leq i \leq 4$. Let $W := W_1 \cup W_2 \cup W_3 \cup W_4$. By Claim 5, $|W \cap B| \leq 4$. Claim 4 implies that $w_k \leq 2$ for every k with $5 \leq k \leq t$. By Corollary 4.13, we have

$$|B| = \sum_{k=1}^{4} w_k + \sum_{k=5}^{t} w_k \leq 4 + (t-3) = t+1 \leq |A|.$$

The case $t-3 \leq i \leq t$ follows by symmetry.

Finally, we assume that $5 \leq i \leq t-4$. Let $W := W_{i-3} \cup W_{i-2} \cup W_{i-1} \cup W_i$. Note that W is convex, since $2 \leq i-3$ and $i < t$. By Lemma 4.18(ii), we have $w_{i-3} + w_{i-2} + w_{i-1} + w_i \leq 3$ and $w_i + w_{i+1} + w_{i+2} + w_{i+3} \leq 3$. By Claim 4, $w_k \leq 2$ for all k with $1 \leq k \leq i-4$. Thus, by Corollary 4.13, $\sum_{k=1}^{i-4} w_k \leq i-3$. Similarly, we have $\sum_{k=i+4}^{t} w_k \leq t-i-2$. Altogether, we obtain that

$$|B| = \sum_{k=1}^{i-4} w_k + \sum_{k=i-3}^{i-1} w_k + w_i + \sum_{k=i+1}^{i+3} w_k + \sum_{k=i+4}^{t} w_k \leq (i-3) + 3 + (t-i-2) = t-2 \leq |A| - 3.$$

□

4.5.5. Finalizing the proof of Theorem 4.2

We are now ready to prove Theorem 4.2. Namely, we show that for every ℓ-divided set $P = A \cup B$ with $|A|, |B| \geq 5$ and with neither A nor B in convex position there is an ℓ-divided 5-hole in P.

Suppose for the sake of contradiction that there is no ℓ-divided 5-hole in P. By the result of Harborth [115], every set P of ten points contains a 5-hole in P. In the case $|A|, |B| = 5$, the statement then follows from the assumption that neither of A and B is in convex position.

So assume that at least one of the sets A and B has at least six points. We obtain an island Q of P by iteratively removing extremal points so that neither part is in convex position after the removal and until one of the following conditions holds.

(i) One of the parts $Q \cap A$ and $Q \cap B$ has only five points.

(ii) Q is an ℓ-critical island of P with $|Q \cap A|, |Q \cap B| \geq 6$.

If (i) holds, we have $|Q \cap A| = 5$ or $|Q \cap B| = 5$. If $|Q \cap A| = 5$ and $|Q \cap B| \geq 6$, then we let Q' be the union of $Q \cap A$ with the six leftmost points of $Q \cap B$. Since $Q \cap A$ is not in convex position, Lemma 4.14 implies that there is an ℓ-divided 5-hole in Q', which is also an ℓ-divided 5-hole in Q, since Q' is an island of Q. However, this is impossible as then there is an ℓ-divided 5-hole in P because Q is an island of P. If $|Q \cap A| \geq 6$ and $|Q \cap B| = 5$, then we proceed analogously.

Otherwise, if (ii) holds, we have $|Q \cap A|, |Q \cap B| \geq 6$. There is no ℓ-divided 5-hole in Q, since Q is an island of P. By Lemma 4.19(i), we can assume without loss of generality that

$|A \cap \partial \operatorname{conv}(Q)| = 2$. Then it follows from Proposition 4.21 that $|Q \cap B| < |Q \cap A|$. By exchanging the roles of $Q \cap A$ and $Q \cap B$ and by applying Proposition 4.22, we obtain that $|Q \cap A| \leq |Q \cap B|$, a contradiction. This finishes the proof of Theorem 4.2.

4.6. Final remarks

At a first glance, it might seem that a similar approach could be used to derive stronger lower bounds also on the minimum number of 6-holes $h_6(n)$. However, since there are point sets of 29 points with no 6-hole [155], one would need to investigate point sets of size at least 30 in order to find an ℓ-divided 6-hole. This task is too demanding for our implementations, since the number of combinatorially different point sets grows too rapidly. Moreover, the case analysis in several steps of our proof would become much more complicated.

4.6.1. Necessity of the assumptions in Theorem 4.2

In the statement of Theorem 4.2 we require that the ℓ-divided set $P = A \cup B$ satisfies $|A|, |B| \geq 5$. We now show that those requirements are necessary in order to guarantee an ℓ-divided 5-hole in P by constructing an arbitrarily large ℓ-critical set $C = A \cup B$ with $|A| = 4$ and with no ℓ-divided 5-hole in C.

Proposition 4.23. *For every integer $n \geq 5$, there exists an ℓ-critical set $C = A \cup B$ with $|A| = 4$, $|B| = n$, and with no ℓ-divided 5-hole in C.*

Proof. First, we consider the case where n is odd. Let $p^+ = (0, 1)$ and $p^- = (0, -1)$ be two auxiliary points and let $\ell^+ = \{(x, y) \in \mathbb{R}^2 : y = x/4\}$ and $\ell^- = \{(x, y) \in \mathbb{R}^2 : y = -x/4\}$ be two auxiliary lines. We place the point $b'_1 = (2, -1/2)$ on the line ℓ^- and the auxiliary point $q = (2, 1/2)$ on the line ℓ^+. For $i = 2, \ldots, n$, we iteratively let b'_i be the intersection of the line ℓ^+ with the segment $p^+ b'_{i-1}$ if i is even and the intersection of ℓ^- with $p^- b'_{i-1}$ if i is odd. We place two points a_1 and a_2 sufficiently close to p^+ so that a_1 is above a_2, the segment $a_1 a_2$ is vertical with the midpoint p^+, and all non-collinear triples (b'_i, b'_j, p^+) have the same orientation as (b'_i, b'_j, a_1) and (b'_i, b'_j, a_2). Similarly, we place two points a_3 and a_4 sufficiently close to p^- so that a_3 is to the left of a_4, the segment $a_3 a_4$ lies on the line $\overline{p^- q}$ and has p^- as its midpoint, the point a_4 is to the left of b'_n, and all non-collinear triples (b'_i, b'_j, p^-) have the same orientation as (b'_i, b'_j, a_3) and (b'_i, b'_j, a_4). Figure 4.13 gives an illustration.

We let A, B', and B'_3 be the sets $\{a_1, a_2, a_3, a_4\}$, $\{b'_1, \ldots, b'_n\}$, and $B' \setminus \{b'_3\}$, respectively. Note that the line $\overline{a_3 a_4}$ intersects the segment $b'_1 b'_3$. Since $\max_{a \in A} x(a) < \min_{b' \in B'} x(b')$, the sets A and B' are separated by a vertical line ℓ.

Next we slightly perturb b'_3 to obtain a point b_3 such that b_3 lies above ℓ^- and all non-collinear triples (b_3, c, d) with $c, d \in A \cup B'_3$ have the same orientation as (b'_3, c, d). Note that the point b_3 lies in the interior of $\operatorname{conv}(B'_3)$, since $n \geq 5$.

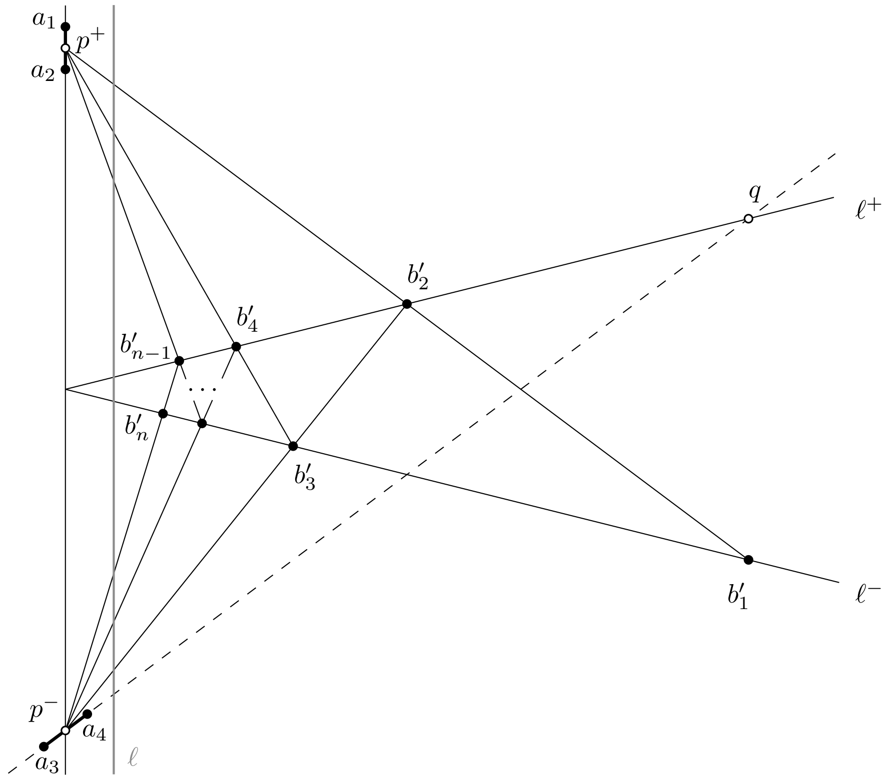

Figure 4.13.: The set C constructed in the proof of Proposition 4.23 for n odd.

To ensure general position, we transform every point $b'_i = (x, y) \in B'_3 \cap \ell^+$ to $b_i = (x, y - \varepsilon x^2)$ and every point $b'_i = (x, y) \in B'_3 \cap \ell^-$ to $b_i = (x, y + \varepsilon x^2)$ for some $\varepsilon > 0$. The remaining points in $A \cup \{b_3\}$ remain unchanged. We choose ε sufficiently small so that all non-collinear triples of points from $A \cup B'_3 \cup \{b_3\}$ have the same orientations as their images after the perturbation. Finally, let B be the set $\{b_1, \ldots, b_n\}$ and set $B_3 := B \setminus \{b_3\}$.

Since the points from B_3 lie on two parabolas, the set B is in general position. In particular, points from B_3 are in convex position and the point b_3 lies inside $\operatorname{conv}(B_3)$. Also observe that the line ℓ separates A and B and that a_1, a_3, and b_1 are the extremal points of $C := A \cup B$. Since neither of the sets A and B is in convex position, and removal of any of the extremal points a_1, a_3, b_1 leaves either A or B in convex position, the set $C = A \cup B$ is ℓ-critical.

We now show that C contains no ℓ-divided 5-hole. Suppose for contradiction that there is an ℓ-divided 5-hole H in C. We set $A^+ := \{a_1, a_2\}$, $A^- := \{a_3, a_4\}$, $B^+ := \{b_2, b_4, \ldots, b_{n-1}\}$, and $B^- := \{b_1, b_3, \ldots, b_n\}$. First we assume that H contains points from both A^+ and A^-. Then $H \cap B \subseteq \{b_{n-1}, b_n\}$, since if there is a point b_i in H with $i < n - 1$, then b_n lies in the interior of $\operatorname{conv}(H)$. Note that if $H \cap B = \{b_{n-1}, b_n\}$, then neither a_4 nor a_1 lies in H and thus $|H| < 5$. Hence $|H \cap B| = 1$, which is again impossible, as H cannot contain all points from A. Therefore we either have $H \cap A \subseteq A^+$ or $H \cap A \subseteq A^-$ and, in particular, $1 \le |H \cap A| \le 2$. We now distinguish these two cases.

Case 1: $|H \cap A| = 2$. If $H \cap A = A^+$, then the hole H can contain only the point b_n from B^-. This is because if there is a point b_i in $H \cap B^-$ with $i < n$, then the point b_{i+1} lies in the interior of $\mathrm{conv}(H)$. Additionally, H contains at most two points from B^+, since otherwise H is not in convex position. Consequently, b_n lies in H and $|H \cap B^+| = 2$, which is impossible, as H would not be in convex position.

If $H \cap A = A^-$, then the hole H contains no point from B^+. This is because if there is a point b_i in $H \cap B^+$, then the point b_{i+1} lies in the interior of $\mathrm{conv}(H)$. The point b_1 cannot lie in H because otherwise H is not in convex position as the line $\overline{a_3a_4}$ separates b_1 from $B \setminus \{b_1\}$. Additionally, H contains at most two points from B^-, since otherwise H is not in convex position. Thus H contains at most four points of C, which is impossible.

Case 2: $|H \cap A| = 1$. Assume first that $H \cap A \subseteq A^+$. Note that for $b_i, b_j \in B^-$ with $i < j \leq n$, the point b_{i+1} lies inside the triangle $\triangle(a_1, b_i, b_j)$ and, if $j < n$, the point b_{j+1} lies inside $\triangle(a_2, b_i, b_j)$. Thus H contains at most one point from B^- or we have $H \cap B^- = \{b_{n-2}, b_n\}$ and $H \cap A = \{a_2\}$. The latter case does not occur, since for every $b_i \in B^+$ with $i < n-1$ the point b_{n-1} lies in the interior of $\mathrm{conv}(\{a_2, b_i, b_{n-2}, b_n\})$. Therefore we consider the case $|H \cap B^-| \leq 1$. However, $|H \cap B^+| \geq 3$ is impossible since H would not be in convex position. Altogether, we obtain $|H| < 5$, which is impossible.

Now we assume that $H \cap A \subseteq A^-$. Note that for $b_i, b_j \in B^+$ with $i < j < n$, the point b_{i+1} lies inside the triangle $\triangle(a_4, b_i, b_j)$ and the point b_{j+1} lies inside $\triangle(a_3, b_i, b_j)$. Thus H contains at most one point from B^+. Consequently, H contains at least three points from B^-, which is possible only if $H \cap B^- = \{b_1, b_3, b_5\}$. However, then H contains a point b_i from B^+ and b_3 lies in the interior of $\mathrm{conv}(H)$.

Thus, in both cases H is not an ℓ-divided 5-hole in C, a contradiction.

To finish the proof, we consider the case in which n is even. Let $\tilde{C} = A \cup \tilde{B}$ be the set constructed above with $|A| = 4$ and $|\tilde{B}| = n + 1$. We set $B := \tilde{B} \setminus \{b_2\}$ and $C := A \cup B$. Note that C is ℓ-critical.

It remains to show that C contains no ℓ-divided 5-hole. Suppose for contradiction that there is an ℓ-divided 5-hole H in C. There is no ℓ-divided 5-hole in $\tilde{C}$ and thus b_2 lies in the interior of $\mathrm{conv}(H)$. Since b_1 is the only point from C to the right of b_2, the point b_1 lies in H. Since a_1 is the only point of C to the left of $\overline{b_2b_1}$, all other points of H lie to the right of $\overline{b_2b_1}$. Then, however, the set $(H \setminus \{a_1\}) \cup \{b_2\}$ is a 5-hole in $\tilde{C}$, which gives a contradiction. $\square$

4.6.2. Necessity of the assumptions in Lemmas 4.14 to 4.17

We remark that all the assumptions in the statements of Lemmas 4.14 to 4.17 are necessary; Figure 4.14(a) shows that the conditions $|B| = 5$ in Lemma 4.16 and the convexity of A in Lemma 4.17 are both necessary. The horizontal reflection of Figure 4.14(a) also shows the necessity of the assumption $|A| = 5$ in Lemma 4.14. It follows from the example in Figure 4.14(b) that the condition $|B| = 4$ cannot be omitted in Lemma 4.17, since there is an a-wedge with three points of B. The same point set without the point a' shows that the assumption $|B| \geq 4$ in Lemma 4.15 is necessary. The example from Figure 4.14(c) shows that the conditions $|B| = 6$

in Lemma 4.14, the convex position of A in Lemma 4.15, and $|A| = 6$ in Lemma 4.16 are all necessary. The same set without the point a shows that $|A| = 5$ in Lemma 4.15 is also needed and, if we remove the points a and a', then the resulting point set shows that we need $5 \leq |A|$ in Lemma 4.17. We can make statements only about convex a-wedges in Lemmas 4.15 and 4.16, as there are counterexamples for the corresponding statements without the convexity condition. It suffices to consider so-called *double-chains*, which are point sets obtained by placing n points on each of the two branches of a hyperbola. Double-chains also show that A cannot be in convex position in Lemma 4.14 and that the non-convex a-wedge must be empty of points in B in Lemma 4.17.

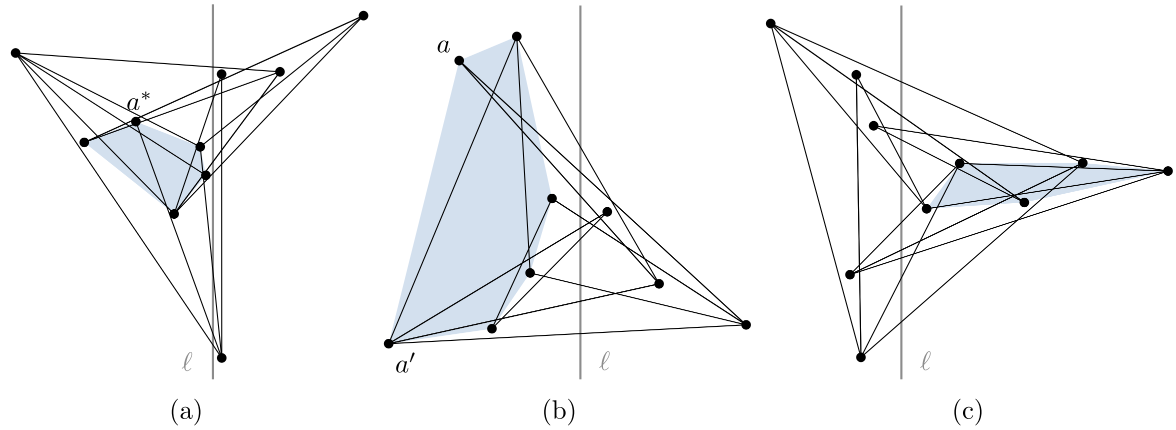

Figure 4.14.: Examples of points sets that witness tightness of Lemmas 4.14 to 4.17. All k-holes in these sets with $k \geq 5$ are shaded in blue. The set of edges drawn to visualize the order type is the set of exit edges defined in Chapter 3.

4.7. Flow summary

In Figure 4.15 we present the flow summary of the results presented.

4.8. Chapter summary

In this chapter we have presented the first superlinear lower bound on the number of 5-holes in any set of n points in the plane in general position. The main ingredient in the proof of this bound is the following structural result, which might be of independent interest: If a finite set P of points in the plane in general position is partitioned by a line ℓ into two subsets, each of size at least five then either one side is in convex position or ℓ intersects the convex hull of a 5-hole in P. The proof of this result integrates the assistance of the computer for dealing with point sets of small cardinality.

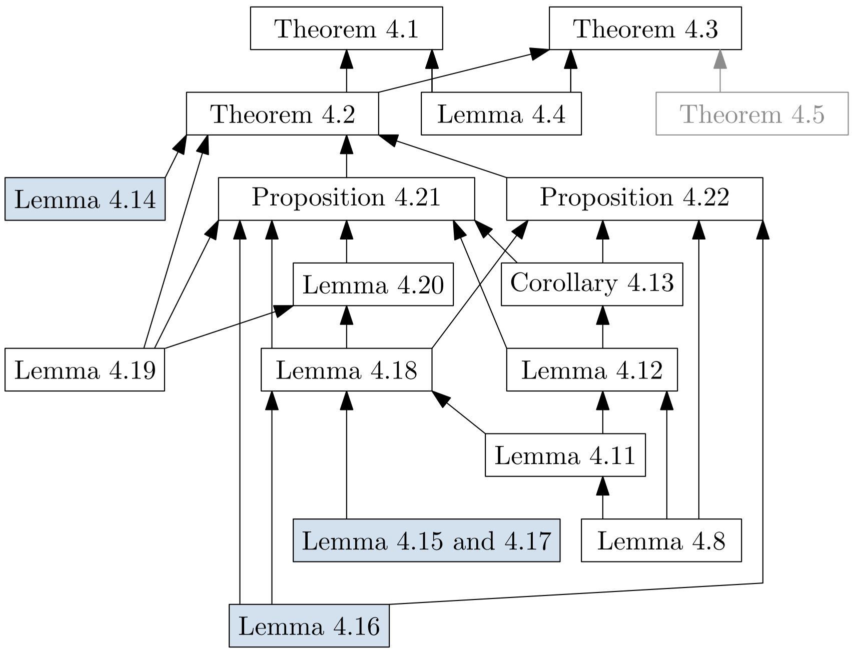

Figure 4.15.: Flow summary. The shaded boxes correspond to computer-assisted results and the result in gray is part of a paper by García [99].

Part II.

Topological drawings of graphs

5. Extending simple drawings

Part of the results in Sections 5.2, 5.3, 5.4 and 5.6 have been accepted for publication [35]. Very recently, we obtained the results presented in Sections 5.5 and 5.7. They are part of a preprint [37].

5.1. Introduction

A *simple drawing* of a graph G (also known as *good drawing* or as *simple topological graph* in the literature) is a drawing $D(G)$ of G in the plane such that every pair of edges share at most one point that is either a proper crossing (no tangent edges allowed) or an endpoint. Moreover, no three edges intersect in the same point and edges must neither self-intersect nor contain other vertices than their endpoints. Simple drawings, despite often considered in the study of crossing numbers, have basic aspects that are yet unknown.

The long-standing conjectures on the crossing numbers of K_n and $K_{n,m}$, known as the Harary-Hill and Zarankiewicz's conjectures, respectively, have drawn particular interest in the study of simple drawings of complete and complete bipartite graphs. The intensive study of these conjectures has produced deep results about simple drawings of K_n [138, 158] and $K_{n,m}$ [61].

In contrast to our knowledge about K_n, little is known about simple drawings of general graphs. In [140] it was observed that, when studying simple drawings of general graphs, it is natural to try to extend them, by inserting the missing edges between non-adjacent vertices. One of the main results in this chapter suggests that there is no hope for efficiently deciding when such operation can be performed.

The complement $\overline{G}$ of a graph G is the graph with the same vertex set as G and where two distinct vertices are adjacent if and only if they are not adjacent in G. Given a simple drawing $D(G)$ of a graph $G = (V, E)$ and a subset M of *candidate edges* from $\overline{G}$, an *extension* of $D(G)$ with M is a simple drawing $D'(G^+)$ of the graph $G^+ = (V, E \cup M)$ that contains $D(G)$ as a subdrawing. If such an extension exists, then we say that M can be *inserted* into $D(G)$.

Given a simple drawing, an extension with one given edge is not always possible, as shown by Kynčl [136]: In Figure 5.1(a) the edge uv cannot be inserted, because uv would cross an edge incident either to u or to v. We can extend this example to a simple drawing of $K_{2,4}$; see Figure 5.1(b). Moreover, we can use it to construct drawings of $K_{m,n}$ with larger values of m and n in which an edge uv cannot be inserted; see Figure 5.1(d). Kynčl's drawing can also be extended to a simple drawing of K_6 minus one edge where the only missing edge cannot be

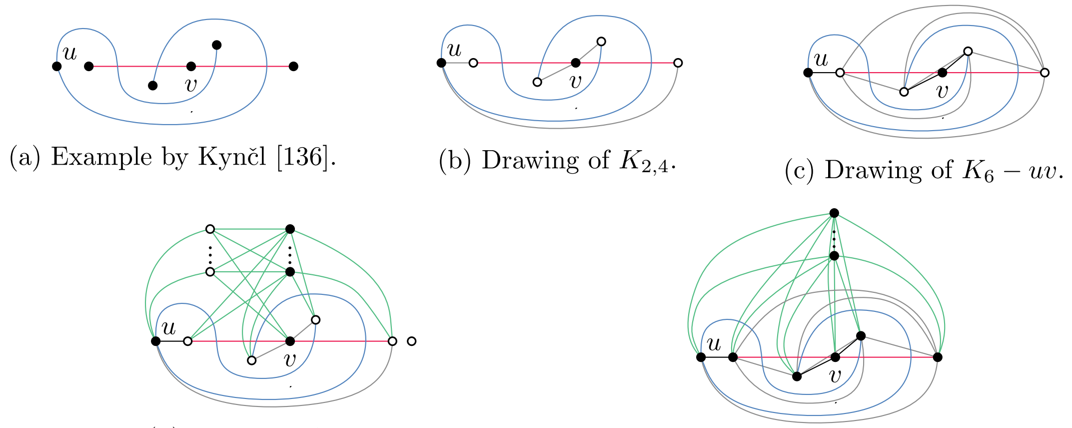

(a) Example by Kynčl [136]. (b) Drawing of $K_{2,4}$. (c) Drawing of $K_6 - uv$.

(d) Drawing of $K_{m,n}$. (e) Drawing of $K_n - uv$.

Figure 5.1.: Drawings in which the edge uv cannot be inserted.

inserted; see Figure 5.1(c). From this drawing one can construct drawings of K_n with $n \geq 6$ minus one edge where the only missing edge cannot be inserted; see Figure 5.1(e).

Extensions, by inserting both vertices and edges, have received a great deal of attention in the last decade, especially for (different classes of) plane drawings [33, 40, 57, 72, 123, 146, 167]. It has also been of interest to study crossing number questions on planar graphs with one additional edge [59, 109, 176]. We remark that the term *augmentation* has also been used in the literature for the similar problem of inserting edges and/or vertices to a graph [86]. Extensions of simple drawings have been previously considered in the context of *saturated* drawings, that is, drawings in which no edge can be inserted [112, 140]. In that context, the main interest lies in finding the minimum number of edges that a saturated drawing on n vertices can have. This minimum was shown to be at most $17.5n$ [140] and later this bound was improved to $7n$ [112].

In this chapter we study the computational complexity of extending a simple drawing $D(G)$ of a graph G. Note that if $D(G)$ is a straight-line drawing the problem is trivial. Pseudolinear drawings are those in which the m curves represented the edges can be extended to an arrangement of m pseudolines, with each edge lying in a distinct pseudoline. Levi's enlargement lemma states that given a pseudoline arrangement and two points u, v that do not lie on the same pseudoline, then we can insert a new pseudoline into the arrangement that contains both u and v. Thus, we can always (iteratively) insert any set of edges into a pseudolinear drawing.

Outline. In Section 5.2, we show that deciding whether a simple drawing $D(G)$ can be extended with a set M of candidate edges is NP-complete. Moreover, in Section 5.3, we prove that finding the largest subset of edges from M that extend $D(G)$ is APX-hard. In the rest of this chapter we focus on the problem of deciding whether an edge can be inserted into a simple drawing. In Section 5.4 we reformulate this problem in the dual graph of the planarization of the drawing. In Section 5.5 we present our new result, that improves the ones in Sections 5.2 and 5.3. It shows that, surprisingly, deciding whether one candidate edge can be inserted into

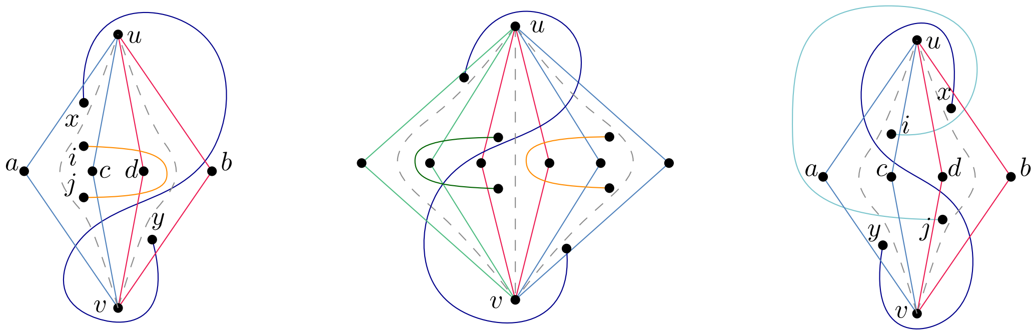

(a) Variable gadget $\mathcal{X}$. (b) Clause gadget $\mathcal{C}$. (c) Wire gadget $\mathcal{W}$.

Figure 5.2.: Basic gadgets for the proof of Theorem 5.1.

a simple drawing is NP-complete. In Section 5.6, we present a polynomial-time algorithm to decide whether an edge uv can be inserted into $D(G)$ when $\{u, v\}$ is a dominating set for G In Section 5.7 we show that that the problem of deciding whether an edge can be inserted into $D(G)$ is FPT with respect to the number of crossings of $D(G)$. Finally, we present a summary and conclusions in Section 5.8.

5.2. Inserting a given set of edges is NP-complete

In this section we prove the following result:

Theorem 5.1. *Given a simple drawing $D(G)$ of a graph $G = (V, E)$ and a set M of edges of the complement of G, it is* NP*-complete to decide whether $D(G)$ can be extended with the set M.*

Notice first that the problem is in NP, since it can be described combinatorially. Our proof of Theorem 5.1 is based on a reduction from *monotone* 3SAT [52]. An instance of that problem consists of a Boolean formula ϕ in 3-CNF with a set of variables $X = \{x_1, \ldots, x_n\}$ and a set of clauses $K = \{C_1, \ldots, C_m\}$. An occurrence of a variable in a clause is called a *literal*. Monotonicity means that in each clause either all the literals are positive (*positive clause*) or they are all negative (*negative clause*). The *bipartite graph $G(\phi)$ associated to ϕ* is the graph with vertex set $X \cup K$ and where a variable x_i is adjacent to a clause C_j if and only if $x_i \in C_j$ or $\overline{x_i} \in C_j$.

We now show how to construct a simple drawing from a given formula. We start by introducing our three basic gadgets, the *variable gadget*, the *clause gadget*, and the *wire gadget*, shown in Figure 5.2. The variable gadget contains two nested cycles, $avbu$ on the outside and $cvdu$ on the inside, drawn in the plane without any crossings. Two additional vertices x and y are drawn in the interior of $avcu$ and $dvbu$, respectively. They are connected with an edge that, starting in x, crosses the edges au, ub, dv, cv, av, and vb, in this order, and ends in y. Another two vertices i and j are drawn inside the region in the interior of $avcu$ that is incident to x. They are connected with an edge that, starting in i, crosses the edges uc, ud, vd, and vc, in this order, and ends in j; see Figure 5.2(a). Notice that the edge uv can be inserted only in two possible regions: either inside the cycle $avcu$ or inside the cycle $dvbu$. Drawing the edge uv in any other region would force it to cross uj or xy more than once. The clause gadget and the wire gadget are similarly defined; see Figure 5.2(b)–(c).

In each of these three gadgets shown in Figure 5.2, the edge uv can only be inserted in the regions where the dashed arcs are drawn. In the rest of the section, when we refer to the *regions* in a gadget we mean these regions where the edge uv can be inserted.

In a variable gadget, these regions encode the truth assignment of the corresponding variable x_i: Inserting the edge uv in the left region corresponds to the assignment $x_i = \texttt{true}$, while inserting it in the right region corresponds to $x_i = \texttt{false}$. We call these left and right regions in a variable gadget the `true` and `false` regions, respectively. In a clause gadget, each of the three regions is associated to a literal in the corresponding clause. Wire gadgets propagate the truth assignment of the variables to the clauses. They are drawn between the gadgets corresponding to clauses and variables that are incident in $G(\phi)$. The idea is that if an assignment makes a literal not satisfy a clause, then the edge uv in the wire gadget blocks the region in the clause gadget corresponding to that literal by forcing uv to cross that region twice.

Let $w^{(\mathcal{G})}$ denote vertex w in gadget $\mathcal{G}$. The following lemma shows that we can get the desired behavior with a wire gadget connecting a variable gadget and a clause gadget. The precise placement of a wire gadget with respect to the variable gadget and the clause gadget that it connects is illustrated in Figure 5.3.

Lemma 5.2. *We can combine a variable gadget $\mathcal{X}$, a clause gadget $\mathcal{C}$, and a wire gadget $\mathcal{W}$ to produce a simple drawing with the following properties.*

- *If $u^{(\mathcal{X})}v^{(\mathcal{X})}$ is inserted in the `false` region in $\mathcal{X}$, then inserting $u^{(\mathcal{W})}v^{(\mathcal{W})}$ prevents $u^{(\mathcal{C})}v^{(\mathcal{C})}$ from being inserted in one specified target region in $\mathcal{C}$.*
- *If $u^{(\mathcal{X})}v^{(\mathcal{X})}$ is inserted in the `true` region in $\mathcal{X}$, then we can insert $u^{(\mathcal{W})}v^{(\mathcal{W})}$ in a way such that $u^{(\mathcal{C})}v^{(\mathcal{C})}$ can then be inserted in any region in $\mathcal{C}$.*

Proof. We start with a drawing of the variable gadget $\mathcal{X}$ and the clause gadget $\mathcal{C}$ such that the two gadgets are drawn on a line and they are disjoint. A representation of how the wire gadget is then inserted is shown in Figure 5.3. In this proof we focus on the wire gadget drawn with blue edges and vertices.

In Figure 5.3, gadget $\mathcal{X}$ lies to the left of gadget $\mathcal{C}$. The `true` and `false` regions in $\mathcal{X}$ are shaded in green and red, respectively. We assume that the target region in $\mathcal{C}$ is the leftmost one, shaded in yellow. The left and right regions in the wire gadget are shaded in red and yellow, respectively.

If the edge $u^{(\mathcal{X})}v^{(\mathcal{X})}$ is inserted in the `false` region in $\mathcal{X}$ then the edge $u^{(\mathcal{W})}v^{(\mathcal{W})}$ cannot be inserted in the yellow region in $\mathcal{W}$, since it would cross $u^{(\mathcal{X})}v^{(\mathcal{X})}$ twice. Thus, $u^{(\mathcal{W})}v^{(\mathcal{W})}$ can only be inserted in the red region in $\mathcal{W}$. If inserted in that region, $u^{(\mathcal{C})}v^{(\mathcal{C})}$ cannot be inserted in the yellow region in $\mathcal{C}$, since it would cross $u^{(\mathcal{W})}v^{(\mathcal{W})}$ twice. In contrast, if the edge $u^{(\mathcal{X})}v^{(\mathcal{X})}$ is inserted in the `true` (green) region in $\mathcal{X}$, then $u^{(\mathcal{W})}v^{(\mathcal{W})}$ can be inserted in either of the two

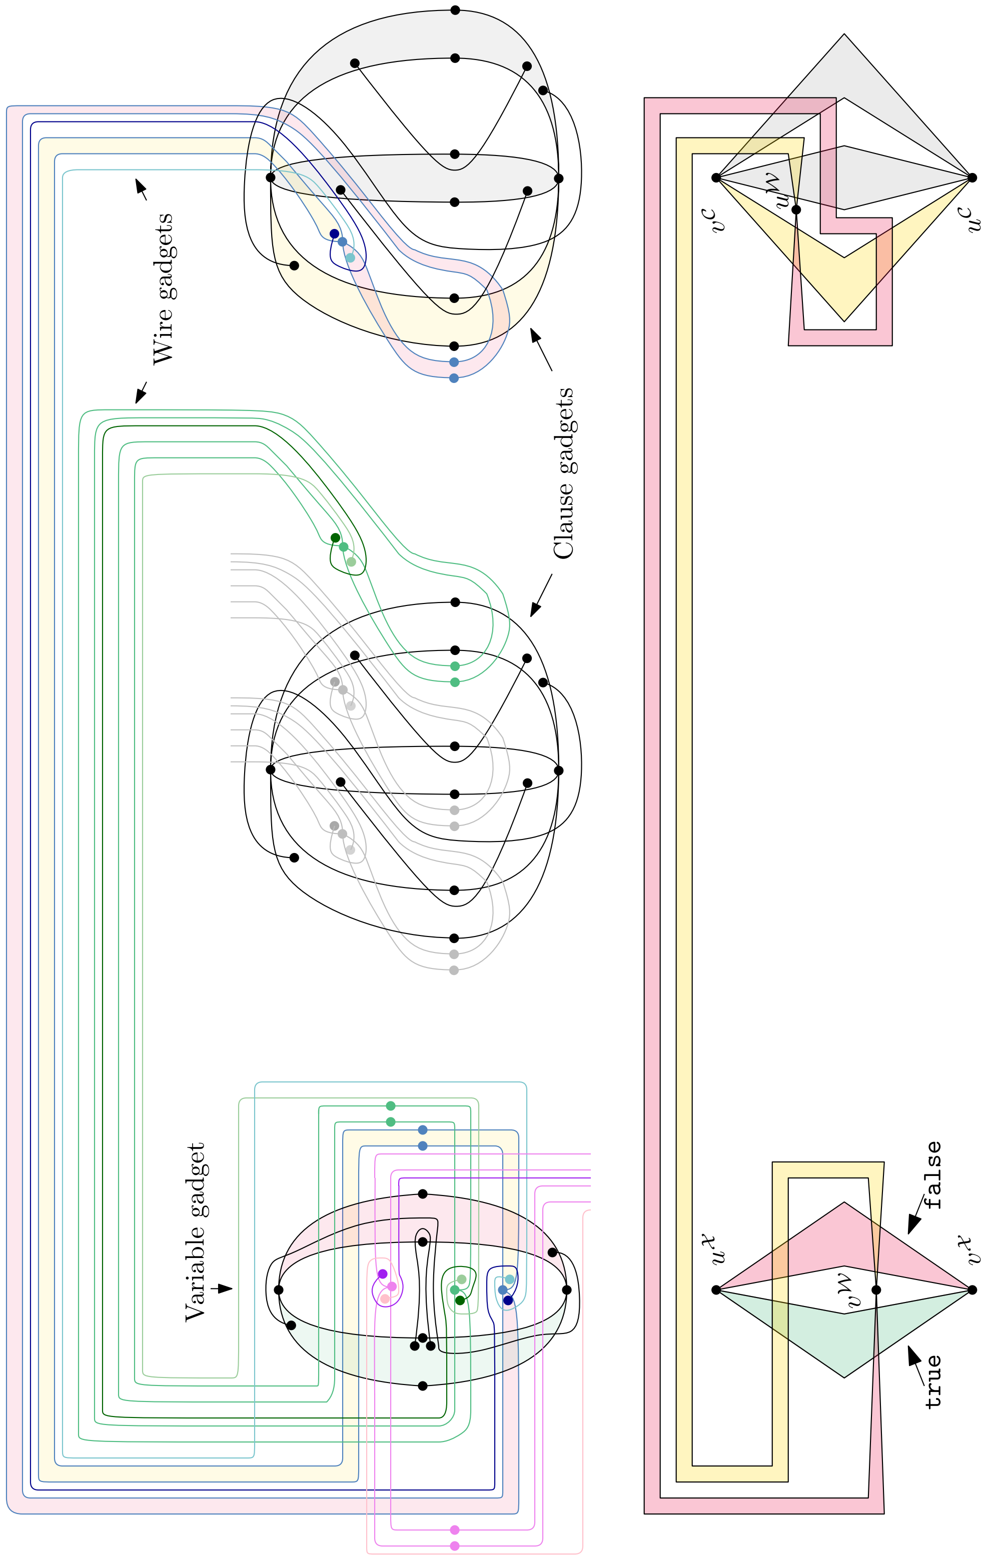

Figure 5.3.: Reduction from monotone 3SAT.

regions in $\mathcal{W}$. In particular, it can be inserted in the yellow region in a way such that $u^{(\mathcal{C})}v^{(\mathcal{C})}$ can then be inserted in any region in $\mathcal{C}$.

Finally, notice that if the target region in $\mathcal{C}$ is not the leftmost one, we can adapt the construction by leaving the region(s) to the left in $\mathcal{C}$ uncrossed by the wire gadget $\mathcal{W}$; see the clause gadget in the middle of Figure 5.3. □

Let ϕ be an instance of monotone 3SATand let $G(\phi)$ be the bipartite graph associated to ϕ. Let $D(\phi)$ be a 2-page book drawing of $G(\phi)$ in which (i) all vertices lie on an horizontal line, and from left to right, first the ones corresponding to negative clauses, then to variables, and finally to positive clauses; and (ii) the edges incident to vertices corresponding to positive clauses are drawn as circular arcs above that horizontal line, while the ones incident to vertices corresponding to negative clauses are drawn as circular arcs below it. In an slight abuse of notation, we refer to the vertices in $D(\phi)$ corresponding to variables and clauses simply as variables and clauses, respectively.

We construct a simple drawing D' from $D(\phi)$ by first replacing the variables and clauses by variable gadgets and clause gadgets, respectively, and drawn in disjoint regions. Moreover, the clause gadgets corresponding to negative clauses are rotated 180°. We then insert the wire gadgets. The edges in $D(\phi)$ connecting variables to positive clauses are replaced by wire gadgets drawn as in the proof of Lemma 5.2; see Figure 5.3. Similarly, the edges in $D(\phi)$ connecting variables to negative clauses are replaced by wire gadgets drawn as the ones before, but rotated 180°.

We now describe how to draw the wire gadgets with respect to each other, so that the result is a simple drawing; see Figure 5.3 for a detailed illustration. First, we focus on the drawing locally around the variable gadgets. Consider a set of edges in $D(\phi)$ connecting a variable with some positive clauses. The drawing $D(\phi)$ defines a clockwise order of these edges around the common vertex starting from the horizontal line. We insert the corresponding wire gadgets locally around the variable gadget following this order. Each new gadget is inserted shifted up and to the right with respect to the previous one (as the blue and green gadgets depicted in Figure 5.3). Edges in $D(\phi)$ connecting a variable with some negative clauses are replaced by wire gadgets in an analogous manner with a 180° rotation. We assign the three different regions in a clause gadget to the target regions in the wire gadgets following the rotation of the edges around the clause in $D(\phi)$. (Note that we can assume without loss of generality, by possibly duplicating variables, that each clause in ϕ contains three literals.) Thus, locally around a clause gadget, it is then possible to draw the different wire gadgets connecting to it without crossing. Since $D(\phi)$ is a 2-page book drawing, the constructed drawing D' is a simple drawing.

Let M be the set of uv edges of all the gadgets. The fact that ϕ is satisfiable if and only if M can be inserted into D' follows now from Lemma 5.2, finishing the proof of Theorem 5.1.

5.3. Maximizing the number of edges inserted is APX-hard

In this section we show that the maximization version of the problem of inserting missing edges from a prescribed set into a simple drawing is APX-hard. This implies that, if P $\neq$ NP, then no PTAS[1] We start by showing that this maximization problem is NP-hard.

Theorem 5.3. *Given a simple drawing $D(G)$ of a graph $G = (V, E)$ and a set M of edges in the complement $\overline{G}$, it is NP-hard to find a maximum subset of edges $M' \subseteq M$ that extends $D(G)$.*

Our proof of Theorem 5.3 is based on a reduction from the maximum independent set problem (MIS). By showing that the reduction when the input graph has vertex degree at most three is actually a PTAS-reduction we will then conclude that the problem is APX-hard.

An independent set of a graph $G = (V, E)$ is a set of vertices $S \subseteq V$ such that no two vertices in S are incident with the same edge. The problem of determining the maximum independent set (MIS) of a given graph is APX-hard even when the graph has vertex degree at most three [32]. We first describe the construction of a simple drawing $D'(G')$ from the graph G of a given MIS instance. Then we argue that for a well-selected set of edges M that are not present in $D'(G')$, finding a maximum subset $M' \subseteq M$ that can be inserted into $D'(G')$ is equivalent to finding a maximum independent set of G.

5.3.1. Constructing a drawing from a given graph

We begin by introducing our two basic gadgets, the vertex gadget $\mathcal{V}$ and the edge gadget $\mathcal{E}$, shown in Figure 5.4. They are reminiscent of the gadgets in the previous section, but adapted to this different reduction. Similarly as in the previous gadgets, there is only one region in which the edge uv can be inserted into $\mathcal{V}$ and only two regions in which the edge uv can be inserted into $\mathcal{E}$. These regions are the ones in which the dashed arcs in Figure 5.4(b) are drawn.

In Figure 5.4(c) we combined an edge gadget and two vertex gadgets. This figure shows a copy $\mathcal{E}^{(e)}$ of the gadget $\mathcal{E}$ (that corresponds to an edge $e = wz$) drawn over two different copies, $\mathcal{V}^{(w)}$ and $\mathcal{V}^{(z)}$, of the gadget $\mathcal{V}$ (that correspond to vertices w and z, respectively). We relabel the vertices in the copies of these gadgets by using the vertex or edge to which they correspond as their superscripts. Since there is only one region in which $v^{(w)}u^{(w)}$ and $v^{(z)}u^{(z)}$ can be drawn, inserting both of these edges prevents $v^{(e)}u^{(e)}$ from being inserted. Inserting either only $v^{(w)}u^{(w)}$ or only $v^{(z)}u^{(z)}$ leaves exactly one possible region where $v^{(e)}u^{(e)}$ can be inserted.

We have all the ingredients needed for our construction. Suppose that we are given a simple graph $G = (V, E)$. This graph admits a 1-page book drawing $D(G)$ in which the vertices are placed on a horizontal line and the edges are drawn as circular arcs in the upper halfplane. Since the edge gadget does not interlink the vertex gadgets symmetrically, we consider the edges in $D(G)$ with an orientation from their left endpoint to their right one.

[1]For any fixed $\varepsilon > 0$, a polynomial-time approximation scheme (PTAS) produces a solution to any given instance within an approximation factor of $1 + \varepsilon$ in polynomial time in the input size. However, the polynomial algorithm might be exponential in ε.

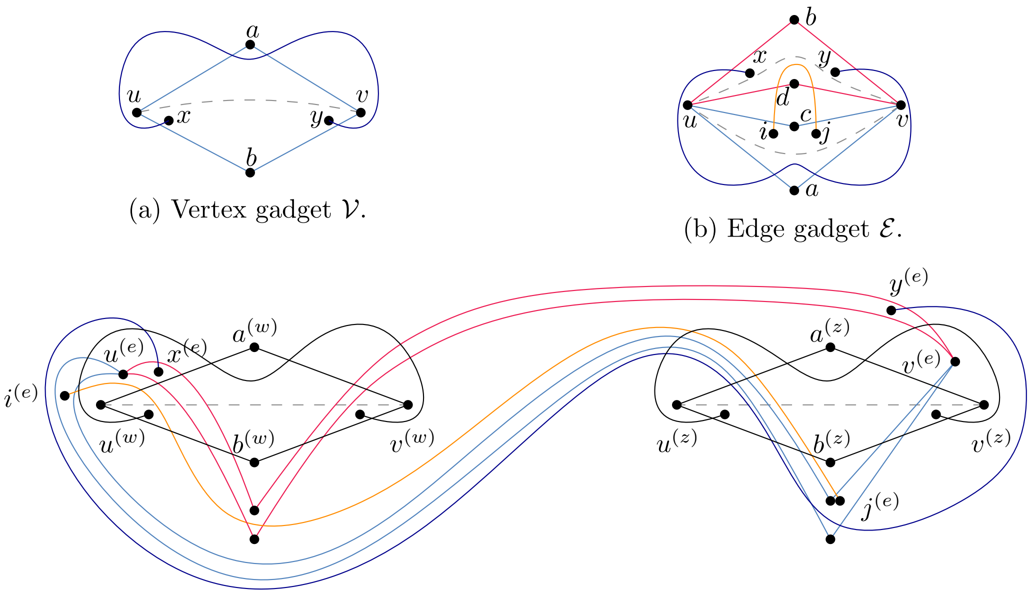

(a) Vertex gadget $\mathcal{V}$.

(b) Edge gadget $\mathcal{E}$.

(c) Two vertex gadgets interlinked by an edge gadget.

Figure 5.4.: Basic gadgets and drawings for the proof of Theorem 5.3.

The following lemma shows that is possible to replace each vertex $w \in V$ in the drawing by a vertex gadget $\mathcal{V}^{(w)}$ and each edge $e \in E$ by an edge gadget $\mathcal{E}^{(e)}$, and obtain simple drawing $D'(G')$ (where G' is the disjoint union of the underlying graphs of the vertex- and edge gadgets).

Lemma 5.4. *Given a 1-page book drawing $D(G)$ of a graph $G = (V, E)$, then we can replace every vertex by a vertex gadget and every edge by an edge gadget to obtain a simple drawing.*

Proof. We show that the copies $\{\mathcal{E}^{(e)} : e \in E\}$ can be inserted into $\bigcup_{w \in V} \mathcal{V}^{(w)}$ such that such that vertex gadgets corresponding to different vertices are drawn in disjoint regions and for every edge $e = wz \in E$, $\mathcal{V}^{(w)} \cup \mathcal{V}^{(z)} \cup \mathcal{E}^{(e)}$ is as in Figure 5.4(c) (up to interchanging the indices w and z), and such that the resulting drawing is simple.

First, for each vertex $w \in V$ we place the gadget $\mathcal{V}^{(w)}$ in its position, so all the copies of $\mathcal{V}$ lie (equidistant) on a horizontal line and do not cross each other. For the edges of G, since the drawing in Figure 5.4(c) is not symmetric, we choose an orientation. We orient all the edges in the 1-page book drawing $D(G)$ of G from left to right. We start by inserting the corresponding $\mathcal{E}$ gadgets from left to right and from the shortest edges in $D(G)$ to the longest. For an edge wz, the intersections of the gadget $\mathcal{E}^{(wz)}$: (i) with the edges $u^{(w)}a^{(w)}$ and $u^{(w)}b^{(w)}$ are placed to the left of all the previous intersections of other edge gadgets with that edge; (ii) with the edge $v^{(w)}b^{(w)}$ are placed to the right of all the previous intersections with that edge; (iii) with the edge $v^{(w)}a^{(w)}$ are placed to the right of previous intersections with gadgets $\mathcal{E}^{(wt)}$ and to the left of previous intersections with gadgets $\mathcal{E}^{(tw)}$; (iv) with the edges $u^{(z)}a^{(z)}$ and $u^{(z)}b^{(z)}$ are placed to the left of the previous intersections with gadgets $\mathcal{E}^{(tz)}$ (v) with the edge $v^{(z)}b^{(z)}$ are placed

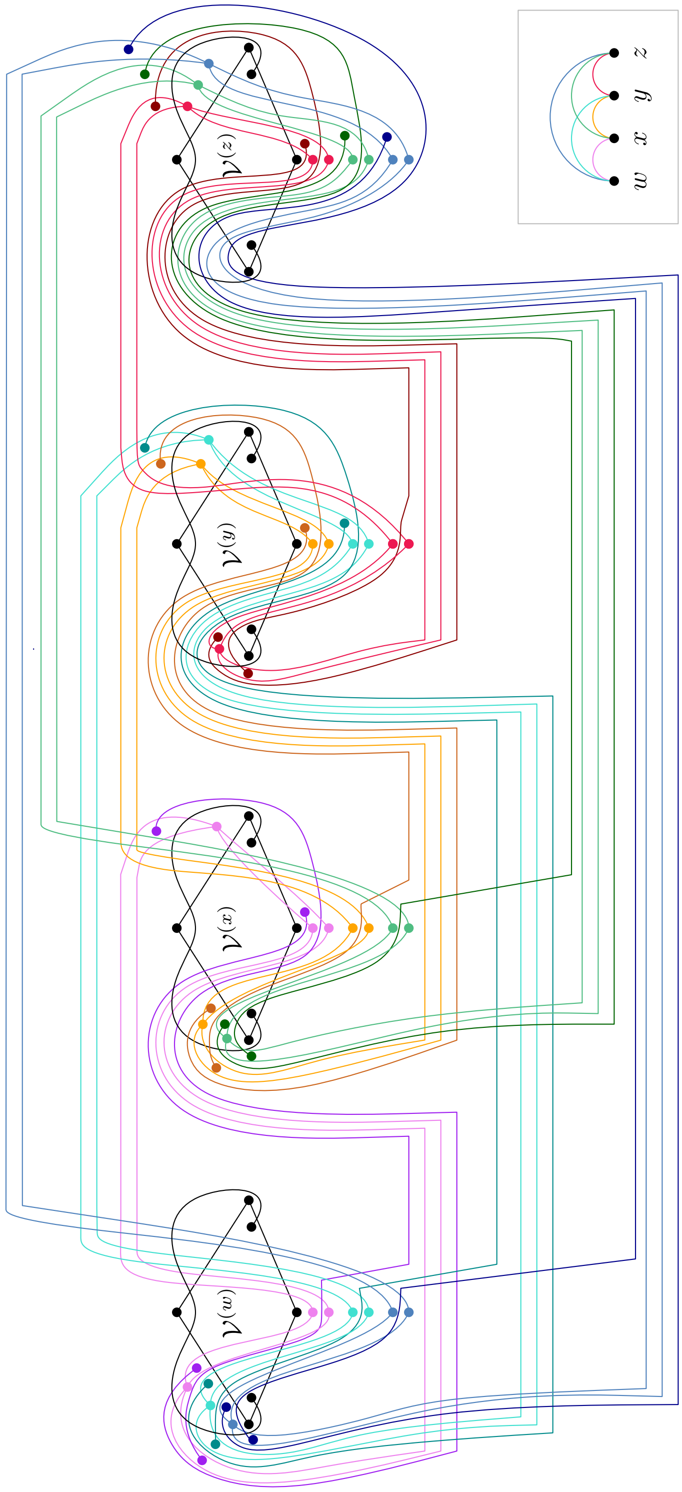

Figure 5.5.: Drawing obtained by a reduction from K_4.

to the left of all previous intersections; and (vi) with the edge $v^{(z)}a^{(z)}$ are placed to the left of all previous intersections with gadgets $\mathcal{E}^{(tz)}$; see Figure 5.5.

Moreover, the arcs of an edge gadget connecting two vertex gadgets are drawn either completely in the upper half-plane or completely in the lower one with respect to the horizontal line and two arcs cross at most twice. If they are part of edges in edge gadgets connected to the same vertex gadget, they might cross locally around this vertex gadget. However, after this crossing, they follow the circular-arc routing induced by $D(G)$ (or its mirror image) and do not cross again. Otherwise, with respect to each other, they follow the circular-arc routing induced by $D(G)$ (or its mirror image) and thus cross at most once; see Figure 5.5.

Since in none of the gadgets two incident edges cross, and edges of different gadgets are vertex-disjoint, we only have to worry about edges from different gadgets crossing more than once. By construction, no edge in an edge gadget intersects more than once with an edge in a vertex gadget. Thus, it remains to show that any two edges from two distinct edge gadgets cross at most once. Such two edges are included in a subgraph H of G with exactly four vertices. The drawing induced by the four vertex gadgets and the at most six edge gadgets is homeomorphic to a subdrawing of the drawing in Figure 5.5. It is routine to check that it is a simple drawing, and thus any two edges cross at most once. □

5.3.2. Reduction from maximum independent set

Proof of Theorem 5.3. Given a graph $G = (V, E)$, we reduce the problem of deciding whether G has an independent set of size k to the problem of deciding whether the simple drawing $D'(G')$ constructed as in Lemma 5.4 with a candidate set of edges $M = \{u^{(w)}v^{(w)} : w \in V\} \cup \{u^{(e)}v^{(e)} : e \in E\}$ can be extended with a set of edges $M' \subseteq M$ of cardinality $|M'| = |E| + k$.

To show the correctness of the (polynomial) reduction, we first show that if G has an independent set I of size k, then we can extend $D'(G')$ with a set M' of $|E| + k$ edges of M. Clearly, the k edges $\{u^{(w)}v^{(w)} : w \in I\}$ can be inserted into $D'(G')$ by the construction of the drawing. Since I is an independent set, each edge has at most one endpoint in I. Thus, in every edge gadget $\mathcal{E}^{(e)}$ at most one of the two possibilities for inserting the edge $u^{(e)}v^{(e)}$ is blocked by the previous k inserted edges. We therefore can also insert the $|E|$ edges $\{u^{(e)}v^{(e)} : e \in E\}$.

Conversely, let $M' \subset M$ be a set of $|E|+k$ edges can be inserted into $D'(G')$ and that contains the minimum number of uv edges from vertex gadgets. If the set of vertices $\{w \in V : u^{(w)}v^{(w)} \in M'\}$ is an independent set of G, then we are done, since at most $|E|$ edges of M' can be from edge gadgets, so at least k are from vertex gadgets. Otherwise, there are two edges $u^{(w)}v^{(w)}$ and $u^{(z)}v^{(z)}$ in M' such that the corresponding vertices $w, z \in V$ are connected by the edge $wz \in E$. By the construction of $D'(G')$ this implies that the edge $u^{(wz)}v^{(wz)}$ belongs to M, but it cannot be in M'. By removing the edge $u^{(w)}v^{(w)}$ and inserting the edge $u^{(wz)}v^{(wz)}$ into $D'(G')$, we obtain another valid extension with the same cardinality but one less uv edge from a vertex gadget. This contradicts our assumption. □

The presented reduction can be further analyzed to show that the problem is actually APX-hard. Note that the problem we are reducing from, maximum independent set in simple graphs, is

APX-hard [32] even in graphs with vertex degree at most three. Our reduction can be shown to be an L-reduction in that case, implying a PTAS-reduction. This shows the the problem is APX-hard and implies that, unless P = NP, there is no PTAS for the problem, and thus, there is a constant C such that finding an approximation with ratio better than C is as hard as finding the optimal solution.

Corollary 5.5. *Given a simple drawing $D(G)$ of a graph G and a set of edges M of the complement of G, finding the size of the largest subset of edges from M extending $D(G)$ is* APX-*hard.*

Proof. Since the MIS problem for graphs with vertex degree at most three is APX-hard [32], it suffices to show that the reduction proving Theorem 5.3 is an L-reduction. This type of reductions was introduced by Papadimitriou and Yannakakis [166]. In order to provide a formal definition, we present some notation.

Given an NP-optimization problem P, we denote by $I(P)$ the set of instances of P. For example, the set of all graphs is $I(\text{MIS})$. The NP-optimization problem P has associated an objective function $cost_P$ that we would like to either maximize or minimize (in our case maximize). For each instance $x \in I(P)$ we denote by $\text{opt}_P(x)$ the optimal value of a feasible solution with respect to $cost_P$. (For the MIS problem, the feasible solutions are the independent sets of the instance graph and *cost* measures the size of a set.)

Let A and B be a pair of NP-optimization problems. There is an *L-reduction* from A to B if there are polynomial-time computable functions f and g and positive constants c_1 and c_2 such that,

(i) f maps every instance $x \in I(A)$ to an instance $x' = f(x) \in I(B)$;

(ii) g maps every feasible solution y' of $x' = f(x)$ to a feasible solution $y = g(x, y')$ of $x \in I(A)$;

(iii) for every instance $x \in I(A)$, $\text{opt}_B(f(x)) \leq c_1 \cdot \text{opt}_A(x)$; and

(iv) for every instance $x \in I(A)$ and for every feasible solution y' of $x' = f(x)$, $|\text{opt}_A(x) - cost_A(y)| \leq c_2 \cdot |\text{opt}_B(x') - cost_B(y')|$, where $y = g(x, y')$.

Given a simple graph $G = (V, E)$, we construct a simple drawing $D'(G')$ as in Lemma 5.4. This construction plays the role of f in (i). We denote by M the candidate set of edges consisting of all the uv edges of the gadgets used to construct $D'(G')$, that is, $M = \{u^{(w)}v^{(w)} : w \in V\} \cup \{u^{(e)}v^{(e)} : e \in E\}$. Then, as argued in the proof of Theorem 5.3, G has an independent set of size k if and only if we can insert $|E| + k$ from M into $D'(G')$. Moreover, suppose that $M' \subseteq M$ is a subset of $|E| + k$ edges that can be inserted into $D'(G')$. Using the ideas of the proof of Theorem 5.3, if the set of vertices $\{w \in V : u^{(w)}v^{(w)} \in M'\}$ is an independent set of G, then they are an independent set of G of size k. Otherwise, there are two edges $u^{(w)}v^{(w)}$ and $u^{(z)}v^{(z)}$ in M' and then the edge $u^{(wz)}v^{(wz)}$ cannot be in M'. By removing the edge $u^{(w)}v^{(w)}$ and inserting the edge $u^{(wz)}v^{(wz)}$ into $D'(G')$, we obtain another set of candidate edges that can be inserted with the same cardinality but with one less uv edge from a vertex gadget. Iterating this process we obtain a subset of $|E|+k$ edges $M'' \subseteq M$ such that the set of k vertices $\{w \in V : u^{(w)}v^{(w)} \in M''\}$ is an independent set of G. This defines the function g mapping a feasible subset $M' \subseteq M$ of at least $|E|$ edges that we can insert into $D'(G')$ to an independent set in G of size $|M'| - |E|$. We extend g, so that every feasible subset $M' \subseteq M$ with $|M'| \leq |E|$ is mapped to the empty set. This proves (ii).

Let $\alpha = \alpha(G)$ be the size of the maximum independent set of G. We now show (iii). First, observe that the handshaking lemma and the fact that the vertex degrees in G are at most three imply $|E| \leq 3/2|V|$. We now bound $|V|$ in terms of $\alpha(G)$. Wei [198] and Caro [62] independently showed that $\alpha(G) \geq \sum_{v \in V} 1/(d(v) + 1)$, where $d(v)$ is the degree of vertex v. Thus, in our case $|V| \leq 4\alpha$. This bound also follows from Turán's theorem [191]; five nice proofs of this theorem can be found in [30]. Plugging this bound $|V| \leq 4\alpha$ into the equation obtained by the handshaking lemma we get $|E| \leq 3/2|V| \leq 6\alpha$. Since an optimal solution for the problem of inserting the largest subset of candidate edges into $D'(G')$ has size $\alpha + |E| \leq 7\alpha$, we have proven (iii) for a constant $c_1 = 7$.

Finally, we show (iv) for the constant $c_2 = 1$. Let $M' \subseteq M$ be a set of l edges that can be inserted into $D'(G')$. If $l \leq |E|$, then g maps M' to the empty set and we have that $\alpha - 0 \leq |E| + \alpha - l$. Otherwise, if $l = |E| + l'$ for $l' \geq 0$ we have that $\alpha - l' = |E| + \alpha - |E| - l'$. Thus, the absolute errors are in the worst case the same, as desired. □

5.4. Dual graph for inserting one edge

In this section, we consider the problem of extending a simple drawing of a graph by inserting exactly one edge uv for a given pair of non-adjacent vertices u and v. We start by rephrasing our problem as a problem of finding a certain path in the dual of the planarization of the drawing.

Given a simple drawing $D(G)$ of a graph $G = (V, E)$, the *dual graph* $G^*(D)$ has a vertex corresponding to each cell of $D(G)$ (where a cell is a component of $\mathbb{R}^2 \setminus D(G)$). There is an edge between two vertices if and only if the corresponding cells are separated by the same segment of an edge in $D(G)$. Notice that $G^*(D)$ can also be defined as the plane dual of the planarization of $D(G)$, where crossings are replaced by vertices so that the resulting drawing is plane.

We define a coloring χ of the edges of $G^*(D)$ by labeling the edges of the original graph G using numbers from 1 to $|E|$, and assigning to each edge of $G^*(D)$ the label of the edge that separates the cells corresponding to its incident vertices. Given two vertices $u, v \in V$, let $G^*(D, \{u, v\})$ be the subgraph of $G^*(D)$ obtained by removing the edges corresponding to connections between cells separated by an (arc of an) edge incident to u or to v, and let χ' be the coloring of the edges coinciding with χ in every edge. The problem of extending $D(G)$ with one edge uv is equivalent to the existence of a heterochromatic path in $G^*(D, \{u, v\})$ (i.e., no color is repeated) with respect to χ, between two vertices that corresponds to a cell incident to u and a cell incident to v, respectively.

We remark that, from this dual perspective, it is clear that the problem of deciding whether a simple drawing can be extended with a given set of edges is in NP.

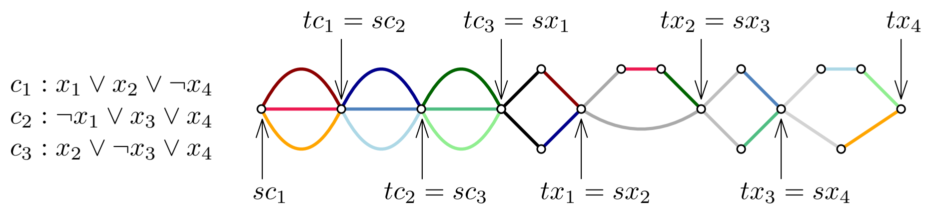

Figure 5.6.: Reduction from 3SAT: Example with four variables and three clauses.

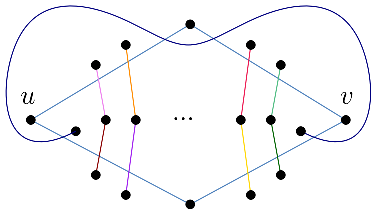

Figure 5.7.: The edge uv can be inserted in an exponential number of ways.

The general problem of finding an heterochromatic path in an edge-colored graph is NP-complete, even when each color is assigned to at most two edges.

Theorem 5.6. *Given a (multi)graph G with an edge-coloring χ and two vertices x and y, it is NP-complete to decide whether there is a heterochromatic path in G from x to y, even when each color is assigned to at most two edges.*

Proof. We reduce from 3SAT. Given a formula in 3-CNF with n variables x_1, ..., x_n and m clauses $C_1, \dots, C_m$, we construct an edge-colored (multi)graph G as the one depicted in Figure 5.6. For each clause C_j, we construct a subgraph that consists of two vertices sc_j and tc_j joined by three different edges with colors j_1, j_2, and j_3, respectively, corresponding to the (without loss of generality) three literals in the clause.

For each variable x_i, we construct a subgraph that consists of two vertices sx_i and tx_i, and two disjoint paths connecting them. The first path has its initial edge colored with color i, while the rest of the edges correspond to the literals x_i in the clauses. The second path also has its initial edge colored i, and the rest of the edges correspond to the literals $\neg x_i$ in the clauses. If an edge corresponds to the k-th literal of the clause C_j we assign color j_k to this edge.

We now join all the clause subgraphs by identifying tc_j with sc_{j+1}, for $j = 1, \dots, m-1$. We also join all the variable subgraphs by identifying tx_i with sx_{i+1}, for $i = 1, \dots, n-1$. Finally, we identify tc_m with sx_1.

It is easy to see that there is a heterochromatic path in G from sx_1 to tc_m if an only if the 3SAT instance is satisfiable. Finally, notice that we can easily modify the reduction to construct a simple graph instead of a multigraph by subdividing edges and using new colors. □

However, in our setting the multigraph and the coloring come from a simple drawing. Thus, the above reduction does not show hardness of the problem of inserting one edge into a simple drawing. In the next section we present a reduction showing that.

5.5. Inserting one edge is NP-complete

We start by noting that there can be an exponential number of ways to insert an edge; see Figure 5.7. The construction in Figure 5.7, though with a trivial solution, shows the spirit of the reduction that we will present in this section. In order to connect u and v, some "barriers" consisting of (parts of) edges in the drawing must be passed, and for that we have to decide which edges to cross.

Theorem 5.7. *Given a simple drawing $D(G)$ of a graph $G = (V, E)$ and an edge uv of the complement of G, it is* NP*-complete to decide whether uv can be inserted into $D(G)$, even if $V \setminus \{u, v\}$ induces a matching in G and u and v are isolated vertices.*

We show NP-hardness via a reduction from 3SAT. Let ϕ be a 3SAT-formula with variables $x_1, \ldots, x_n$ and with set of clauses $\mathcal{C} = \{C_1, \ldots, C_m\}$. To make the reduction easier to describe, we assume that in ϕ each clause has three literals (possibly with duplicated literals). In a preprocessing step we transform ϕ into an equivalent formula in which no clause has three positive or three negative literals.

Claim 1. *The following transformation of the clauses in a formula preserves satisfiability of the formula:*

$$x_i \vee x_j \vee x_k \Rightarrow \begin{cases} x_k \vee y \vee \texttt{false} & \textit{(i)} \\ x_i \vee x_j \vee \neg y & \textit{(ii)} \end{cases} \qquad \neg x_i \vee \neg x_j \vee \neg x_k \Rightarrow \begin{cases} \neg x_i \vee \neg x_j \vee y & \textit{(iii)} \\ \neg x_k \vee \neg y \vee \texttt{false} & \textit{(iv)} \end{cases}$$

where y is a new variable for each transformed clause and `false` is the constant truth value `false`.

Proof. We prove the statement for the case in which the original clause has three positive literals, the other case is analogous. Assume x_i or x_j satisfies the original clause. Then it also satisfies Clause (ii) and y can be set to `true` to satisfy Clause (i). If x_k satisfies the original clause, then it also satisfies Clause (i) and y can be set to `false` to satisfy Clause (ii). If none of x_i, x_j, and x_k satisfies the original clause, then to satisfy Clause (i) we have to set y to `true`, which implies that case Clause (ii) is not satisfied. □

We remark that without using the constant truth value `false`, a formula in which every clause has a positive and a negative literal can be trivially satisfied by setting all variables to `true` or to `false`.

After transforming a formula, the clauses are of four types depending on the number of positive and negative literals (and `false` constants). Clauses (i)–(iv) in Claim 1 are each of one of these

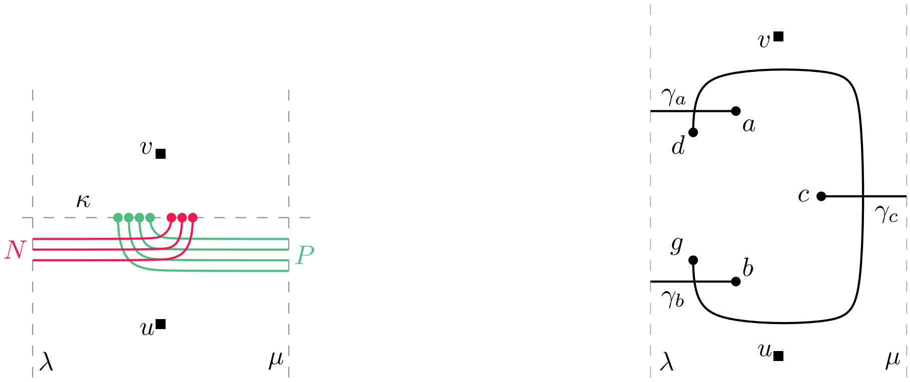

(a) Variable gadget. The red arcs belong to N, the green ones to P.

(b) Clause gadget.

Figure 5.8.: The two main gadgets for the reduction.

types. Consequently, we denote these types Type (i)–(iv). This means that clauses of Type (i) have two positive literals and one constant `false`, clauses of Type (ii) have two positive and one negative literal, Type (iii) clauses contain two negative and one positive literal, and finally, a Type (iv) clause has two negative literals and one constant `false`.

Given a transformed 3SAT-formula ϕ be a formula with variables $x_1, \ldots, x_n$ and with set of clauses $\mathcal{C} = \{C_1, \ldots, C_m\}$, the reduction uses gadgets consisting of simple drawings to represent the variables and clauses. Satisfiability of ϕ will correspond to being able to insert a given edge uv into a simple drawing D of a matching. The main idea of the reduction is that the variable and clause gadgets act as barriers inside a simple closed region region Γ of D, in which we need to insert an arc γ from one side to the other to complete the connection between u and v.

To simplify the description, we first restrict our attention to the inside of the simple closed region Γ. We assume that γ cannot cross the boundary of Γ. In the following we use two lines, named λ and μ, to bound the regions in which a variable and clause gadget will be placed. In particular, these lines will be identified with opposite segments on the boundary of Γ.

Variable gadget. A variable gadget W includes two sets of arcs (parts of later-defined edges), P and N, that correspond to positive and negative appearances of a variable, respectively. The gadget is bounded on the left by a line λ and on the right by a line μ. Arcs in P and N have one endpoint on a horizontal line κ such that the endpoints of arcs in P are to the left of the endpoints of arcs in N and all these arcs lie between λ and μ. The other endpoint of arcs in P and N lies below κ and on μ and λ, respectively. Notice that an arc in P intersects every arc in N, and vice versa; see Figure 5.8(a) for an illustration. Finally, we choose two points u and v such that u is below all arcs in W and v is above them.

Lemma 5.8. *Let W be a variable gadget. Any arc between the vertical lines λ and μ that connects u and v crosses either all arcs in P or all arcs in N.*

Proof. Assume there is an arc connecting u and v that neither crosses all the arcs in P nor crosses all the arcs in N. Hence, there are two arcs $p \in P$ and $n \in N$ such that this arc neither crosses p nor n. By the construction of the gadget, p and n cross. Thus, their union together with λ and μ separates u from v. It follows that the arc has to cross either p or n, a contradiction. $\square$

Clause gadget. A clause gadget K includes three arcs γ_a, γ_b, and γ_c (parts of later-defined edges) incident to three points a, b, and c, respectively, and an arc (edge) dg incident to two other points d and g. As in the variable gadget, the clause gadget is bounded on the left by a line λ and on the right by a line μ. The arcs γ_a and γ_b have their other endpoint on λ and γ_c has its other endpoint on μ. None of these three arcs intersect. The arc dg is placed such that it crosses γ_a, γ_c, and γ_b in that order as we traverse it from d to g; see Figure 5.8(b) for an illustration. Notice that we do not require any specific rotation of the crossings of dg with γ_a and γ_b (where the rotation is the clockwise order of the endpoints of the crossing arcs). Finally, we choose two points u and v such that u is below all arcs in K and v is above them.

Lemma 5.9. *Let K be a clause gadget. Any arc uv between the vertical lines λ and μ that connects u and v crosses either dg twice or at least one of the arcs γ_a, γ_b, and γ_c.*

Proof. Let $\times$ be the crossing point of γ_c and dg. This point splits the arc dg into two arcs $d\times$ and $g\times$. Assume that the arc uv does not cross the arcs γ_a, γ_b, and γ_c. The union of γ_a and γ_c together with $d\times$ and the lines λ and μ separates u from v. Since the arcs γ_a and γ_c are not crossed by uv, uv must cross $d\times$ in a point that is not $\times$. Analogously, the union of γ_b, γ_c together with $g\times$ and the lines λ and μ separates u from v. Thus, uv has to cross $g\times$ in a point that is not $\times$. This implies that uv crosses dg twice. $\square$

Reduction. Let ϕ be a 3SAT-formula transformed as described in Claim 1 and with set of clauses $\mathcal{C} = \{C_1, \ldots, C_m\}$ (each clause being of one of the four types described above). To build our reduction we need one more gadget. First, we introduce the following simple drawing described by Kynčl et al. [140, Figure 11] and depicted in Figure 5.9(a). Here, we denote this drawing by $\circledcirc$. It might be of independent interest to note that it can be modified to a drawing of an arbitrarily large matching in which a specific edge uv cannot be inserted but such that if we remove any edge or any other vertex then uv can be inserted; see Figure 5.10 (right). Figure 5.10 (left) shows a different drawing of a matching where the edge uv cannot be inserted.

Following the notation by Kynčl et al., we denote the six arcs in $\circledcirc$ by a_1, a_2, a_3, b_1, b_2, and b_3 and its eight cells by X, A_1, A_2, A_3, B_1, B_2, B_3, and Y; see Figure 5.9(a). The core property $\mathcal{P}$ of $\circledcirc$ is that it is not possible to insert an edge between a point in cell X and another point in cell Y such that the result is a simple drawing [140, Lemma 15].

For our reduction we first choose two arbitrary points u and v in the cells X and B_2 and insert them as vertices to $\circledcirc$. Let $\circledcirc'$ be the simple drawing in which we inserted the vertices u and v into $\circledcirc$. Finally, let b_2^* be the part of the arc b_2 between the crossing point of b_2 and a_2 and the crossing point of b_2 and b_3.

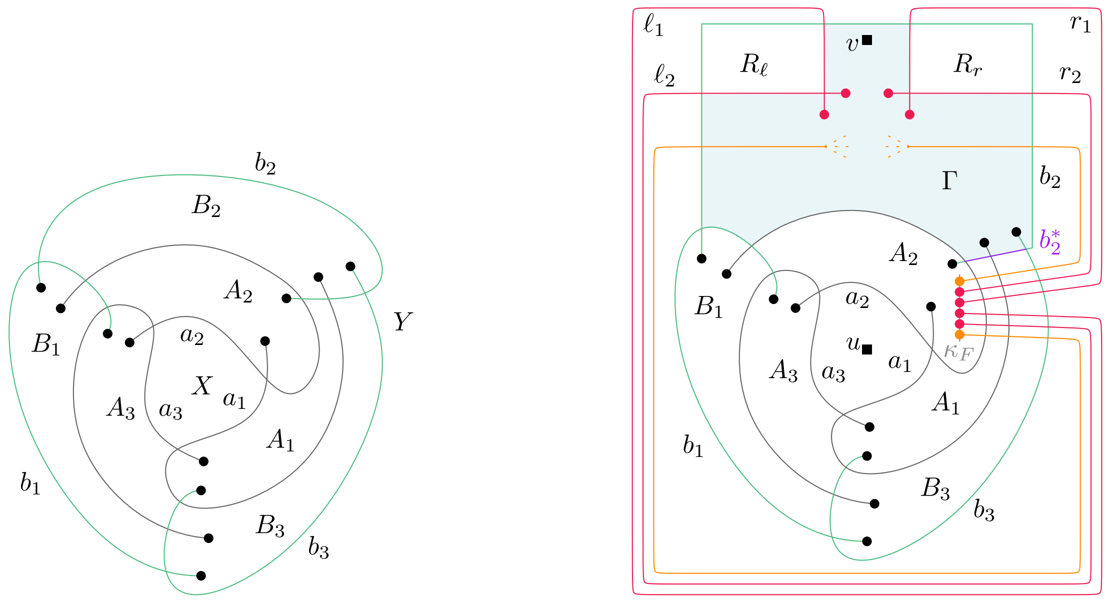

(a) Drawing ◎ by Kynčl et al. [140, Figure 11]. It is not possible to insert an edge between a point in X and one in Y.

(b) A schematic overview of the edges in F (red and orange) and how they are combined with the drawing by Kynčl et al.

Figure 5.9.: Last gadget for the reduction acting as a frame.

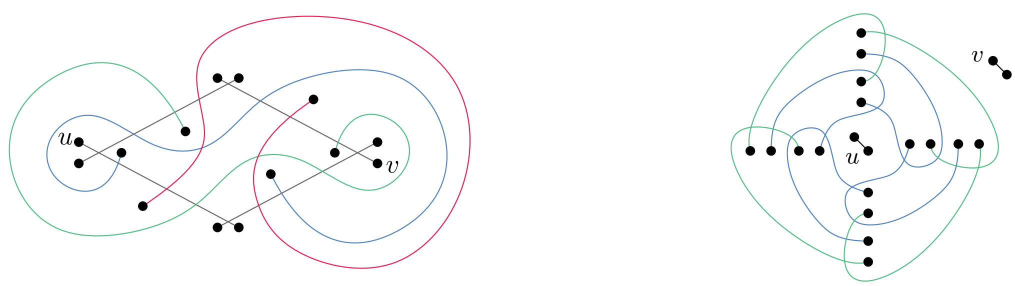

Figure 5.10.: Matchings in which the edge uv cannot be inserted. The drawing on the right is a modification of the drawing ◎ by Kynčl et al. [140, Figure 11].

Lemma 5.10. *The edge uv cannot be inserted into $\circledcirc'$ without crossing b_2^*.*

Proof. Assume for contradiction that uv can be inserted not crossing b_2^* and let γ_{uv} be such an arc. If γ_{uv} did not cross b_2, then we would be able to prolong it and cross b_2 to reach Y, a contradiction to property $\mathcal{P}$. Thus, γ_{uv} crosses b_2. Further, we may assume without loss of generality that γ_{uv} does not cross b_2 inside A_2 or B_1, as otherwise it would be possible to modify γ_{uv} to not cross b_2. Thus, γ_{uv} intersects B_2 on one side of the crossing with b_2. Since γ_{uv} cannot intersect Y, this crossing must be on b_2^*. □

The final piece we need for our a reduction is a set F of $m^I + m^{IV} + 4$ arcs that we insert into $\circledcirc'$, where m^I is the number of clauses of Type (i) and m^{IV} the number of clauses of Type (iv). For an arc $f \in F$ we will place one of its endpoints on a vertical line κ_F inside A_2 and the other one inside B_2. The only crossings of f with $\circledcirc'$ are with the arcs a_2, a_1, b_3, and b_2, in that order when traversing f from its endpoint on κ_F to its endpoint in B_2. Furthermore, f, traversed in that direction, crosses from A_2 to A_1, from A_1 to B_3, from B_3 to Y, and from Y to B_2.

Consider the $m^I + m^{IV} + 4$ endpoints on κ_F sorted from top to bottom. We denote by f_j the arc in F incident to the j-th such endpoint. When traversing b_2 from its endpoint in A_2 to its endpoint in B_1, the crossings of arcs in F with b_2 appear in the same order as their endpoints on κ_F. More precisely, the crossings of b_2 when traversed in that direction are with a_2, a_1, b_3, f_1, f_2, $\dots$, $f_{|F|}$, and b_1.

The arcs f_{m^I+1}, f_{m^I+2}, f_{m^I+3}, and f_{m^I+4} will behave differently than the other arcs in F. In the following, we denote these four arcs by r_2, r_1, ℓ_1, and ℓ_2, respectively. There are only two crossings between arcs in F, namely of r_1 and r_2, and of ℓ_1 and ℓ_2, and both these crossings are inside B_2. These four crossing arcs divide B_2 into three regions. We denote the region by b_2^* on its boundary with Γ, the (other) region with the crossing of r_1 and r_2 on its boundary with Γ_r, and the (other) region with the crossing of ℓ_1 and ℓ_2 on its boundary with Γ_ℓ. Arcs r_1, r_2, ℓ_1, and ℓ_2 must be drawn such that the vertex v lies in Γ; see Figure 5.9(b) for an illustration. The precise endpoints of the edges in $F \setminus \{r_1, r_2, \ell_1, \ell_2\}$ will be fixed when we insert the clause gadgets.

Lemma 5.11. *The edge uv cannot be inserted into $\circledcirc'$ without crossing every arc in F inside A_1 or B_3.*

Proof. Assume for contradiction that there is an arc $f \in F$ such that uv does not cross f. From Lemma 5.10 we know that uv has to cross b_2^*. Consider the region bounded by b_2^*, b_3, f, and a_2. Observe that, since b_2^* is fully contained on the boundary of this region, uv has to cross at least one of the three other arcs as well. By assumption, uv does not cross f. Crossing b_3 is impossible by property $\mathcal{P}$, as the part contained on this region's boundary separates B_3 from Y. Finally, crossing the arc which is part of a_2 is not possible, since this would imply the existence of a point v' in A_2 such that uv passes through v' without having crossed a_2. Hence, we could prolong the arc uv' that is part of uv by crossing a_2 such that it reaches B_2 without having crossed b_2^*, a contradiction to Lemma 5.10. Furthermore, as we do not allow more than two arcs to cross in one point, the statement follows. □

It remains to insert inside Γ the clause and variable gadgets and precisely define the endpoints of arcs in $F \setminus \{\ell_1, \ell_2, r_1, r_2\}$. For simplicity, we first insert the variable gadgets and then the clause gadgets. The idea is that each clause and variable gadget is inserted in Γ separating b_2^* from v. This is done by identifying the endpoints that were lying on λ or μ with points on ℓ_1, ℓ_2, r_1, r_2, or b_2. As a result, Lemmas 5.8 and 5.9 can be applied to the arc that we aim to insert connecting u and v in the final simple drawing, since it has to cross b_2^* by Lemma 5.10.

We insert now the variable gadgets into Γ. Let $W^{(i)}$ be the variable gadget corresponding to variable x_i. For a gadget $W^{(i)}$, the arcs in N are drawn such that the endpoints on λ, lie on the part of ℓ_1 that bounds Γ. The arcs in P are drawn similarly, but with the endpoints on μ lying on the part of r_1 that bounds Γ. Moreover, we identify vertex v in the gadget with vertex v in $\circledcirc'$. Gadgets corresponding to different variables are inserted without crossing each other. We now specify how they are inserted relative to each other. As we traverse ℓ_1 from its endpoint on κ_F to its endpoint in Γ we encounter the endpoints of arcs in $W^{(i)}$ before the endpoints of arcs in $W^{(i+1)}$. Analogously, as we traverse r_1 from its endpoint on κ_F to its endpoint in Γ we encounter the endpoints of arcs in $W^{(i)}$ before the endpoints of arcs in $W^{(i+1)}$. An illustration is presented in Figure 5.11.

The clause gadgets are inserted in a similar way. Let $K^{(j)}$ be the clause gadget corresponding to clause C_j. If C_j is of Type (i), $K^{(j)}$ is inserted such that the endpoints on λ lie on the part of ℓ_2 that bounds Γ. If C_j is the j'-th clause of Type (i), we identify c with the endpoint of the arc $f_{j'}$. Similarly, if C_j is of Type (iv), $K^{(j)}$ is inserted such that the endpoints on λ lie on the part of r_2 that bounds Γ. If C_j is the j'-th clause of Type (iv), we identify c with the endpoint of the arc $f_{m^I+4+j'}$. If C_j is of Type (ii), $K^{(j)}$ is inserted such that the endpoints on λ lie on the part of ℓ_2 that bounds Γ and the endpoint on μ lies on the part of r_2 that bounds Γ. Similarly, if C_j is of Type (iii), $K^{(j)}$ is inserted such that the endpoint on μ lies on the part of ℓ_2 that bounds Γ and the endpoints on λ lie on the part of r_2 that bounds Γ. The crossings in Γ of arcs from different clause gadgets are of arcs with an endpoint in r_2 with arcs in $\{f_j : 1 \leq j \leq m^I\}$.

We now specify how different clause gadgets are inserted relative to each other. As we traverse ℓ_2 from its endpoint on κ_F to its endpoint in Γ we encounter the endpoints of arcs corresponding to clauses of Type (iii) before the ones corresponding to clauses of Type (ii), and those before the ones corresponding to clauses of Type (i). Analogously, as we traverse r_2 from its endpoint on κ_F to its endpoint in Γ we encounter the endpoints of arcs corresponding to clauses of Type (iv) before the ones corresponding to clauses of Type (iii), and those before the ones corresponding to clauses of Type (ii). Moreover, as we traverse ℓ_2 and r_2 in the specified directions, the endpoints of arcs corresponding to the j'-th clause of a certain type are encountered before the endpoints of arcs corresponding to the $(j'-1)$-st clause of this type. An illustration is presented in Figure 5.11.

Finally, we connect arcs from variable and clause gadgets inside the regions Γ_ℓ and Γ_r. This is done such that if a literal in a clause is x_k then the corresponding arc in the clause gadget, that has an endpoint on ℓ_2, is connected with an arc in N of the gadget $W^{(k)}$, that has an endpoint on ℓ_1. Thus, these connections can lie in Γ_ℓ. Analogously, if a literal in a clause is $\neg x_k$ then the corresponding arc in the clause gadget, that has an endpoint on r_2, is connected with an arc in P of the gadget $W^{(k)}$, that has an endpoint on r_1. Thus, these connections can lie in Γ_r. Since, without loss of generality, we can assume that Γ_ℓ and Γ_r are convex regions and the endpoints we want to connect are in general position (no three on the same line), the connections can be drawn as straight-line segments. (For clarity, in Figure 5.11 these connections have one bend per arc.) Therefore, there is at most one crossing between each pair of connecting arcs.

Each connecting arc can be concatenated with the arcs in a variable and in a clause gadget that it joins. These concatenated arcs are edges in our drawing that have one endpoint in a variable gadget and the other in a clause gadget. By construction, each of them corresponds to a literal in the formula ϕ and each pair of these edges crosses at most once. Similarly, the arcs in $F \setminus \{\ell_1, \ell_2, r_1, r_2\}$ have one endpoint in a clause gadget and also define a set of edges in our final drawing that we denote by the same name as the corresponding arcs.

We now have all the pieces that constitute our final drawing. It consists of

(i) the simple drawing $\circledcirc'$;

(ii) the edges $f_i \in F$ drawn as the described arcs (with their endpoints as vertices);

(iii) the edges corresponding to literals (with their endpoints as vertices); and

(iv) the edges dg in each clause gadget (with d and g as vertices).

Observe that the constructed drawing is a simple drawing, as it is the drawing of a matching (plus the vertices u and v) and, by construction, two edges cross at most once.

Correctness. It is now straight-forward to show that the presented construction is a valid reduction.

Proof of Theorem 5.7. We show that the above construction is a polynomial-time reduction from 3SAT to the problem of deciding whether an edge can be inserted into a simple drawing. Given a 3SAT-formula ϕ with variables $x_1, \dots, x_n$ and with clauses $C_1, \dots, C_m$ we construct a simple drawing D as above and aim to insert the edge uv into it. This construction can clearly be computed in polynomial time and space, since only the combinatorial description of the drawing is needed.

Assume that the edge uv can be inserted into D and let uv denote the resulting arc. By Lemmas 5.10 and 5.11 we know that uv has to cross b_2^* and every arc in F. Let u^* be the point where uv crosses b_2^*. Each clause and variable gadget separates u^* from v and thus, Lemmas 5.8 and 5.9 can be applied. This means that in a variable gadget $W^{(i)}$ either all arcs in P or all arcs in N are crossed. In the former case we assign to variable x_i the value `true`, and otherwise the value `false`. Assume that this truth assignment does not satisfy ϕ. Then there exists a clause C_j for which all three literals evaluate to `false`. Consider the clause gadget $K^{(j)}$. By Lemma 5.9 we must cross in it an edge corresponding to one of its literals. However, by Lemma 5.11 an edge corresponding to the constant value `false` cannot be crossed (again) in a clause gadget. By construction and the truth assignment of the variables, the edges corresponding to the other literals of C_j cannot be crossed either.

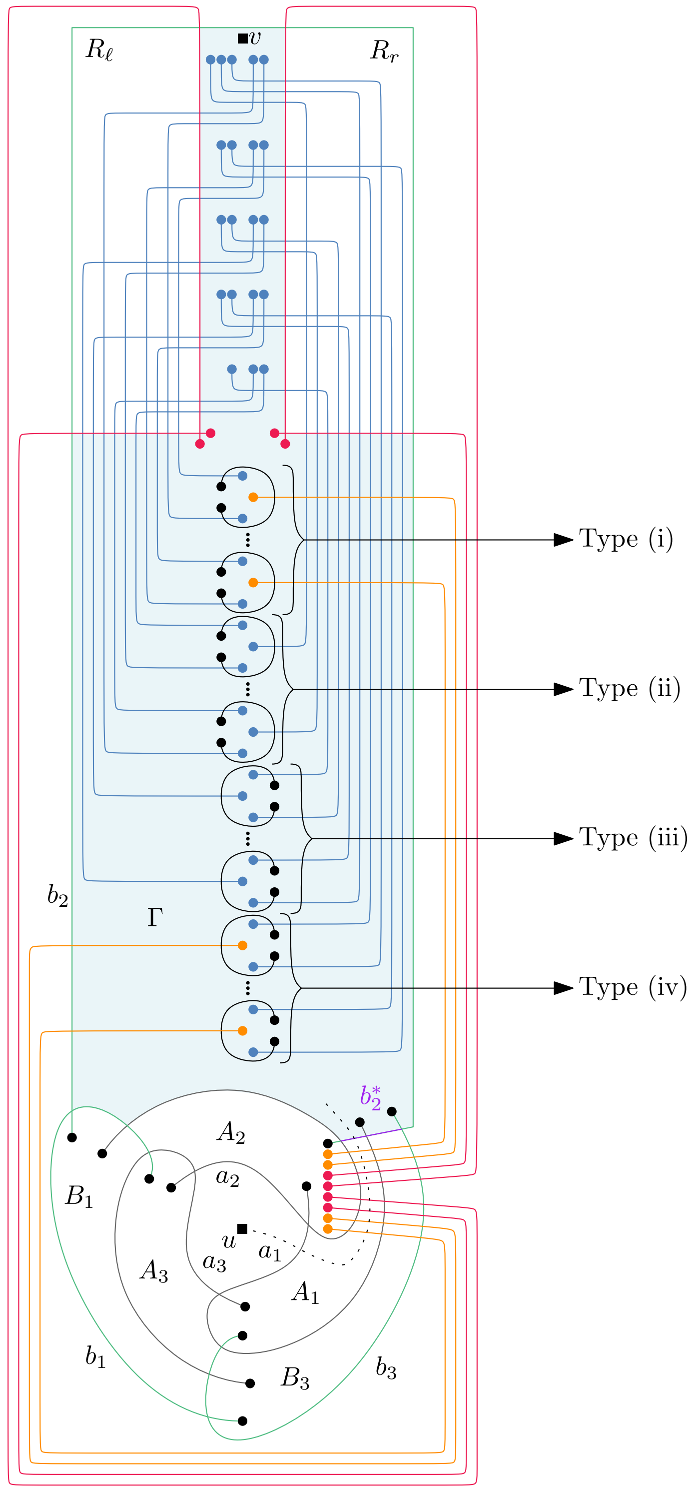

Figure 5.11.: Illustration of the reduction.

Conversely, assume we are given a satisfying assignment of ϕ. We then can insert uv into D as follows. Starting from u, edge uv crosses a_1 to enter region A_1, then crosses all arcs in F, and crosses b_2^* to enter Γ; see also the dotted line in Figure 5.11. In each clause gadget, edge uv crosses one edge corresponding to a literal evaluating to `true`, none corresponding to a literal evaluating to `false`, and the edge dg in the gadget if necessary. By construction, this leaves in each variable gadget all arcs either in P or in N free to be crossed by uv. Moreover, this allows us to connect u and v without crossing any edge twice. □

Remarks. The presented reduction from 3SAT constructs a simple drawing of a matching, and thus, the problem remains NP-hard when G is as sparse as possible (isolated vertices that are not the starting or ending vertices of the edge that we aim to insert can be disregarded for our problem, and thus we can restrict our attention to graphs without such vertices). Moreover, we only make use of the fact that in simple drawings an edge cannot cross itself and two edges can cross at most once. We do not use that two incident edges are not allowed to cross. In fact, our proof shows that deciding whether an arrangement of pseudosegments (a family of simple curves of finite length pairwise intersecting in at most one point) can be extended by one more pseudosegment is NP-hard. (The existance of such an extension can be seen as the analogue of Levi's enlargement lemma for pseudosegments.)

We remark that if we do not require G to be a matching, our variable gadget can be simplified by identifying all the vertices on κ and removing the crossings between edges in N and P. Furthermore, the disconnectedness of the produced instance is not a restriction. If an instance $D(G)$ is a simple drawing of a disconnected graph G we can transform it to an equivalent instance consisting of a simple drawing of a connected graph by inserting an apex vertex into any cell of the drawing and subdividing its incident edges that connect to all the vertices of $D(G)$.

5.6. A case in which inserting one edge is in P

The following theorem shows that we can decide in polynomial time whether we can extend a simple drawing $D(G)$ of a graph G with the edge uv from $\overline{G}$, when $\{u, v\}$ is a dominating set for G.

Theorem 5.12. *Let $D(G)$ be a simple drawing of a graph $G = (V, E)$ and let u, $v \in V$ be non-adjacent vertices. If $\{u, v\}$ is a dominating set for G, that is, every vertex in $V \setminus \{u, v\}$ is a neighbor of u or v, then the problem of extending $D(G)$ with the edge uv can be decided in polynomial time.*

The rest of this section is devoted to prove Theorem 5.12. The first step is to reduce the problem of inserting the edge uv to the problem of finding a valid path crossing some colored arcs at most once in a plane with some forbidden regions (holes). This new problem has the advantage of being a more suitable ground for inductive proofs. The main ingredients needed in our algorithm are a series of lemmas describing sufficient conditions for which this problem has a solution.

For an integer $k \geq 1$, a *plane with k holes* is a set $\Gamma \subseteq \mathbb{R}^2$ obtained from considering k disjoint simple closed curves in $\mathbb{R}^2$, all bounding a common cell, and removing for each curve C the cell bounded by C that is disjoint from the rest of the curves. If $k = 1$, then only one side of C is removed. The closure of each removed cell is a *hole of* Γ.

Path problem with holes. Given a plane with holes Γ and a set of colored Jordan arcs $\mathcal{J}$ drawn in Γ, the path problem with holes asks whether there is a Jordan arc connecting two points p, $q \in \Gamma \setminus \mathcal{J}$, called *terminals*, that crosses at most one arc in $\mathcal{J}$ of each color. If such a pq-arc exists, then it is a *valid pq-arc* for the instance $(\Gamma, \mathcal{J}, p, q)$.

We assume that every instance of the path problem with holes that we consider meets the following properties:

(i) Every two arcs of $\mathcal{J}$ share at most one point.

(ii) Pairs of distinct arcs in $\mathcal{J}$ having the same color are disjoint.

(iii) Each arc in $\mathcal{J}$ starts and ends on the boundary of Γ, that is, no arc has an endpoint in the interior of Γ.

Reduction. Let $D(G)$ be a drawing of a graph G, and let $\{u, v\}$ be a dominating set of vertices in G such that uv is an edge of $\overline{G}$. We now reduce the problem of deciding whether uv can be inserted into $D(G)$ to the path problem with at most two holes.

If G' is a subgraph of G, then we denote by $D[G']$ the subdrawing of D induced by the vertices and edges of G'. In a slight abuse of notation, if G' consists only of a vertex v or of an edge $e = uv$, then we will write $D[v]$ and $D[e]$ (or $D[uv]$), respectively.

For a vertex v of G, the *star of v* consists of v, its adjacent vertices, and its incident edges. Let S_u and S_v be the subdrawings of D induced by the stars of u and v, respectively. Moreover, let H be the subgraph of G that is the union of the stars of u and v. Then, S_u and S_v are plane stars whose union is $D[H]$. If an extension with uv exists, then the arc connecting u and v representing the edge uv cannot cross any of those edges and must lie in the closure of a cell F of $D[H]$ with u and v on its boundary. Thus, our problem reduces to testing the existence of a valid uv-arc in each cell F of $D[H]$ with both u and v on its boundary.

We can assume without loss of generality that u and v are incident to at least one edge by maybe inserting small segments incident to them. Let F be a cell of $D[H]$ with both u and v on its boundary. Notice that it might be bounded or unbounded. Moreover, the part of $D[H]$ that is in the closure of F can be connected of disconnected.

If it is connected, we consider a simple closed curve C in the interior of F, closely following the part of $D[H]$ that is in the closure of F. We slightly modify C so that, at a certain occurrence of u and of v on ∂F, the curve C touches ∂F; see the dashed curve in Figure 5.12 (right). In our reduction we consider all possible modifications of C, differing on where we decide to make C touch u and v. The number of possible resulting curves is at most the degree of u times the degree of v. In this case we define $C' = C$.

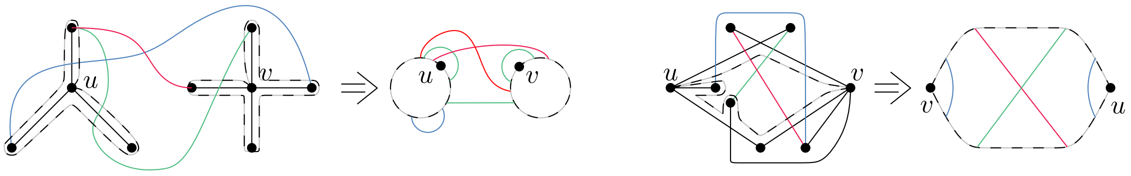

Figure 5.12.: Reduction to the path problem with holes.

If the part of $D[H]$ that is in the closure of F is not connected it must consist of two connected components containing u and v, respectively. We consider two simple curves C and C' in the interior of F each one closely following one of these connected components. As before, we slightly modify the curves so that, at a certain occurrence of u and of v on ∂F, they touch ∂F; see the dashed curves in Figure 5.12 (left).

In both cases, we consider the inside of the curves C and C' to be the regions bounded by them and such that the union of their closures contains $S_u \cup S_v$. Let Γ be the closure of the region consisting of F with the inside of the curves C and C' removed. Then, Γ is a plane with at most two holes (the closures of the inside of the curves C and C').

To finish our reduction, we need to identify the set of colored Jordan arcs and the two terminals in the path problem with holes. The set $\mathcal{J}$ is the defined as the union of the arcs of $D[e] \cap \Gamma$, for each edge $e \in E$. In order to assign colors to the arcs in $\mathcal{J}$, we first assign a different color to each edge of G. Each arc of $D[e] \cap \Gamma$ then inherits the color of e; see Figure 5.12. Finally, the terminals p and q are points in the two cells of $C \cup \mathcal{J}$ in Γ having $D[u]$ and $D[v]$ on their boundary, respectively.

Notice that a reduction from the problem of inserting an edge into a simple drawing to the path problem with holes results in an instance satisfying properties (i) and (ii). Moreover, if $\{u, v\}$ is a dominating set for G, then the instance of the path problem with holes also meets property (iii). The discussion above leads to the following statement:

Observation 5.13. *Let $D(G)$ be a simple drawing of a graph $G = (V, E)$ and let u, $v \in V$ be non-adjacent vertices such that $\{u, v\}$ is a dominating set for G. The problem of deciding whether uv can be inserted into $D(G)$ can be reduced to the path problem with at most two holes.*

We now prepare the tools for solving in polynomial time an instance of the path problem with at most two holes with properties (i)–(iii). Apart from introducing the notation and operations used in the algorithm solving that problem, we will show that if all arcs are of different colors, then there is always a solution.

Given a plane with holes Γ and a set of Jordan arcs $\mathcal{J}$ in Γ, a *cell* of $(\Gamma, \mathcal{J})$ is the interior of a component of $\Gamma \setminus \mathcal{J}$. For any arc $\alpha \in \mathcal{J}$, a *segment of* α is the closure of a component of $\alpha \setminus (\mathcal{J} \setminus \{\alpha\})$. If the set of arcs has one element, $\mathcal{J} = \{\alpha\}$, then, we abuse notation by writing (Γ, α) instead of $(\Gamma, \{\alpha\})$. Two cells of $(\Gamma, \mathcal{J})$ are *adjacent* if they share a segment of an arc in $\mathcal{J}$. Given two points $p, q \in \Gamma$ and a Jordan arc α, α is *pq-separating* if every pq-arc in Γ intersects α.

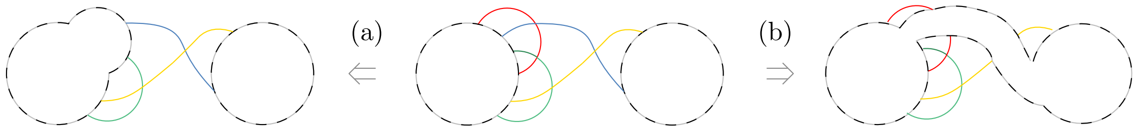

Figure 5.13.: The two operations transforming an instance of the path problem with at most two holes: (a) enlarging a hole along an arc and (b) cutting through an arc.

In the following, let $(\Gamma, \mathcal{J}, p, q)$ be an instance of the path problem with at most two holes and properties (i)–(iii). Then, a pq-separating arc $\alpha \in \mathcal{J}$ has its ends on the same hole of Γ and p and q are in different cells of (Γ, α). Moreover, each arc $\alpha \in \mathcal{J}$ is one of the following three types:

T1: α has its ends on two different holes of Γ;

T2: α has its ends on the same hole of Γ and is pq-separating; and

T3: α has its ends on the same hole of Γ and is not pq-separating.

We say that two instances of a problem are *equivalent* if the lead to the same output of a decision problem. The following operation shows how to transform any instance $(\Gamma, \mathcal{J}, p, q)$ into another equivalent one where no arcs of Type T3 occur.

Enlarging a hole along an arc. If there is an $\alpha \in \mathcal{J}$ such that α is of Type T3, having both its ends on the same hole h, then the operation of *enlarging a hole along* α converts $(\Gamma, \mathcal{J}, p, q)$ into a new instance $(\Gamma', \mathcal{J}', p, q)$, where Γ' is obtained from Γ by removing the cell of (Γ, α) disjoint from p and q and $\mathcal{J}' = \mathcal{J} \cap \Gamma'$; see Figure 5.13 for an illustration.

Lemma 5.14. *Let $(\Gamma, \mathcal{J}, p, q)$ be an instance of the path problem with at most two holes and that meets properties (i)–(iii) and let $(\Gamma', \mathcal{J}', p, q)$ be the instance obtained from $(\Gamma, \mathcal{J}, p, q)$ by enlarging a hole along an arc α of Type T3. Then, for every arc $\beta \in \mathcal{J}$ there is at most one arc $\beta' \in \mathcal{J}'$ and it is of the same type as β. Thus, $|\mathcal{J}'| < |\mathcal{J}|$. Moreover, there is a valid pq-arc in $(\Gamma, \mathcal{J}, p, q)$ if and only if there is a valid pq-arc in $(\Gamma', \mathcal{J}', p, q)$.*

Proof. To see that the first part holds, consider an arc $\beta \in \mathcal{J} \setminus \{\alpha\}$ with $\beta \cap \Gamma' \neq \emptyset$. Let F be the cell of (Γ, α) disjoint from p and q. If $\alpha \cap \beta = \emptyset$, then $\beta \cap \Gamma' = \beta'$. Thus, the remaining case is that α and β cross, and because they can only cross once, $\beta \setminus \alpha$ has two components: one is included in F, while the other is in $\Gamma \setminus F = \Gamma'$ and its closure is β'. Moreover, since p and q are not in F, p and q belong to the same cell of (Γ, β) if and only if they belong to the same cell of (Γ', β'). Also β has its ends on different holes if and only β' has its ends on different holes of Γ'. Now the second part of the lemma follows from the fact that any valid pq-arc in $(\Gamma, \mathcal{J})$ does not cross α, since p and q are in the same cell of (Γ, α). $\square$

The operation of enlarging a hole along an arc allows us to eliminate all arcs of Type T3. Thus, if our instance has only one hole, then we can transform it to one where there are only arcs of Type T2. If there are two arcs of Type T2 of the same color, then it is clear that there cannot

Figure 5.14.: Examples of instances with two holes.

be a solution. The following result shows that this condition is also sufficient for instances with only one hole.

Lemma 5.15. *Let $(\Gamma, \mathcal{J}, p, q)$ be an instance of the path problem with one hole that meets properties (i)–(iii). Then a valid pq-arc exists if and only if there are no two pq-separating arcs of the same color.*

Proof. Suppose that $\mathcal{J}$ has at most one pq-separating arc of each color. To show that there is a valid pq-arc, we proceed by induction on $|\mathcal{J}|$. The base case $|\mathcal{J}| = 0$ clearly holds. Henceforth, we assume $|\mathcal{J}| \geq 1$.

If an arc in $\mathcal{J}$ is not pq-separating, then we apply Lemma 5.14 to reduce $(\Gamma, \mathcal{J}, p, q)$ into an instance $(\Gamma', \mathcal{J}', p, q)$ with fewer arcs and satisfying the same conditions as $(\Gamma, \mathcal{J}, p, q)$. The induction hypothesis implies the existence of valid pq-arc in $(\Gamma', \mathcal{J}', p, q)$, that, by Lemma 5.14, also implies the existence of a valid one for $(\Gamma, \mathcal{J}, p, q)$.

Suppose now that every arc $\mathcal{J}$ is pq-separating. Since $|\mathcal{J}| \geq 1$, p and q are in different cells of $(\Gamma, \mathcal{J})$. Let F_p be the cell containing p and let $\alpha \in \mathcal{J}$ be an arc with a segment σ on the boundary of F_p. Consider a point p' in the other cell of $(\Gamma, \mathcal{J})$ having σ on its boundary.

With the exception of α, all the arcs in $\mathcal{J}$ are $p'q$-separating. From the preceding discussion it follows that a valid $p'q$-arc not intersecting α exists. No $p'q$-separating arc has the same color as α and therefore, we can extend this valid $p'q$-arc to a valid pq-arc. □

With Lemma 5.15 in hand, we can now focus on instances with two holes. In this context, the condition of not having two pq-separating arcs of the same color is not sufficient to imply the existence of a valid pq-arc, as Figure 5.14 (a) shows. However, using the following operation, we can transform an instance with two holes into an instance with only one hole when there is an arc of Type T1 that cannot be crossed by a valid arc.

Cutting through an arc. Let $\alpha \in \mathcal{J}$ be an arc of Type T1 having its ends on distinct holes. The transformed instance $(\Gamma', \mathcal{J}', p, q)$ obtained from $(\Gamma, \mathcal{J}, p, q)$ by *cutting through* α is defined as follows. Consider a thin open strip Σ in Γ covering α and neither containing p nor q. Then, $\Gamma' = \Gamma \setminus \Sigma$ (this merges the two holes of Γ into one hole) and $\mathcal{J}' = \mathcal{J} \cap \Gamma'$; see Figure 5.13 (b) for an illustration.

Observation 5.16. *Let $(\Gamma, \mathcal{J}, p, q)$ be an instance of the path problem with two holes that meets properties (i)–(iii), and let $(\Gamma', \mathcal{J}', p, q)$ be the instance obtained from $(\Gamma, \mathcal{J}, p, q)$ by cutting through an arc $\alpha \in \mathcal{J}$ of Type T1. Then, there is a valid pq-arc in $(\Gamma, \mathcal{J}, p, q)$ not crossing α if and only if there is a valid pq-arc in $(\Gamma', \mathcal{J}', p, q)$.*

Suppose that $\alpha, \beta \in \mathcal{J}$ are two crossing arcs of Type T1. Then, $(\Gamma, \{\alpha, \beta\})$ has exactly three cells; see Figure 5.14 (b). Moreover, if the terminals p and q are located in the pair of non-adjacent cells, then any valid pq-arc is forced to cross both α and β. The next result shows that if, for an arc α of Type T1, there is no arc β of Type T1 producing this situation and all arcs are of different colors, then there is a valid pq-arc not crossing α.

Lemma 5.17. *Let $(\Gamma, \mathcal{J}, p, q)$ be an instance of the path problem with two holes that meets properties (i)–(iii). Suppose that every arc in $\mathcal{J}$ is either of Type T1 or of Type T2 and that all the arcs in $\mathcal{J}$ are of different colors. Let $\alpha \in \mathcal{J}$ be any arc of Type T1. If, for every Type T1 arc $\beta \in \mathcal{J} \setminus \{\alpha\}$ crossing α, p and q are in adjacent cells of $(\Gamma, \{\alpha, \beta\})$, then there is a valid pq-arc not intersecting α.*

Proof. Let $(\Gamma', \mathcal{J}', p, q)$ be the instance obtained from cutting Γ along α. Let h be the hole of Γ' obtained from merging the two holes h_1 and h_2 of Γ with a thin strip covering α. We decompose the boundary of h as the union of four arcs α_1, γ_1, α_2 and γ_2, where α_1 and α_2 bound the strip covering α and, for $i = 1, 2$, γ_i is the arc on the boundary of h_i connecting α_1 and α_2.

From Observation 5.16, it is enough to show the existence of a valid pq-arc in $(\Gamma', \mathcal{J}', p, q)$. Assume for contradiction that there is no valid pq-arc for $(\Gamma', \mathcal{J}', p, q)$. Lemma 5.15 shows that then there are two separating pq-arcs β_1, $\beta_2 \in \mathcal{J}'$ of the same color. Since all arcs in $\mathcal{J}$ are of different colors, if $\beta \in \mathcal{J} \setminus \{\alpha\}$, then the arc components (one or two, depending on whether α and β cross or not) of $\Gamma' \cap \beta$ induce one chromatic class of arcs in $\mathcal{J}'$. Thus, there is an arc $\beta \in \mathcal{J}$ that crosses α and with two arc components β_1 and β_2 of $\Gamma' \cap \beta$.

Since β crosses α, each of β_1 and β_2 has exactly one endpoint on a different arc of α_1 and α_2. By possibly relabeling β_1 and β_2, we may assume that, for $i = 1, 2$, β_i has an endpoint a_i in α_i. For $i = 1, 2$, let b_i be the endpoint of β_i that is not a_i.

First, we suppose that both b_1 and b_2 are on the same hole of Γ, say h_1 (so β is of Type T2). Then, $(\Gamma', \{\beta_1, \beta_2\})$ has three cells. As both β_1 and β_2 are pq-separating, p and q are in the two cells of $(\Gamma', \{\beta_1, \beta_2\})$ that do not have γ_2 on the closure of their boundaries. However, these two cells are included in the same cell of (Γ, β), contradicting that β is pq-separating (and thus of Type T2).

Second, suppose that b_1 and b_2 are on different holes (so β is of Type T1). By symmetry, we may assume $b_1 \in h_1$ and $b_2 \in h_2$. There are three cells of $(\Gamma', \{\beta_1, \beta_2\})$, and, since β_1 and β_2 are pq-separating, p and q are in the cells that have exactly one of β_1 and β_2 on their boundary. However, this implies that p and q are in non-adjacent cells of $(\Gamma, \{\alpha, \beta\})$, contradicting our hypothesis. □

In fact, when all the arcs in $\mathcal{J}$ are of different colors there is always a valid pq-arc:

Lemma 5.18. *Let $(\Gamma, \mathcal{J}, p, q)$ be an instance of the path problem with at most two holes that meets properties (i)–(iii). If all the arcs in $\mathcal{J}$ are of different colors, then there exists a valid pq-arc.*

Proof. If Γ has only one hole, then the result follows from Lemma 5.15, so we assume that Γ has two holes. We proceed by Induction on $|\mathcal{J}|$. The base case $|\mathcal{J}| = 1$ clearly holds. Suppose that $|\mathcal{J}| \geq 2$.

If $\mathcal{J}$ has an arc of Type T3, then we can apply Lemma 5.14 to obtain an instance with fewer arcs that satisfies the same conditions as $(\Gamma, \mathcal{J}, p, q)$. The induction hypothesis shows that there is a valid pq-arc for the transformed instance, and thus, there is also a valid pq-arc in $(\Gamma, \mathcal{J}, p, q)$. Henceforth, we assume that $\mathcal{J}$ has only arcs of types T1 and T2.

Let F_p be the cell of $(\Gamma, \mathcal{J})$ containing p and let $\alpha \in \mathcal{J}$ be an arc having a segment σ on the boundary of F_p. Consider a point p' in the cell adjacent to F_p that has σ on its boundary.

If α is of Type T2 (with respect to terminals p and q), then, as α is not $p'q$-separating, applying Lemma 5.14 as before shows that there is a valid $p'q$-arc (not crossing α) that can be extended to a valid pq-arc.

Thus, the only remaining case is that α is of Type T1, so we assume that α has its ends on two different holes h_1 and h_2.

Claim 1. *Either there is a valid pq-arc not intersecting α or there is a valid $p'q$-arc not intersecting α.*

Proof. Assume for contradiction that there are no valid pq- and $p'q$-arcs disjoint from α. Lemma 5.17 implies that there is a Type T1 arc $\beta \in \mathcal{J} \setminus \{\alpha\}$ crossing α, such that the two non-adjacent cells F_p^β and F_q^β of $(\Gamma, \{\alpha, \beta\})$ contain p and q, respectively. Let F^β be the other cell of $(\Gamma, \{\alpha, \beta\})$ neither including p nor q. Likewise, there exists $\beta' \in \mathcal{J} \setminus \{\alpha\}$ crossing α, such that the two non-adjacent cells $F_{p'}^{\beta'}$ and $F_q^{\beta'}$ of $(\Gamma, \{\alpha, \beta'\})$ contain p' and q, respectively. Let $F^{\beta'}$ be the other cell of $(\Gamma, \{\alpha, \beta'\})$.

Let $\times$ and $\times'$ be the crossings between α and β and between α and β', respectively. By symmetry, we may assume that when we traverse α from h_1 to h_2, we encounter $\times$ before $\times'$. Also, by possibly relabeling h_1 and h_2, we may assume that F_p^β has a subarc of the boundary of h_1 on its boundary, while F_q^β has a subarc of the boundary of h_2 on its boundary.

Since $p \in F_p$ and $F_p \subseteq F_p^\beta$, the segment $\sigma \subseteq \alpha$ shared by F_p and $F_{p'}$ is located on α between the endpoint of α in h_1 and $\times$. As $p' \in F_{p'}$, both p' and $F_{p'}$ are contained in F^β.

The boundary of $F_{p'}^{\beta'}$ is a simple closed curve C made of three arcs: The first one connects $\times'$ to the boundary of h_1 along α; the second one is an subarc of the boundary of h_1 connecting the endpoint of α on h_1 to the endpoint of β' on h_1; and the third one is a subarc of β' connecting the endpoint of β' on h_1 to $\times'$. Since $F_{p'}^{\beta'}$ contains $F_{p'}$, the points on $C \cap \beta'$ near $\times'$ are on the side of α that contains points in $F_{p'}$. As $\times'$ comes after $\times$ when we traverse α from h_1 to h_2, the points on $C \cap \beta'$ near $\times'$ are in F_q^β. Since the endpoint of $C \cap \beta'$ on h_1 is not in F_q^β, the arc $C \cap \beta'$ crosses β at some point $\times_{\beta,\beta'}$.

Since $\times_{\beta,\beta'}$ is the only crossing between β and β', the subarc of β' from $\times'$ to h_2 is disjoint from β. The points on this subarc near $\times'$ are in F^β, and thus, the cell $F_q^{\beta'}$ is included in F^β. However, this shows that $F_q^\beta \cap F_q^{\beta'} = \emptyset$, contradicting that $q \in F_q^\beta \cap F_q^{\beta'}$. $\square$

Algorithm 1 PPH$(\Gamma, \mathcal{J}, p, q)$: outputs whether there is a valid pq-arc.

1: **while** $\mathcal{J} \neq \emptyset$ **do**
2:   **if** $\mathcal{J}$ has an arc α of Type T3 (has its ends on same hole and is not pq-separating) **then**
3:     $(\Gamma, \mathcal{J}, p, q) \leftarrow$ ENLARGE$((\Gamma, \mathcal{J}, p, q), \alpha)$
4:   **else**
5:     **if** all arcs in $\mathcal{J}$ are of different colors **then**
6:       **return** `True`
7:     **else**
8:       find two arcs α and $\alpha' \in \mathcal{J}$ of the same color
9:       **if** both α and α' are of Type T2 (pq-separating) **then**
10:         **return** `False`
11:       **else if** α is of Type T2 (pq-separating) and α' is of Type T1 (has its ends on two holes) **then**
12:         **return** PPH(CUT$((\Gamma, \mathcal{J}, p, q), \alpha')$))
13:       **else if** both α and α' are of Type T1 (have their ends on two holes) **then**
14:         **return** PPH(CUT$((\Gamma, \mathcal{J}, p, q), \alpha)$) $\vee$ PPH(CUT$((\Gamma, \mathcal{J}, p, q), \alpha')$)
15:       **end if**
16:     **end if**
17:   **end if**
18: **end while**
19: **return** `True`

From the previous claim, either there is a valid pq-arc not crossing α or there is a valid $p'q$-arc not crossing α. In the former case we are done. In the later, we extend the valid $p'q$-arc to a valid pq-arc by crossing σ. □

With all the previous results we can now show the polynomial-time algorithm that proves Theorem 5.12. From Observation 5.13, it is enough to solve the path problem with at most two holes for instances meeting properties (i)–(iii) in polynomial time. To show this we consider Algorithm 1.

In Algorithm 1, $(\Gamma, \mathcal{J}, p, q)$ is an instance of the path problem with at most two holes (PPH) meeting properties (i)–(iii). ENLARGE$((\Gamma, \mathcal{J}, p, q), \alpha)$ is a shorthand for the instance obtained from $(\Gamma, \mathcal{J}, p, q)$ by enlarging a hole of Γ along α and CUT$((\Gamma, \mathcal{J}, p, q), \alpha)$ is a shorthand for the instance obtained from $(\Gamma, \mathcal{J}, p, q)$ by cutting through α. We now show the correctness of Algorithm 1.

Theorem 5.19. *Let* $(\Gamma, \mathcal{J}, p, q)$ *be an instance of the path problem with at most two holes that meets properties (i)–(iii). Then, Algorithm 1 decides whether there is a valid pq-arc in polynomial time in the number of arcs in* $\mathcal{J}$*.*

Proof. Step 1 primarily checks if our current instance $(\Gamma, \mathcal{J}, p, q)$ is trivial (i.e. $\mathcal{J} = \emptyset$). If not, the algorithm moves towards Step 2, where it verifies if $\mathcal{J}$ has an arc of Type T3. If it has one, it uses this arc to enlarge a hole and applies Lemma 5.14 to update our instance to one with fewer arcs.

Otherwise, if $\mathcal{J}$ has no arcs of Type T3, the process continues with Step 4. The first possibility is that all arcs in $\mathcal{J}$ are of different colors, and in this case the conditions of Lemma 5.18 apply, so there is a valid pq-arc (Steps 5–6).

The second possibility is that $\mathcal{J}$ has two arcs α and α' of the same color. If both α and α' are pq-separating, then clearly no valid pq-arc exists (Steps 9–10). Otherwise, one of them, say α', is of Type T1. If α is pq-separating (Type T2), then any valid pq-arc must cross α, and thus, it does not cross α'. Therefore, it is enough to look for a valid pq-arc not crossing α'. Observation 5.16 translates that into finding a valid pq-arc for the instance with one hole that we obtain with the operation $\mathrm{CUT}((\Gamma, \mathcal{J}, p, q), \alpha')$ (Steps 11–12).

The third and last alternative is that both α and α' are of Type T1. In this case, any valid pq-arc crosses only one of α and α'. Thus, by Observation 5.16, it is enough verify both instances obtained by applying the transformations $\mathrm{CUT}((\Gamma, \mathcal{J}, p, q), \alpha)$ and $\mathrm{CUT}((\Gamma, \mathcal{J}, p, q), \alpha')$ (Steps 13–14). An attentive reader may notice how, in principle, an iterative occurrence of Step 14 may lead into an exponential blow-up of the running time. However, the fact that both instances that we obtain applying the transformations $\mathrm{CUT}((\Gamma, \mathcal{J}, p, q), \alpha)$ and $\mathrm{CUT}((\Gamma, \mathcal{J}, p, q), \alpha')$ are instances of the path problem with one hole, guarantees that the algorithm goes through Step 14 at most once. □

5.7. FPT-algorithm for bounded number of crossings

In this section we show that for simple drawings with a bounded number of crossings it can be decided in FPT-time[2] if an edge can be inserted.

Theorem 5.20. *Given a simple drawing $D(G)$ of a graph $G = (V, E)$ and an edge uv of the complement of G, there is an* FPT*-algorithm in the number k of crossings in $D(G)$ for deciding whether uv can be inserted into $D(G)$.*

Proof. Let $n = |V|$ and $m = |E|$ be the number of vertices and edges of G, respectively. We consider a subdrawing $D'(G')$ of $D(G)$ consisting of the edges incident to u and v, the (at most $2k$) edges which are crossed, and the vertices incident to all these edges together with u and v if they are isolated (and thus, not yet included). We first show that we can decide in FPT-time in k whether uv can be inserted into $D'(G')$. We then argue that uv can be inserted into $D'(G')$ if and only if it can be inserted into $D(G)$.

As we saw in Section 5.4, the problem of extending $D(G)$ with one edge uv is equivalent to the existence of a heterochromatic path in $G^*(D, \{u, v\})$ between a vertex corresponding to a cell incident to u and a vertex corresponding to a cell incident to v.

[2]An FPT-algorithm for a problem with respect to the parameter k runs, for an instance of size n, in time $O(f(k) \cdot n^c)$, where f is a computable function that does not depend on n and c is an absolute constant.

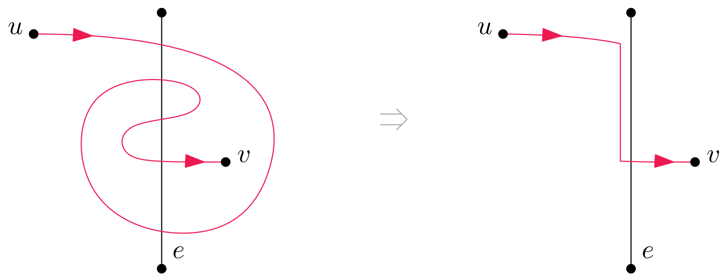

Figure 5.15.: Rerouting uv when it crosses an edge uncrossed in $D(G)$ more than once.

The number of segments of crossed edges in $D'(G')$ is at most $4k$. Thus, $G^*(D', \{u, v\})$ has at most $4k$ edges (but the number of vertices might not be bounded by a function of k). There are $O(n)$ cells with u on their boundary, and we consider the vertices of $G^*(D', \{u, v\})$ corresponding to them all as the possible starting cells of a valid heterochromatic path. Thus, the algorithm checking whether uv can be inserted into $D'(G')$ runs in $O(nk2^{4k})$ time.

We now argue that uv can be inserted into $D'(G')$ if and only if it can be inserted into $D(G)$. Since $D'(G')$ is a subdrawing of $D(G)$, it is clear that if uv cannot be inserted into $D'(G')$ then it cannot be inserted into $D(G)$. Suppose that uv can be inserted into $D'(G')$ and let γ be a *valid arc* for uv in $D'(G')$ resulting in a simple drawing. We orient γ from u to v. If γ is not a valid arc for uv in $D(G)$ then it must cross more than once an edge e uncrossed in $D(G)$. We can modify γ such that it is routed close to e between its first and last crossings with e, producing at most one intersection; see Figure 5.15 for an illustration. Repeating this process for every edge uncrossed in $D(G)$ and crossed by γ more than once we obtain a valid arc for uv in $D(G)$. □

5.8. Chapter summary

In this chapter we have shown that, given a simple drawing $D(G)$ of a graph $G = (V, E)$ and two non-adjacent vertices $u, v \in V$, it is NP-complete to decide whether uv can be inserted into $D(G)$. Moreover, the problem remains NP-complete when G is a matching. We presented a polynomial algorithm for the case in which $\{u, v\}$ is a dominating set for G and an FPT-algorithm in the number of crossings of $D(G)$.

6. Extending 1-plane drawings

An extended abstract containing the results in this chapter has appeared in [131].

6.1. Introduction

This chapter, as the previous one, studies extensions of drawings with certain properties. In general, the (maximization version of this) problem is the following: Given a drawing $D(G)$ of a graph $G = (V, E)$ that has a set of properties, and a set of *candidate edges* M of the complement graph $\overline{G}$ of G, find a maximum subset $M' \subseteq M$ that can be *inserted* into $D(G)$ such that the result is a drawing $D'(G^+)$ of the graph $G^+ = (V, E \cup M')$ with a desired set of properties and containing $D(G)$ as a subdrawing. For a broader introduction on this type of problems, we refer the reader to the introduction of Chapter 5.

Here we focus on maximal extensions of 1-plane drawings. We call a graph G *1-planar* if there exists a simple drawing $D(G)$ in which no edge is crossed more than once. Such a drawing is called *1-plane*. The class of 1-planar graphs is widely studied in graph theory and graph drawing in the context of so called "beyond planarity graphs" [77]. For general results on 1-planar graphs see the annotated bibliography by Kobourov et al. [132]. Recognizing if a given graph is 1-planar is NP-hard [107] and it stays hard even if the graph consists of a planar graph plus one edge [59]. It also stays hard if the rotation system is fixed [39] and it is also NP-hard for graphs of bounded bandwidth [44].

Note that from these results it follows that deciding whether we can insert all the edges such that the result is a 1-plane drawing is NP-hard, since the vertices in a simple drawing can be arbitrarily placed. (And if the decision version asking whether a set of edges can be inserted is NP-hard, then the maximization version asking how many edges can be inserted is NP-hard too.) However, the problem we study here includes the natural assumption that the initial drawing is connected. The proofs showing that recognition of 1-planarity is hard [39, 44, 59, 107, 134] cannot be directly applied to this setting, since they all rely on the fact that certain vertices can be freely placed. Nevertheless, we believe that the reduction presented in [39] can be adapted to our setting.

Outline. Given a (1-)plane drawing $D(G)$ of a graph G and candidate set of edges M of $\overline{G}$, we study the problem of deciding whether there is a subset $M' \subseteq M$ of cardinality k that can be inserted into $D(G)$. In Section 6.2 we show with a simple reduction that this problem is NP-complete, even if $D(G)$ is plane, connected, and orthogonal. In Section 6.3 we present an

FPT-algorithm with respect to k for deciding whether k edges in $\overline{G}$ can be inserted. Finally, we present a summary and conclusions in Section 6.4.

6.2. NP-hardness

In this section we prove that extending a connected (1-)plane drawing by a maximum set of given candidate edges is NP-hard, even if the initial drawing is plane and orthogonal.

Our reduction, as the one in Section 5.3 in Chapter 5, is from the maximum independent set problem and the intuition behind it is similar. However, in this case we require the input graph to be planar and with maximum vertex degree three. (In the previous chapter we did not require the input graph to be planar.) For planar graphs the maximum independent set problem is NP-hard, but not hard to approximate (there is a PTAS). Moreover, the gadgets we use are substantially different. When extending simple drawings, the gadgets profile the ways in which an edge can be inserted, and rely on edges that can cross multiple other edges. In our case, the gadgets are plane, straight-line drawings, and instead, rely on pairs of edges that cannot be inserted simultaneously. We also give an FPT-time algorithm in the size of the set M', when M is the set of all edges in $\overline{G}$. Our algorithm works even if the initial drawing is not connected and 1-plane.

Theorem 6.1. *Given a connected plane drawing $D(G)$ of a graph G, an integer k, and a subset M of the edges of the complement graph of G, it is* NP*-complete to decide whether there is a subset $M' \subseteq M$ of cardinality k extending $D(G)$ to a 1-plane drawing.*

Note that the problem is in NP, since it can be encoded combinatorially. To prove Theorem 6.1 we reduce from (a variant of) maximum independent set (MIS). Recall that a set of vertices of a graph is an *independent set* if no pair of vertices in the set are adjacent. The problem of determining the maximum independent set of a given graph is NP-hard in general, even when the given graph is planar and has degree at most three [101, Lemma 1].

Planar graphs with degree at most three admit a 2-page book embedding [51, 118]. A *2-page book embedding* is a plane drawing in which all vertices are placed on a horizontal line, the *spine* of the book, and the edges lie completely either in the upper or lower half-plane. Thus, we can construct a 2-page book embedding $D(G)$ from an MIS instance consisting of a graph G with degree at most three. By replacing the vertices of $D(G)$ with vertex gadgets we construct a plane drawing $D'(G')$ of a graph $G' = (V', E')$. Then the edges in each half-plane of $D(G)$ define a set of edges M of the completement graph of G' such that finding a maximum subset of edges $M' \subseteq M$ extending $D'(G')$ to a 1-plane drawing is equivalent to finding a maximum independent set of G.

Vertex gadget. We now introduce the main gadget for our reduction. The vertex gadget is a orthogonal plane drawing symmetric with respect to a horizontal line (the spine) by a 180° rotation.

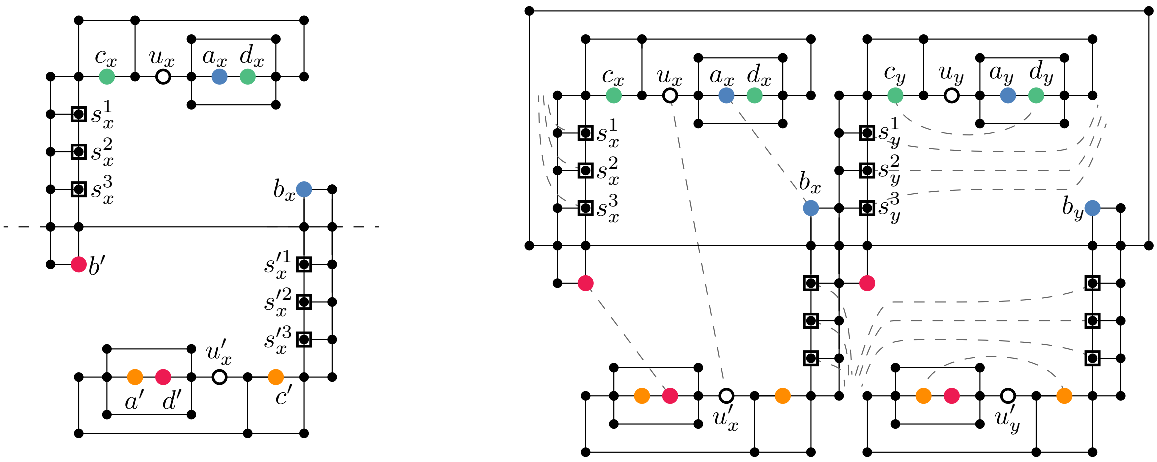

Figure 6.1.: Reduction. Left: Vertex gadget. Right: Concatenation of two vertex gadgets and possible extensions.

For one gadget, there are four vertices placed on the spine and connected by a path. Assume the leftmost one is placed at (0,0). The other three are then placed at (1,0), (9,0), and (10,0), respectively. The two rightmost vertices form a 4-cycle together with two vertices placed at (9,1) and (10,1). We denote by b the vertex at (9,1). The two leftmost vertices on the spine are the bottom two vertices of a 2 by 5 grid that includes all unit-length edges. We denote the vertices at (1,3), (1,2), and (1,1) by s^1, s^2 and s^3, respectively. The top right vertex of that grid, drawn at (1,4), is the leftmost vertex of a horizontal path with nine vertices and unit-length edges. We denote the vertices at (2,4), (4,4), (6,4), and (7,4) by c, u, a, and d, respectively. Finally, we insert the four following paths (and vertices, when not already inserted) into our drawing: (1,4)-(1,6)-(3,6)-(9,6)-(9,4), (3,6)-(3,4), (5,4)-(5,5)-(8,5)-(8,4), and (5,4)-(5,3)-(8,3)-(8,4). The bottom part of the drawing is a copy of the top half, rotated by 180° such that the leftmost vertex on the spine before the rotation corresponds to the rightmost one of the top part; see Figure 6.1 for an illustration. For any specially named vertex in the top half we add an apostrophe to the name in the bottom-half.

Concatenating vertex gadgets. Let $D(G)$ be a 2-page book embedding of a graph $G = (V, E)$ with max-degree three. For each vertex $u \in V$ we create a vertex gadget at the position of u as above. We insert the vertex label of the corresponding vertex in V as a subscript to the named vertices in the gadget. From left to right we connect the gadgets as follows. For two consecutive gadgets on the spine, we identify the rightmost vertex on the spine of the first vertex gadget with the leftmost vertex on the spine of the second vertex gadget. In that way, we join all vertex gadgets into one plane drawing. Finally, we insert an orthogonal path connecting the leftmost vertex on the spine with the rightmost one such that it surrounds the top-half of the drawing; see Figure 6.1 (right). We refer to this path as the *surrounding path*.

Candidate edges. For each vertex in the 2-page-book drawing $D(G)$, we sort the incident (at most three) edges in the top half-plane in clockwise order. Consider an edge xy drawn in

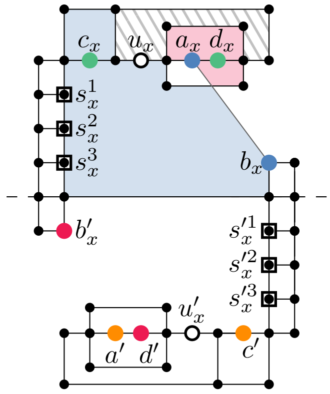

(a) Proof of Claim 1.

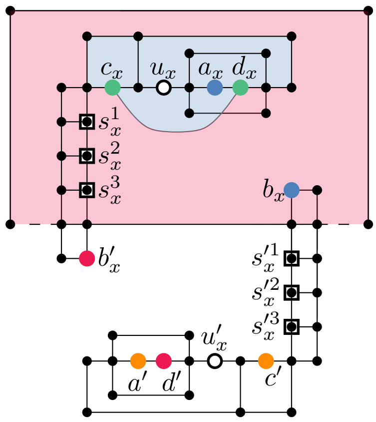

(b) Proof of Claim 2.

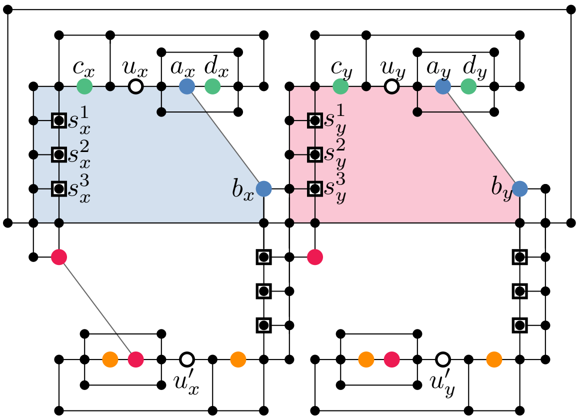

(c) Proof of Claim 3.

Figure 6.2.: Closed curves separating the vertices that we aim to connect.

the top half-plane of $D(G)$. Without loss of generality, let this edge be the i-th one incident to x and the j-th one incident to y, we then add the *candidate* edge $s_x^i s_y^j$ to M. For an edge in the bottom half-plane of $D(G)$ we proceed analogously. After adding all candidate edges for the top and bottom half-planes, we add for all $v \in V$ the edges $a_v b_v$, $a'_v b'_v$, $c_v d_v$, $c'_v d'_v$, and $u_v u'_v$ to M.

Lemma 6.2. *The construction is a polynomial-time reduction from maximum independent set in planar graphs with degree at most three.*

Proof. We begin by showing in the three following claims which of the candidate edges can be inserted simultaneously into the vertex gadgets. The proofs rely on the following strategy. To show that an edge cannot be inserted, we identify two simple closed curves whose intersection consists at most of vertices and intersection points of the drawing. These curves strictly separate the endvertices of the edge we want to insert.

Claim 1. *Edges ab and cd (analogously $a'b'$ and $c'd'$) cannot be inserted simultaneously.*

Proof. Consider the smallest cycle C around a and d in the vertex gadget drawing; see Figure 6.2(a). If the candidate edge ab is inserted, it must cross C since it does separate a from b. Moreover, it must cross one of the three lower edges of C, since if the edge starts from a and crosses one of the three upper edges of C it gets in a region not containing d; see the stripped region in Figure 6.2(a). Now, to see that the edges ab and cd cannot be inserted simultaneously, suppose that ab is inserted. Then there are two closed curves both strictly separating c and d and only intersection in two points. The first one, filled in red in Figure 6.2(a), goes from a to its left neighbor along the edge, then follows clockwise the boundary of C until its intersection with the inserted edge ab, and finally follows ab until it reaches a. The second one, filled in blue in Figure 6.2(a), goes from b to its bottom neighbor along the edge, then follows the edge on the spine to the left, goes up along five edges, then one to the right, one down, and two to the right, and, finally, follows C counterclockwise until its intersection with ab and ab until b. Thus, ab and cd cannot be both inserted into a vertex gadget. Symmetric arguments show that $a'b'$ and $c'd'$ cannot be both inserted simultaneously. □

Claim 2. *If the edge uu' is inserted into a horizontal concatenation of vertex gadgets (including the surrounding path), neither cd nor $c'd'$ can be inserted.*

Proof. Suppose cd is inserted. Then there are two non-crossing closed curves both strictly separating u and u'. The first one, filled in red in Figure 6.2(b), follows all the edges along the spine and the surrounding path. The second one, filled in blue in Figure 6.2(b), goes from the left neighbor of c to c along the edge, then follows edge cd until d, goes to the right along two edges, then one up, two to the left and one down. Notice that since both c and d both lie (strictly) on the same side of the first curve, then the edge cd is also completely on that same side, and the two curves are disjoint. If instead of cd we assume that $c'd'$ is inserted, then the second curve is symmetric to the one described and again does not intersect the first one. □

Claim 3. *Consider a horizontal concatenation of vertex gadgets including one with x as subscript and one with y as subscript. Then, the edge $s_x^i s_y^j$ cannot be inserted if both $a_x b_x$ and $a_y b_y$ are inserted.*

Proof. Assume both $a_x b_x$ and $a_y b_y$ are inserted. As in the proof of Claim 1, in a vertex gadget edge ab, if inserted, must cross the smallest cycle around a and d. Thus, we know the cycles where $a_x b_x$ and $a_y b_y$ spend their only crossing. Moreover, there are two disjoint cycles strictly separating s_x^i and s_y^j. The first one, filled in blue in Figure 6.2(c), follows the edge $a_x b_x$ from a_x to b_x and then goes down one edge, two to the left, four up, and six to the right until a_x. The second one, filled in red in Figure 6.2(c), is the analogous one in the other vertex gadget. □

With these three claims at hand, we proceed to tackle Lemma 6.2. Let $G = (V, E)$ be a planar graph with degree at most three and $k \in \mathbb{N}$. We reduce from the problem of deciding whether G has an independent set of size k. First, we construct a plane drawing $D'(G')$ from G and a set M of candidate edges as explained above. Note that this set consists of $5|V| + |E|$ edges: $a_v b_v$, $a'_v b'_v$, $c_v d_v$, $c'_v d'_v$, and $u_v u'_v$ for each $v \in V$ and one $s_x^i s_y^j$ or $s_x'^i s_y'^j$ with $i, j \in \{1, 2, 3\}$ per edge $xy \in E$. Moreover, it is a polynomial construction.

The problem we want to reduce to is deciding whether $D'(G')$ can be extended to a 1-plane drawing by inserting a set of edges $M' \subseteq M$ with cardinality $|M'| = |E| + 2|V| + k$. The rest of the proof goes along the lines of the proof of Theorem 5.3 in the previous chapter.

First, we show that if G has an independent set I of size k we find a subset $M' \subseteq M$ of the candidate edges of size $|E| + 2|V| + k$. Consider Figure 6.1 (right). For every vertex $x \in I$ we insert the edges $a_x b_x$, $a'_x b'_x$, and $u_x u'_x$ as shown in the left half of the figure. For every vertex $y \in V \setminus I$ we insert the edges $c_y d_y$ and $c'_y d'_y$ as shown in the right half. Finally, we insert all candidate edges of the form $s^i_x s^j_y$ and $s'^i_x s'^j_y$ with $i, j \in \{1, 2, 3\}$ and $xy \in E$ to M'. The way we draw a candidate edge $s^i_x s^j_y$ when inserting it depends on whether x or y are part of the independent set I. In general we draw these segments with three arcs: The first and last ones locally around the endpoints, respectively, and the middle one as a (deformed) arc of the 2-page-book drawing $D(G)$; see Figure 6.1 (right). Since I is an independent set, each edge has at most one endpoint in I. Thus, each edge $s^i_x s^j_y$ and $s'^i_x s'^j_y$ with $i, j \in \{1, 2, 3\}$ and $xy \in E$ that we inserted into the drawing has at most one crossing. Therefore, we obtained a 1-plane extension with $|E| + 2|V| + k$ edges.

Conversely, let $M' \subset M$ be a set of $|E| + 2|V| + k$ candidate edges that can be inserted into $D'(G')$ and that contains the minimum possible amount of uu' edges. We then find an independent set of size k of G in the following way: By Claim 1, at most $2|V|$ candidate edges of the form ab and cd can simultaneously be inserted. Thus, at least k of the inserted candidate edges are uu' edges. Therefore, if the set of vertices $\{v : u_v u'_v \in M'\}$ is an independent set of G we are done. Assume on the contrary that $\{v : u_v u'_v \in M'\}$ is not an independent set. Then there exists an edge $xy \in E$ such that $u_x u'_x$ as well as $u_y u'_y$ were inserted. By Claim 3 $a_x b_x$ or $a_y b_y$ cannot be inserted. Moreover, by Claim 2, neither $c_x d_x$ nor $c_y d_y$ were inserted. If $a_x b_x$ is not in M', we could remove $u_x u'_x$ from M' and add $c_x d_x$. That edge can always be inserted by drawing it below (and as close as needed to) the horizontal path from c_x to d_x, since neither u_x nor a_x have an incident edge in M' after $u_x u'_x$ was removed, and all other edges cannot enter a simple closed curve that they would have to leave. Symmetrically, if $a_y b_y$ is not in M' we could remove $u_y u'_y$ from M' and add $c_y d_y$ to M'. This contradicts the fact that M' has the minimum amount of uu' edges. □

6.3. FPT for inserting arbitrary edges

In this section we show that for a 1-plane drawing $D(G)$ of a graph G one can decide in FPT-time[1] in k if there exists a set of k edges in $\overline{G}$ that extend $D(G)$ to a 1-plane drawing.

We prove this result using a series of technical lemmas and observations. The goal is to obtain conditions checkable in polynomial time that can lead to a positive answer, and that, if not met, imply a bound that is polynomial in k on for the size of the structures where edges can be non-trivially inserted.

[1] We recall that an FPT-algorithm for a problem with respect to the parameter k runs, for an instance of size n, in time $O(f(k) \cdot n^c)$, where f is a computable function that does not depend on n and c is an absolute constant.

For a 1-plane drawing $D(G)$ of a graph G we construct a plane drawing by placing a vertex on every intersection point of two edges in $D(G)$ and consider the faces of this planarized drawing as the *cells* of the 1-plane drawing $D(G)$. Note that every cell has some vertex of $D(G)$ on its boundary.

Observation 6.3. *Let $D(G)$ be a 1-plane drawing of a graph G. Assume there is an edge e of $\overline{G}$ which can be drawn in one cell (two adjacent cells) of $D(G)$ in a 1-plane way such that no other edge of $\overline{G}$ can be inserted into $D(G)$ intersecting the interior of that cell (either of the two cells). Then, if $D(G)$ can be extended with k edges, there is an extension with k edges in which e is drawn into that cell (those two cells).*

The first lemma considers a case in which we can insert k edges into a single cell of $D(G)$.

Lemma 6.4. *Let $D(G)$ be a 1-plane drawing of a graph G. If $D(G)$ contains a cell with at least $6k+1$ vertices on its boundary, then we can extend $D(G)$ by k edges to a 1-plane drawing.*

Proof. It suffices to find a vertex v on the boundary of the cell that has at least k non-neighbors (excluding the vertex itself) on the same boundary. If such a vertex exists we can extend $D(G)$ by the edges between v and all its non-neighbors on the boundary of the cell without introducing new crossings.

Assume such a vertex does not exist and denote by V_B the set of vertices of G on the boundary of the cell under consideration and by E_B the edges of G between them, so $|E_B| \geq \frac{|V_B|\cdot(|V_B|-k)}{2}$. At the same time, as G is 1-planar, every subgraph of G is K_7-free [133]. By Turán's Theorem [191] the number of edges in a K_7-free graph with n vertices is at most $\frac{5}{12}n^2$. This means that $|E_B| \leq \frac{5}{12}(|V_B|)^2$. Calculation yields that $\frac{|V_B|\cdot(|V_B|-k)}{2} > \frac{5}{12}(|V_B|)^2$ if $|V_B| > 6k$, which is the case by the choice of the cell, and we obtain a contradiction. $\square$

Now we consider cases in which we can extend drawings within many, possibly small cells.

Observation 6.5. *Let $D(G)$ be a 1-plane drawing of a graph G. If $D(G)$ contains at least k cells, each with the endpoints of a distinct edge of $\overline{G}$ on its boundary, we can extend $D(G)$ by these k edges.*

However, it might be necessary to introduce new crossings when extending $D(G)$. For this reason, we consider pairs of adjacent cells whose shared boundary can be crossed by an inserted edge. If one can find sufficiently many disjoint such pairs, one can avoid introducing crossings between new edges. The main challenge is to enforce disjointness of the pairs and distinctness of the inserted edges at the same time.

Formally, we are interested in pairs of cells such that their shared boundary can be crossed in a 1-plane extension of $D(G)$. We call these pairs *crossable pairs*. For these pairs we consider edges of $\overline{G}$ whose insertition into $D(G)$ in the crossable pair requires crossing the shared boundary, which we call its *edge options*. Observe that (i) in a crossable pair, both boundaries of the cells share at least two vertices in V, as their boundaries share at least an edge and (ii) edge options have a vertex on the boundary of each cell that is not on the common boundary. We remark that the common boundary might not be connected, as it might consist of different connected components.

Lemma 6.6. *Let $D(G)$ be a 1-plane drawing of a graph G. If $D(G)$ contains at least k interior-disjoint crossable pairs, each pair having at least two edge options, then we can insert k edges into $D(G)$.*

Proof. Let $\mathcal{F}_2$ be the crossable pairs with at least two edge options and $\mathcal{F}_d \subseteq F_2$ a maximum set of interior-disjoint crossable pairs in $\mathcal{F}_2$. For two crossable pairs $F_1, F_2 \in \mathcal{F}_d$ we say $F_1 \neq F_2$ if they share no cell. We show the lemma via induction on the size of $\mathcal{F}_d$.

In case $|\mathcal{F}_d| = 1$ or $|\mathcal{F}_d| = 2$ the statement holds immediately. For $|\mathcal{F}_d| > 2$, let $F = (f_1, f_2) \in \mathcal{F}_d$ be a crossable pair with at least two edge options consisting of two adjacent cells f_1 and f_2 with edge options x_1y_1 and x_2y_2. Without loss of generality $x_1 \neq x_2$, and both x_1 and x_2 lie on the part of the boundary of f_1 that is not the common boundary of f_1 and f_2, see Figure 6.3 for an illustration.

If we find that F is the only crossable pair in $\mathcal{F}_d$ for which x_1y_1 is an edge option, we insert x_1y_1 in F into $D(G)$ and do not remove an edge option from any other crossable pair in $\mathcal{F}_d$. Symmetrically for x_2y_2.

Furthermore, let $v_1, \dots v_z \in V$, $z \geq 2$ be the vertices on the shared boundary of the cells f_1, f_2. If $z > 2$ and since f_1 and f_2 are cells, there exists an edge v_ix_j with $i \in \{2, \dots, z-1\}, j \in \{1, 2\}$ in $\overline{G}$. Assume this edge was an edge option for some $F' \in \mathcal{F}_d$, $F' \neq F$. Then the cells of F' would require to have v_i as well as x_j on its boundaries. It follows that the boundary of the cells in F' would subdivide f_1 or f_2 or both, but those are cells. Therefore, we can insert v_ix_j into $D(G)$ in (f_1, f_2) without removing an edge option from any other crossable pair. Symmetrically for v_iy_j with $i \in \{2, \dots, z-1\}, j \in \{1, 2\}$.

In the following, assume we could not handle F with either of the two easy cases from above. Then we can assume each edge option of F is also an edge option for at least one other $F' \in \mathcal{F}_d$, $F \neq F'$, and the shared boundary between f_1 and f_2 is only one edge. Let v_sv_t be that edge on the boundary between f_1 and f_2 and x_1v_t, x_2v_s such that they would cross if both drawn inside f_1. It remains to consider three cases, illustrated in Figure 6.3:

Case 1: Both, x_1v_t and x_2v_s, are edge options for crossable pairs $F', F'' \in \mathcal{F}_d$ with $F' \neq F$ and $F'' \neq F$.

Case 2: Without loss of generality, x_1v_t is drawn in $D(G)$ while x_2v_s is edge option of one $F' \in \mathcal{F}_d$ with $F' \neq F$.

Case 3: x_1v_t and x_2v_s are both drawn in $D(G)$.

For Case 1 let $F', F'' \in \mathcal{F}_d$ be crossable pairs such that $F' \neq F$ and $F'' \neq F$, x_1v_t is an edge option for F', and x_2v_s is an edge option for F''. Observe, that two edge options of two interior-disjoint crossable pairs, can both be inserted into $D(G)$ without intersecting each other. Furthermore, the cells of F' and F'' lie outside of the cells of F, hence if x_1v_t and x_2v_s are inserted in F' and F'' into $D(G)$, their images cross. By the previous observation, this can only

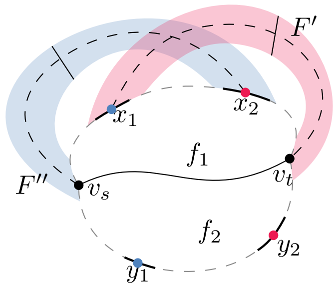

(a) Case 1.

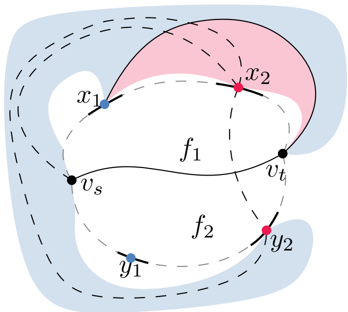

(b) Case 2.

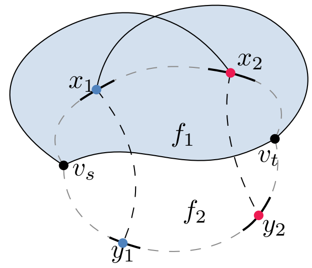

(c) Case 3.

Figure 6.3.: An illustration of Lemma 6.6. The colored vertices are the endpoints of the two edge options x_1y_1 and x_2y_2.

be the case if $F' = F''$. It follows that we can simply insert x_1v_t in F and x_2v_s in F' without removing an edge option from any other crossable pair in $\mathcal{F}_d$.

In Case 2 we assume that, without loss of generality, x_1v_t is already drawn in $D(G)$. Since f_1 is a cell we know that x_1v_t is not drawn inside f_1. Let $F' \in \mathcal{F}_d$ be a crossable pair with $F \neq F'$ and x_2v_s is an edge option of F'. Note that any drawing of x_2v_s which is not inside f_1, f_2 in $D(G)$ has to cross the image of x_1v_t in $D(G)$ by definition, it follows that x_1v_t is the shared boundary of the cells in F'.

Now, consider the edge option x_2y_2 of F. By assumption there exists a crossable pair $F'' \in \mathcal{F}_d$ with $F \neq F''$ such that x_2y_2 is an edge option for F''. Furthermore, any drawing of x_2y_2 that is not inside f_1, f_2 in $D(G)$ has to cross the image of x_1v_t in $D(G)$. It follows that F' and F'' have the same shared boundary and are hence the same crossable pair. Thus, inserting x_2v_s in F and x_2y_2 in F' does not remove any edge option from any other crossable pair in $\mathcal{F}_d$.

It remains to consider Case 3. Here, edges x_1v_t and x_2v_s are already drawn in $D(G)$. Since f_1 is a cell, both are drawn completely outside of f_1. Consider the simple closed curve that is defined as follows: Starting from v_s it follows the edge v_sx_2 until its intersection with the edge v_tx_1, then it follows v_tx_1 until v_t, and finally it follows the edge v_sv_t until v_s. This curve separates x_1 from y_1; see Figure 6.3 for an illustration.

The only part of that curve that an inserted edge can cross is the one corresponding to the edge v_sv_t, since v_sx_2 and v_tx_1 already cross. It follows that to connect x_1 to y_1 we must use the interior of (f_1, f_2). Moreover, since v_sx_2 is already crossed, x_1y_1 cannot be an edge option for any other crossable pair. Swapping the roles of v_s with v_t and x_1 with x_2 we can argue symmetrically that also x_2y_2 is unique to (f_1, f_2). It follows that both edge options can be inserted in F into $D(G)$ without removing an edge option from any other crossable pair in $\mathcal{F}_d$. $\square$

Under the condition that no edge as in Observation 6.3 exists in $D(G)$, the next lemma bounds the number crossable pairs in which we can find a particular edge option. This lemma can then be used to bound the number of crossable pairs with one edge option when Observation 6.3 does not apply.

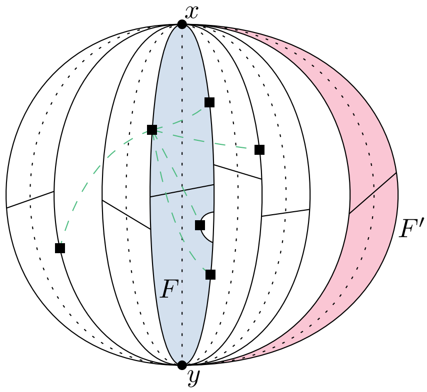

Figure 6.4.: Illustration of the first part of Lemma 6.7. The edge option e shared by at least $3k$ crossable pairs is xy. The green edges illustrate possible edge options e' for the crossable pair F.

Lemma 6.7. *Let $D(G)$ be a 1-plane drawing of a graph $G = (V, E)$ whose complement is $\overline{G} = (V, \overline{E})$. If $e = xy \in \overline{E}$ is an edge option for at least $36k^2$ crossable pairs and, for each of these crossable pairs, at least one of their cells either allows the insertion of an edge in $\overline{E} \setminus \{e\}$ in a plane manner or is in a crossable pair with another edge option, then we can insert k edges into $D(G)$.*

Proof. Let $\mathcal{F}$ be the set of crossable pairs for which e is an edge option, and let $\mathcal{F}_d \subseteq \mathcal{F}$ be a maximum set of interior-disjoint crossable pairs.

If $|\mathcal{F}_d| \geq 3k$, we order the crossable pairs in $\mathcal{F}_d$ cyclically in a counterclockwise manner around x. Let $\mathcal{F}'_d = \{F_1, \ldots, F_k\}$ be the set that we obtain by choosing every third crossable pair with respect to that order. (That is, we pick the third, sixth, ..., until the $3k$-th crossable pair in the cyclic order.) Then, for every two crossable pairs F_i and F_j with $i \neq j$, every point on the boundary of the cells of F_i except x and y is separated from every point on the boundary of the cells of F_j except x and y by at least two closed curves of edges in $D(G)$.

We now describe how to find an edge to insert for every $F \in \mathcal{F}'_d$. If there is an edge $e' \neq e$ of $\overline{G}$ that can be inserted in a plane manner within one of the cells of F, then we can safely insert this edge. Note that this edge, by construction of $\mathcal{F}'_d$, cannot be inserted intersecting cells of some $F' \in \mathcal{F}'_d \setminus \{F\}$, neither in a plane manner, nor as an edge option involving any such cell, because at least one endpoint of e' is separated by two closed curves from the boundary of the cells of F'.

If there is no such edge, by the condition of the lemma, there is an edge option $e' \in \overline{E} \setminus \{e\}$ involving a cell f of F. Using an argument symmetric to the prior case, by construction of $\mathcal{F}'_d$, the edge option e' cannot be inserted within a cell of some $F' \in \mathcal{F}'_d \setminus \{F\}$. Similarly, e' cannot be an edge option for a crossable pair involving a cell of some $F' \in \mathcal{F}'_d \setminus \{F\}$, since at least one of the endpoints of e' that is not an endpoint of e has to lie on the same side of the closed curves separating the boundary of F and F'.

Otherwise, among the $36k^2$ crossable pairs in $\mathcal{F}$, less than $3k$ are interior-disjoint. Thus, at least $6k+1$ crossable pairs in $\mathcal{F}$ share a cell. This cell has at least $6k+1$ edges on its boundary, and hence, more than $6k$ vertices. Thus, by Lemma 6.4, we can insert k edges into $D(G)$. □

We now describe how to check the conditions of Observation 6.3, Lemma 6.4, Observation 6.5, Lemma 6.6, and Lemma 6.7. We start with Observation 6.3. For each cell f in $D(G)$ let E_f be the set of edges in $\overline{G}$ with at least one endpoint on the boundary of f. This set can be computed in $O(n^2)$ time for each cell. Since there are only $O(n)$ many cells in a 1-plane drawing all these sets can be computed in $O(n^3)$.

For each edge uv in $\overline{G}$ we do the following: Iterate the cells incident to u and v. If there is a cell f with $|E_f| = 1$ and both u and v are on the boundary of f, we insert uv into G and $D(G)$, and return. If there are two cells f and g, u incident to f and v incident to g, such that $|E_f| = |E_g| = 1$ and f and g are adjacent, we insert uv into G and $D(G)$, and return. To check Observation 6.3 we execute the above algorithm k times or until no edge can be inserted.

This is correct, since if a set E_f for a cell f has only one element, this element corresponds to the only edge of $\overline{G}$ which is incident to a vertex on the boundary of f. Hence this is the only edge that has a possible drawing intersecting the interior of f. By precomputing for each cell f the set of cells f is adjacent to in $D(G)$, the above algorithm runs in $O(n^4)$. Consequently, Observation 6.3 can be checked in $O(kn^4)$ time.

The condition of Lemma 6.4 can easily be checked by traversing the boundaries of all cells of $D(G)$ in polynomial time.

The conditions of Observation 6.5 and Lemma 6.6 can be modeled as cardinality maximum matching problems. For Observation 6.5: Consider the graph H whose vertex set is the union of the edge set of $\overline{G}$ and the set of cells of $D(G)$, and whose edges connect any vertex corresponding to a cell to all vertices corresponding to edges whose endpoints lie on the boundary of the cell. Note that H can be obtained in polynomial time from $D(G)$. It is straightforward to see that H has a matching of size $\geq k$ if and only if the condition of Observation 6.5 holds for $D(G)$. Since H has $O(n^2)$ many vertices and, with Lemma 6.4 ,$O(nk^2)$ many edges we can check if such a matching exists in $O(n^2k^2)$ time using the Hopcroft-Karp algorithm [121].

To check the conditions of Lemma 6.6 we need to check if there are k interior-disjoint crossable pairs $\mathcal{F}_d$ with at least two edge options. Let $\mathcal{F}_2$ be the set of all crossable pairs with at least two edge options. This set has size $O(n^2)$ and can be computed in $O(n^2k^2)$. Consider the graph H where for each cell in $D(G)$ we have a vertex and two vertices are connected by an edge if the corresponding cells form a crossable pair in $\mathcal{F}_2$. This graph has $O(n)$ many vertices and, since every cell has size bounded in $O(k)$ by Lemma 6.4, $O(nk)$ many edges. Clearly a matching of size k in H implies the existance of a set of interior-disjoint crossable pairs of size k such that each crossable pair in this set has at least two edge options. Vice versa the existence of such a set implies the existence of a matching of equal size in H. It follows that H has a matching of size k if and only if there exists a set $\mathcal{F}_d$ of interior-disjoint crossable pairs with at least two edge options each with $|\mathcal{F}_d| \geq k$. Using the algorithm by Micali and Vazirani [147] we can check the condition of Lemma 6.6 in $O(nk\sqrt{n})$ time.

It follows that the conditions of Observation 6.3, Lemma 6.4, Observation 6.5, Lemma 6.6, and Lemma 6.7 can be checked in $O(n^4 k)$ time and space.

Our final theorem ties the above observations and lemmas together.

Theorem 6.8. *Given a 1-plane drawing $D(G)$ of a graph G it is* FPT *in k to find a subset of k edges of the complement graph that extend $D(G)$ to a 1-plane drawing.*

Proof. We start by checking if one can insert edges that do not interfere with the insertition of any other edge as in Observation 6.3. This can be done in polynomial time. If we could insert k such edges we are done, otherwise, assume that we have inserted k_0 of them into G and $D(G)$. We want to decide whether we can insert another $k - k_0$ edges. We update k to this new value. Note that after application of Observation 6.3, the second precondition of Lemma 6.7 is always fulfilled.

As argued above, the conditions of Lemma 6.4, Observation 6.5, Lemma 6.6, and Lemma 6.7 can also be checked in polynomial time. In case any of them can be applied they yield polynomial time constructions for a valid set of k edges of $\overline{G}$ that can be inserted into $D(G)$.

If we did not find in that way k edges that can be inserted, we know that:

(i) cells have at most $6k$ vertices in their boundaries,

(ii) there are at most $k - 1$ cells in each of which a distinct edge of $\overline{G}$ can be inserted,

(iii) there are at most $k - 1$ interior-disjoint crossable pairs, each pair having at least two edge options, and

(iv) each edge option is in at most $36k^2 - 1$ crossable pairs.

A 1-planar graph on n vertices has at most $4n - 8$ edges [161]. By property (i), in a cell we can insert at most $24k - 8$ edges. Thus, by the pigeonhole principle, if there are at least $(k - 1)(24k - 8) + 1$ edges that can be inserted into $D(G)$ without intersecting $D(G)$, then those edges belong to at least k different cells. Considering exactly one edge per cell contradicts property (ii). Thus, there are at most $(k - 1)(24k - 8) = 24k^2 - 32k + 8$ edges that can be inserted into $D(G)$ without intersecting any edge in $D(G)$.

By property (i), a cell can be involved in at most $6k$ crossable pairs. We can apply this to, by property (iii), the at most $2k - 2$ cells of a set of disjoint crossable pairs with at least two edge options. Therefore, the number of crossable pairs with multiple edge options is bounded by $12k^2 - 12k$. Furthermore, the number of edge options for one crossable pair can be bounded by $36k^2$ due to property (i). Thus, there are in total at most $(12k^2 - 12k)36k^2 = 432k^4 + 432k^3$ edge options in crossable pairs with multiple edge options.

It remains to bound the number of edge options in crossable pairs with one single edge option each. Assume that there are at least $36k^3 + 12k^2$ crossable pairs with one edge option each. Then, we can insert k edges into $D(G)$ by iteratively choosing one such crossable pair and drawing its edge option. By property (iv), one edge option is in at most $36k^2 - 1$ crossable pairs. Thus, removing a crossable pair and its edge option might lead to a decrease of at most $36k^2 + 12k$ in the number of crossable pairs with one edge option each. Therefore, we can assume that there are in total at most $36k^3 + 12k^2$ edge options in crossable pairs with one edge option.

We have bounded the number of ways in which we can possibly extend $D(G)$ with one edge, that are either inserting an edge not crossing any edge of $D(G)$, or inserting an edge option of a crossable pair. This allows us to branch over all choices, that are at most $432k^4 + o(k^3)$. Computing possible edge drawings can be done in polynomial time in k, using the same method as for checking Observation 6.3.

Since all operations to check the conditions of Observation 6.3, Lemma 6.4, Observation 6.5, Lemma 6.6, and Lemma 6.7 run in $O(n^4k)$, with n being the number of vertices in G, the whole algorithm runs in $O(k^{4k} + n^4k)$. $\square$

6.4. Chapter summary

We showed that the problem of finding a maximum subset of candidate edges for extending connected 1-plane drawings is NP-hard, even if the initial drawing is connected, plane, and orthogonal. Furthermore, we gave an FPT-algorithm in the number of edges to insert, for the case in which the set of candidate edges M consists of all edges in $\overline{G}$. For this special case it remains open if it is NP-hard or polynomial time solvable.

7. Shooting stars in simple drawings of $K_{m,n}$

An extended abstract of the results in Sections 7.2, 7.3, and 7.4 has appeared in [29]. Moreover, the results in Sections 7.3 and 7.4 appeared in the Master's thesis [199], co-supervised by the author of this thesis. Very recently, we obtained the result in Section 7.5. The corresponding manuscript [23] is currently under preparation.

7.1. Introduction

In this chapter we are interested in *plane spanning subdrawings* of a given simple drawing, that is, drawings without crossings that contain all the vertices of the given drawing and a subset of its edges.

The existence of plane subdrawings of simple drawings of the complete graph K_n has received quite a lot of attention. Ruiz-Vargas [178] showed that every simple drawing of K_n contains $\Omega(n^{1/2-\epsilon})$ pairwise disjoint edges for any $\epsilon > 0$, by this improving over many previous bounds [158, 160, 186]. Fulek and Ruiz-Vargas [97] proved that given a simple drawing of K_n, a plane cycle C in the drawing, and any vertex v that is not part of C, at least two edges connecting v to C do not intersect C. Hence, every simple drawing of K_n contains a plane subdrawing with at least $2n-3$ edges. Rafla [171] conjectured that every simple drawing of K_n contains a plane Hamiltonian cycle, a statement that is known to be true for several classes of simple drawings (e.g., 2-page book drawings, monotone drawings, cylindrical drawings), but is still open in the general case. Simple drawings of K_n might not contain a triangulation of the vertices as a subdrawing. In fact, there are simple drawings of K_n with $2n-4$ triangles [114]. The currently best lower bound on the number of empty triangles in simple drawings of K_n is n [24]. For plane trees, Pach et al. [158] proved that every simple drawing of K_n contains a plane drawing of any fixed tree with at most $c \log^{1/6} n$ vertices.

In this chapter we concentrate on the existence of plane spanning trees in simple drawings. For a vertex v in a simple drawing, a *star* rooted at v is a subdrawing that consists of v, a subset of its adjacent vertices, and the edges connecting v to the vertices in this subset. The *star of v* is the star rooted at v including all its adjacent vertices and the edges connecting v to them. Obviously, any simple drawing of the complete graph K_n contains a plane spanning tree: The star of any vertex is one.

For the complete bipartite graph $K_{m,n}$, the situation is not obvious. For straight-line drawings, the existence of plane spanning trees follows from a more general result [124]. However, as a warm

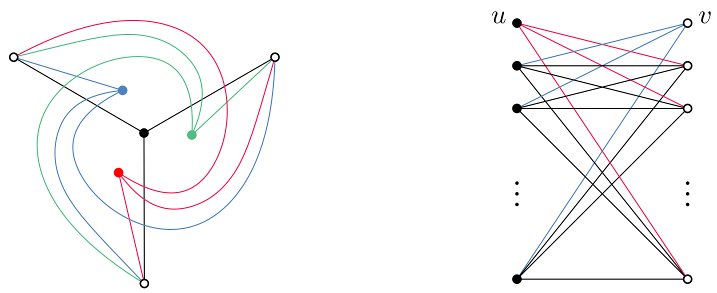

Figure 7.1.: Left: Simple drawing of $K_{3,4}$ by Schröder [182] in which every edge is crossed. Right: Drawing of the complete bipartite graph $K_{m,n}$ minus one edge uv that does not contain a plane spanning tree.

up exercise, let us consider straight-line drawings of $K_{m,n}$ and present a simple construction. Given a straight-line drawing of $K_{m,n}$ with vertex partitions R and B, we pick an arbitrary vertex $r \in R$ and consider its star. For every edge $rb_i, b_i \in B$ in the star of r we draw the ray $\overrightarrow{rb_i}$. These rays partition the plane into wedges, one of which might have an opening angle larger than π. In each wedge, the angle bisector divides the wedge into two parts. We connect the vertices in R lying on each part (the ones on the bisector can be assigned arbitrarily) to the point $b_i \in B$ on the ray that bounds that part of the wedge. These connections together with the star of R and the vertices in $R \setminus r$ define a plane spanning tree in the straight-line drawing of $K_{m,n}$. Actually, the plane spanning tree produced is what we call a *shooting star*. A *shooting star rooted at* v is a plane spanning tree that consists of the star of v and, for each vertex that is not adjacent to v, an edge connecting it to a vertex in the star of v.

For simple drawings of $K_{m,n}$, we are not aware of a similarly easy construction. Simple inductive approaches such as removing an uncrossed edge and one of its incident vertices from the drawing cannot be applied for simple drawings of $K_{m,n}$: The drawing in Figure 7.1 (left) by Schröder [182] is a simple drawing of $K_{3,4}$ in which every edge is crossed at least once. We note that by inserting copies of vertices close to the original ones we can generalize it to a drawing of $K_{m,n}$ with $m \geq 3$, $n \geq 4$ in which every edge is crossed at least once. (For the complete graph K_n with $n \geq 8$, Harborth and Mengersen [117] showed that there are simple drawings in which every edge is crossed.) Moreover, Figure 7.1 (right) shows a drawing of the complete bipartite graph $K_{m,n}$ minus one edge uv that does not contain a plane spanning tree, since every edge in the star of u crosses every edge in the star of v.

Outline. In this chapter we show that every simple drawing of $K_{m,n}$ contains a shooting star rooted at every vertex. In Chapter 7.2 we present a mainly combinatorial proof for simple drawings of $K_{2,n}$ and shooting stars rooted at vertices in the vertex partition of size two. Using this result, we also show the existence of shooting stars in simple drawings of $K_{3,n}$ (Section 7.3) and other classes of simple drawings of $K_{m,n}$ (Section 7.4). In Section 7.5 we show a proof independent of the previous ones and valid for every simple drawing of $K_{m,n}$. Moreover, this proof shows the existence of a shooting star rooted at every vertex. Finally, in Section 7.6 we present a summary and conclusions.

7.2. Shooting stars in simple drawings of $K_{2,n}$

In this section we prove that every simple drawing of $K_{2,n}$ and $K_{3,n}$ contains plane spanning trees of a certain structure. In order to do so, we introduce some notions and provide some auxiliary results.

For a given simple drawing D of a graph $G = (V, E)$ and two fixed vertices $r \neq g \in V$, we define a relation $\rightarrow_{rg}$ on the remaining vertices $V \backslash \{r, g\}$, where $u \rightarrow_{rg} v$ if and only if the edge gu crosses the edge rv. In the following, we simply write $u \rightarrow v$ if the two vertices r and g are clear from the context.

Lemma 7.1. *Let D be a simple drawing of a graph $G = (V, E)$ and let $r \neq g \in V$ be two fixed vertices. The relation $\rightarrow_{rg}$ is asymmetric and acyclic, that is, there are no other vertices $v_1, v_2, \ldots, v_k \in V$ with $k \in \mathbb{N}$ such that $v_1 \rightarrow v_2 \rightarrow \ldots \rightarrow v_k \rightarrow v_1$.*

Proof. We give a proof by induction on k.

Induction basis. The case $k = 1$ is trivial. The case $k = 2$ follows from the fact that there is at most one crossing in every subdrawing induced by four vertices of a simple drawing. Thus, if gu crosses rv then, even when existing as edges of the drawing, gv cannot cross ru. This shows that the relation $\rightarrow$ is asymmetric. For the case $k = 3$, assume there are three vertices u, v, w with $u \rightarrow v \rightarrow w \rightarrow u$. Let $\triangle$ denote the closed region bounded by the edges ru, rv, and gu and not containing the vertex g; see Figure 7.2 for an illustration. We distinguish the following two cases:

Case 1: $w \notin \triangle$. Since $w \rightarrow u$ holds, the edge gw crosses ru, and therefore the boundary of $\triangle$. Since $g \notin \triangle$ and since gw cannot cross gu, gw must also cross rv. Thus, we have $w \rightarrow v$, which is a contradiction to $v \rightarrow w$ since the relation is asymmetric.

Case 2: $w \in \triangle$. Since $u \rightarrow v$, the edge gv cannot cross ru. Moreover, since gv can neither cross gu nor rv, it is therefore entirely outside $\triangle$. Since rw is entirely contained in $\triangle$, gv and rw cannot cross, and therefore, $v \not\rightarrow w$.

Since w can neither be inside nor outside $\triangle$, the statement is proven for the case $k = 3$.

Inductive step. Suppose for a contradiction that there exist $v_1, \ldots, v_k$ with $k \geq 4$ and $v_1 \rightarrow v_2 \rightarrow \ldots \rightarrow v_k \rightarrow v_1$ and that there is no smaller cycle. We write $u = v_1$, $v = v_2$, $w = v_{k-1}$, and $z = v_k$. Let $\triangle$ denote the closed region bounded by the edges ru,rv, and gu that does not contain the vertex r. We distinguish the following two cases:

Case 1: $z \notin \triangle$. We continue analogously to Case 1 of the base case $k = 3$. Since $z \rightarrow u$ holds, gz crosses ru, and therefore the boundary of $\triangle$. Since $g \notin \triangle$ and since gz cannot cross gu, gz must also cross rv. Thus, $z \rightarrow b$.

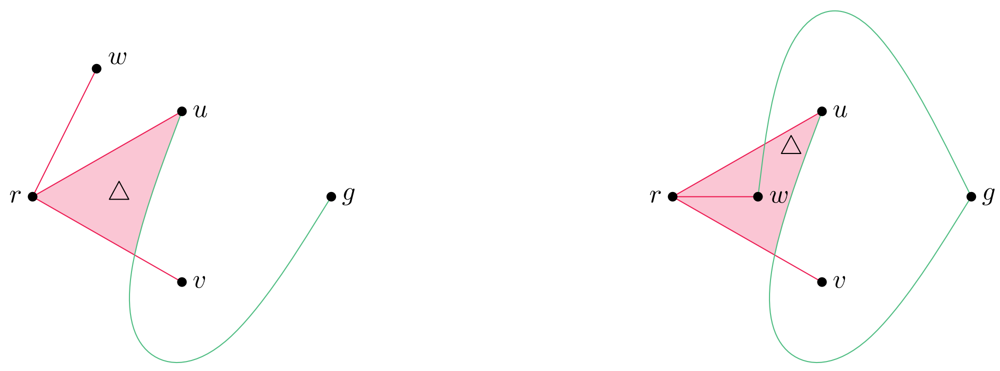

Figure 7.2.: An illustration of the base case $k = 3$ in Lemma 7.1.

Case 2: $z \in \triangle$**.** Since $w \to z$ holds, gw crosses rz at some point inside $\triangle$. Since $g \notin \triangle$ and since gw cannot cross gu, it must cross ru or rv (or both). Thus, we have $w \to u$ or $w \to v$.

In both cases, we can find $v'_1, \ldots, v'_l$ for some $l < k$ with $v'_1 \to \ldots \to v'_l \to v'_1$, which is a contradiction. This completes the proof of the lemma. □

Theorem 7.2. *Let D be a simple drawing of the complete bipartite graph $K_{2,n}$ with vertex partitions $\{r, g\}$ and B. Then, for every $k \in \{1, \ldots, n\}$, D contains a plane spanning tree with k edges incident to r and $n - k + 1$ edges incident to g.*

Proof. According to Lemma 7.1, we can find a labeling $b_1, \ldots, b_n$ of the vertices in B such that $b_i \to_{rg} b_j$ only holds if $i < j$. Let S_1 be the star rooted at r and with edges $rb_1, \ldots, rb_k$ and let S_2 be the star rooted at g and with edges $gv_k, \ldots, gv_n$. By Lemma 7.1 and the definition of relation $\to_{rg}$, the edges of S_1 and S_2 do not cross, and hence we have a plane spanning tree. □

Note that the same proof can be applied to the complete graph, and thus, we have the following result that might be of independent interest.

Corollary 7.3. *Let D be a simple drawing of the complete K_n. Then, for any two vertices r, g of D and for every $k \in \{1, \ldots, n-2\}$, D contains a plane spanning tree with k edges incident to r and $n - k - 1$ edges incident to g.*

Going back to simple drawings of $K_{2,n}$, the following corollary shows that they contain shooting stars rooted at both vertices of the vertex partition with two vertices.

Corollary 7.4. *Let D be a simple drawing of the complete bipartite graph $K_{2,n}$ with vertex partitions $\{r, g\}$ and B. Then for each $v \in \{r, g\}$, D contains a shooting star rooted at v.*

Proof. Consider again the proof of Theorem 7.2. With the according labeling $b_1, \ldots, b_n$ of the vertices in B, no edge gb_i can cross the edge rb_1. Hence, the plane spanning tree consisting of all the edges incident to g together with the edge rb_1 gives the desired shooting star rooted at g. Similarly, the tree with all edges incident to r and the edge gb_n is a shooting star rooted at r. □

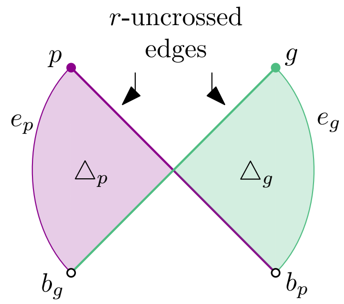

Figure 7.3.: The two crossing r-uncrossed edges together with the edges e_g and e_p bound the triangular regions $\triangle_g$ and $\triangle_p$.

7.3. Shooting stars in simple drawings of $K_{3,n}$

Making use of the result in the previous section, we can show that shooting stars also exist in simple drawings of $K_{3,n}$. The proof is based on a case analysis on the placement of the third vertex in the vertex partition with exactly three vertices.

Theorem 7.5. *Let D be a simple drawing of the complete bipartite graph $K_{3,n}$ with vertex partitions $\{r, g, p\}$ and B. Then for each $v \in \{r, g, p\}$, D contains a shooting star rooted at v.*

Proof. We color the vertices r, g, and p with colors red, green, and purple, respectively, and all edges of D in the color of the vertex in $\{r, g, p\}$ that they are incident to. Without loss of generality, let $v = r$. Starting with the star of r, we construct a shooting star ★ rooted at r.

From Theorem 7.2 it follows that there is a green and a purple edge that do not cross any of the red edges. In the rest of this proof, we call these two edges *r-uncrossed edges*. If they do not cross each other, we can insert them, together with g and p, into ★ and hereby complete our shooting star.

Hence, assume that the r-uncrossed edges cross. Let $b_g, b_p \in B$ be the vertices that are incident to the green and purple r-uncrossed edges, respectively. We denote by e_g the edge gb_p and by e_p the edge pb_g. These two edges e_g and e_p can neither cross the r-uncrossed edges nor each other (since in the subdrawing induced by four edges there can be at most one crossing). Thus, e_g, gb_g, and pb_p bound a (closed) triangular region $\triangle_g$ that does not contain p. Analogously, e_p, gb_g, and pb_p bound a (cloased) triangular region $\triangle_p$ that does not contain g; see Figure 7.3 for an illustration.

If either e_g or e_p does not cross any of the red edges, we obtain our shooting star ★ by inserting g, p, that edge, and the r-uncrossed edge of the other color. So assume now that both edges cross at least one red edge. We will show that then there has to be another green or purple edge which does neither cross a red edge nor the r-uncrossed edges.

Since red edges do not cross the r-uncrossed edges, at least one vertex must lie in the interior of $\triangle_g$ and $\triangle_p$. Furthermore, the red vertex r can either lie inside one of these triangular regions or outside them both. As the case in which r lies in $\triangle_p$ and the case in which it lies in $\triangle_g$ are symmetric, we only distinguish two cases regarding where r lies.

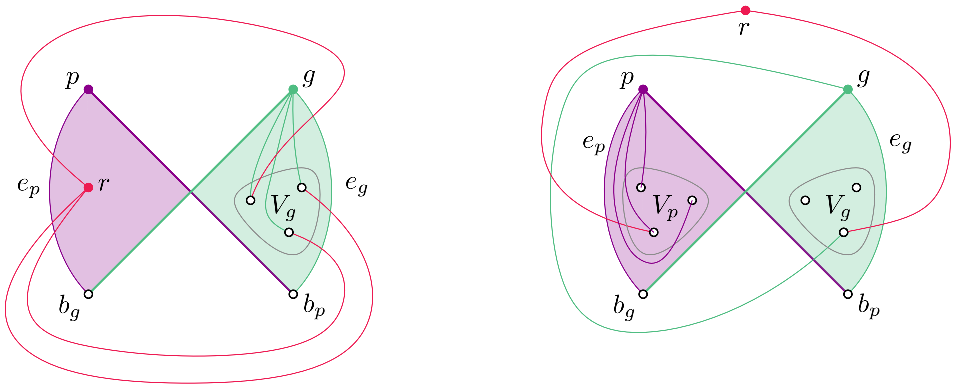

Figure 7.4.: An illustration of the proof of Theorem 7.5. Left: Case 1. Right: Case 2.

Case 1: r lies in $\triangle_p$. In this case, there has to be at least one vertex that lies in $\triangle_g$. Actually, all red edges that lie partly inside $\triangle_g$ have one endpoint in $\triangle_g$ (since they cannot cross the r-uncrossed edges). We denote this set of vertices inside $\triangle_g$ by V_g; see Figure 7.4 (left).

Claim 1. *A green edge incident to g and a vertex in V_g lies entirely in $\triangle_g$.*

Proof. If a green edge that has an endpoint in V_g, starting from this endpoint, leaves $\triangle_g$, it must cross the purple r-uncrossed edge to do so. Then, it cannot enter $\triangle_p$, since it can only cross one of its boundary edges. But, as it cannot cross the red edge incident to the same point in V_g, it can then not reach g. □

Claim 2. *There is a green edge incident to g and a vertex in V_g that neither crosses a red edge nor crosses the purple r-uncrossed edge.*

Proof. Consider the subdrawing D' of D that has V_g, the red vertex r, and the green vertex g as its vertex set and all green and red edges incident to the vertices in V_g as its edge set. Notice that D' is a simple drawing of $K_{2,|V_g|}$. From Theorem 7.2 it follows that there is a green edge in that subdrawing that does not cross any red edge in the subdrawing. By Claim 1, this green edge lies entirely in $\triangle_g$. Thus, it does not cross the purple r-uncrossed edge. Moreover, no green or red edge of D but not part of D' can lie (partly) inside $\triangle_g$. Thus, that green edge fulfills the conditions of the claim. □

Adding to the star of r the green edge guaranteed by Claim 2 and the purple r-uncrossed edge completes our shooting star ★ in Case 1.

Case 2: r lies outside both $\triangle_g$ and $\triangle_p$. As argued before, in the interior of both triangular regions $\triangle_g$ and $\triangle_p$ there must be at least one vertex. As in Case 1, all red edges that lie partly inside one of these triangular regions have one endpoint in the corresponding triangular region.

We denote by V_g and V_p the set of vertices inside $\triangle_g$ and $\triangle_p$, respectively. If all green edges incident to a vertex in V_g lie entirely in $\triangle_g$, we can obtain our shooting star ★ in the same way as in Case 1. Thus, we assume that this is not the case. We will show that all purple edges lie entirely inside $\triangle_p$ and then we can obtain our shooting star $\bigstar$ in the same way as in Case 1.

Note that all the green edges incident to a vertex in V_g and that do not lie entirely inside $\triangle_g$, (i) connect to g from outside $\triangle_g$, (ii) lie outside $\triangle_p$, and (ii) intersect all edges from r to a vertex in V_p; see Figure 7.4 (right). Thus, if a purple edge that has an endpoint in V_p, starting from this endpoint, leaves $\triangle_p$, it must cross the green r-uncrossed edge to do so. Then, since it cannot cross the red edge incident to the same point in V_p, in order to reach p it would need to cross twice all the green edges incident to a vertex in V_g and that do not lie entirely inside $\triangle_g$. $\square$

7.4. Shooting stars in certain classes of drawings of $K_{m,n}$

In this section we study the problem of finding shooting stars in special kinds of simple drawings of complete bipartite graphs.

Outer drawings [61] are simple drawings of $K_{m,n}$ in which all the m vertices of one vertex partition lie on the outer boundary of the drawing. That means that we can draw a simple closed curve through all the vertices in the vertex partition with m vertices that is not crossed by any of the edges in the drawing.

We now proceed to prove that there is a shooting star in every outer drawing of $K_{m,n}$.

Theorem 7.6. *Let D be an outer drawing of the complete bipartite graph $K_{m,n}$ with vertex partitions R and B where the vertices of R lie on the outer boundary. Let r be an arbitrary vertex in R. Then D contains a shooting star rooted at r.*

Proof. First, we label the vertices of R. We start in $r = r_1$ and traverse the outer boundary in clockwise direction. The vertices in R are denoted by $r_2, \ldots, r_m$ following the order in which they occur along the boundary. Let T_1 be the star of r_1. We will insert edges of Ds into T_1 until it becomes a spanning tree. We do so inductively by first inserting an edge incident to r_2, then an edge incident to r_3 and so on until we insert an edge incident to r_m. We denote by T_i the drawing that we get by inserting into T_{i-1} both r_i and the selected edge incident to r_i for $2 \leq i \leq m$. We will show that it is possible to insert vertices and edges such that every T_i is plane. After inserting the last vertex and edge, the statement follows.

In the first step, for T_2, we need to find an edge that is incident to r_2 and does not cross any edge incident to r. We know from Theorem 7.2 that there is at least one such edge. We insert r_2 and this edge into T_1 and get a plane tree T_2. For T_i we need to find an edge that is incident to r_i and does not cross any of the edges of T_{i-1}. We denote by e_{i-1} the edge of T_{i-1} that is incident to r_{i-1} and by b_{i-1} the vertex in B that e_{i-1} is incident to. We also denote by e'_{i-1} the edge that is incident to b_{i-1} and r. The edges e'_{i-1} and e_{i-1} separate two closed regions inside of the outer boundary of the drawing. We call Γ_1 the one including the part of the outer boundary of the drawing that goes from r following a clockwise direction until r_{i-1}, and the other one Γ_2; see Figure 7.5 for an illustration.

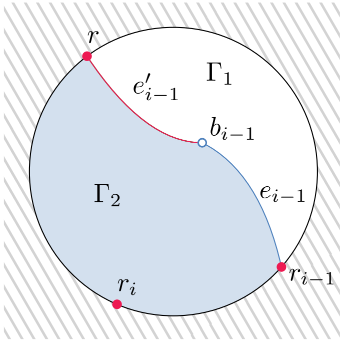

Figure 7.5.: An illustration of the proof of Theorem 7.6. The edges e'_{i-1} and e_{i-1} together with the outer boundary define two regions.

Claim 1. *All the edges of T_{i-1} that are not incident to r lie entirely inside Γ_1.*

Proof. Since the boundary of Γ_1 consists of edges in T_{i-1} and part of the outer boundary, all edges of T_{i-1} that lie partly inside Γ_2 have to lie entirely inside it. The edges of T_{i-1} that are not incident to r are incident with the vertices $r_2, \ldots, r_{i-1}$. As they have to lie on the part of the outer boundary that is also part of the boundary of Γ_1, the edges incident to these vertices have to lie partly inside Γ_1. Thus, these edges have to lie entirely inside Γ_1. □

We now consider the region Γ_2. The subdrawing of D induced by r, r_i, and all vertices of B that lie in Γ_2 is a simple drawing of $K_{2,n'}$. By Theorem 7.2, there is an edge e_i incident to r_i that does not cross any edge incident to r. This edge can neither cross the outer boundary nor e'_{i-1} and it can only cross e_{i-1} once. Since e_i has both end points in Γ_2, it has to lie entirely in Γ_2. From Claim 1 it follows that it does not cross any edge of T_{i-1} that is not incident to r. As it does not cross any edge incident to r either, e_i it does not cross any of the edges of T_{i-1}. Thus, we can insert r_i and e_i into T_{i-1} it and obtain a plane tree T_i. We continue to do so until we inserted an edge for every vertex in B. The plane spanning tree T_m is then a shooting star. □

We apply the strategy of the previous proof to a more general class of drawings. A *circular drawing* is a simple drawing of $K_{m,n}$ in which all the m vertices of one vertex partition lie on a closed curve that is not crossed by any edges, and all other vertices do not lie on this curve. Since we can consider separately the subdrawings in the two regions separated by the closed curve, we get the following corollary.

Corollary 7.7. *Let D be a circular drawing of the complete bipartite graph $K_{m,n}$ and let v be a vertex on the unintersected closed curve though all the vertives of one vertex partition. Then D contains a shooting star rooted at v.*

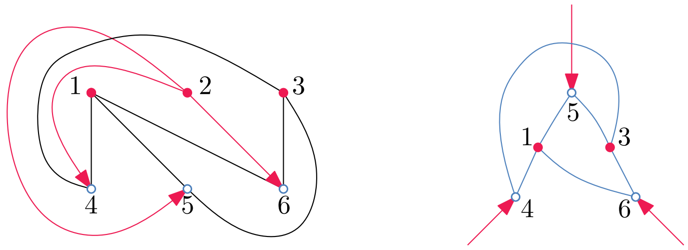

Figure 7.6.: Left: A simple drawing of $K_{3,3}$. Right: Drawing after a stereographic projection from vertex 2.

Figure 7.7.: Left: e_2 crosses e_1 in clockwise direction. Right: e_2 crosses e_1 in counterclockwise direction.

7.5. Shooting stars in simple drawings of $K_{m,n}$

We now present the general result that applies to every simple drawing of $K_{m,n}$ and that does not depend on any of the previous results in this chapter.

Theorem 7.8. *Let D be a simple drawing of the complete bipartite graph $K_{m,n}$ and let r be an arbitrary vertex of $K_{m,n}$. Then, D contains a shooting star rooted at r.*

Proof. We can assume that the vertex partitions are $R = \{r_1, \dots, r_m\}$, and $B = \{b_1, \dots, b_n\}$, and the vertices in them are colored red and blue, respectively. Without loss of generality, let $r = r_1$.

To simplify the figures in the proof, we consider the drawing D in the sphere and apply an stereographic projection from r onto a plane. In that way the edges in the star of r are represented as rays; see Figure 7.6. Moreover, we will depict them in red. In order to specify how two edges cross each other, we introduce some notation. Consider two crossing edges $e_1 = u_1v_1$ and $e_2 = u_2v_2$ oriented from u_1 to v_1 and from u_2 to v_2, respectively. Let $\times$ denote the crossing point and consider the arcs $\times u_1$ and $\times v_1$ on e_1 and the arcs $\times u_2$ and $\times v_2$ on e_2. We say that e_2 crosses e_1 in *clockwise direction* if the clockwise cyclic order of these arcs around the crossing point $\times$ is $\times u_1$, $\times u_2$, $\times v_1$, and $\times v_2$; see Figure 7.7 (left). Otherwise, we say that e_2 crosses e_1 in *counterclockwise direction*; see Figure 7.7 (right). In the following, we assume that all edges are oriented from their red to their blue endpoint.

We prove Theorem 7.8 by induction on n. For $n = 1$, the drawing D is a shooting star rooted at any vertex, and in particular at r.

Assume that the existence of shooting stars rooted at any vertex has been proved for any simple drawing of $K_{m,n'}$ with $n' < n$. Let M be a subset of the edges of D connecting each vertex $r_i \neq r$ to some blue vertex in B, such that (i) $M \cup \{\bigcup_{j=2}^{n} rb_j\}$ does not contain any crossing and (ii) the number of crossings of M with edge rb_1 is the minimum possible. Observe that the

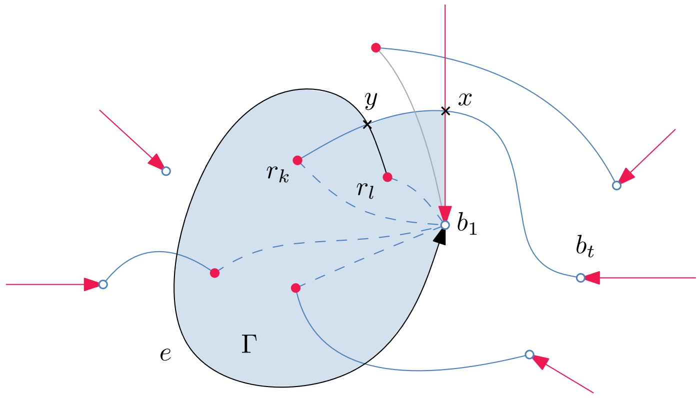

Figure 7.8.: Illustration of Case 1.

set M is well defined, since, by the inductive hypothesis, the drawing of D obtained by deleting the blue vertex b_1 and its incident edges contains a shooting star rooted at r. Thus, there exists a set of edges M_1 from D connecting each vertex $r_i \neq r$ to some blue vertex such that $M_1 \cup \{\bigcup_{j=2}^{n} rb_j\}$ does not contain any crossing. However, some of the edges in M_1 might cross rb_1 in D. We remark that this set M_1 might not be the one fulfilling condition (ii). Actually, we will show that M does not contain any crossing with rb_1.

Assume for a contradiction that rb_1 crosses at least one edge in M. When traversing rb_1 from b_1 to r, let x the first crossing point of rb_1 with an edge r_kb_t in M. Without loss of generality, when orienting rb_1 from r to b_1 and r_kb_t from r_k to b_t, r_kb_t crosses rb_1 in counterclockwise direction (otherwise we can mirror the drawing).

Suppose first that the arc r_kx (on r_kb_t and oriented from r_k to x) is crossed in counterclockwise direction by an edge incident to b_1 (and oriented from the red endpoint to b_1). Let $e = r_lb_1$ be such an edge whose crossing with r_kx at a point y is the closest to x. Otherwise, let e be the edge r_kb_1 and y be the point r_k. In the remaining figures, we represent in red the edges of the star of r, in blue the edges in M, and in black the edge e.

We distinguish two cases depending on whether e crosses an edge of the star of r. The idea in both cases is to define a region Γ and, inside it, redefine the connections between red and blue vertices to reach a contradiction.

Case 1: e does not cross any edge of the star of r. Let Γ be the closed region of the plane bounded by the arcs yb_1 (on e), b_1x (on rb_1), and xy (on r_kb_t); see Figure 7.8. Observe that for all the red vertices r_i inside region Γ the edge r_ib_1 must be in Γ. Let M_Γ denote this set of r_ib_1 edges with $r_i \in \Gamma$ and note that $r_kb_1 \in M_\Gamma$. Consider the set M' of red edges obtained from M by replacing, for each red vertex $r_i \in \Gamma$, the (unique) edge incident to r_i in M by the edge r_ib_1 in M_Γ, and keeping the other edges in M unchanged. In particular, the edge r_kb_t has been replaced by the edge r_kb_1. The edges in M_Γ neither cross each other nor cross any of the red edges rb_j. Moreover, we now show that the non-replaced edges in M must lie entirely outside Γ. These edges can neither cross r_kb_t (by definition of M) nor the arc b_1x (on rb_1). Thus, if they are incident to b_1 they cannot

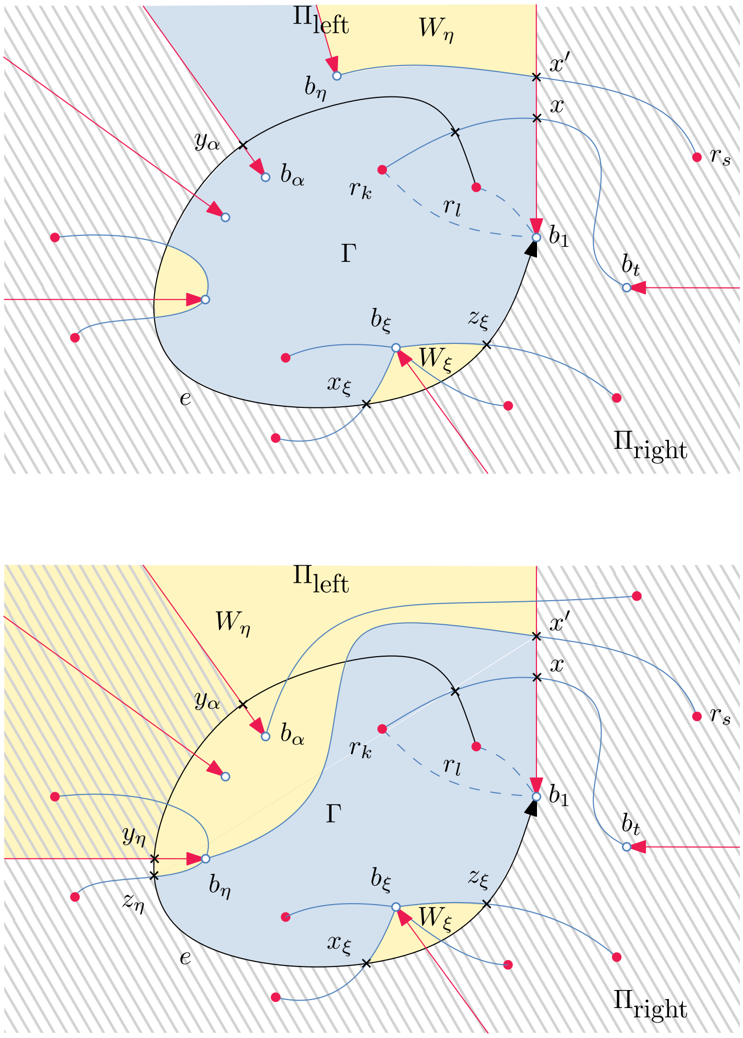

Figure 7.9.: Illustration of Case 2. Region Γ is shaded in blue while regions in $\bigcup_{\xi \in I} W_\xi \cup W_\eta$ are shaded in yellow. Top: $\eta \notin I$. Bottom: $\eta \in I$.

cross the boundary of Γ, and otherwise their endpoints lie outside Γ and they can only cross one arc of the boundary. Therefore, M' satisfies that $M' \cup \{\bigcup_{j=2}^{n} rb_j\}$ does not contain any crossing, and has less crossings with rb_1 than M (at least crossing x is not present and none of the edges in M_Γ crosses rb_1). This contradicts the definition of M as the one with the minimum amount of crossings with rb_1.

Case 2: e crosses the star of r. When traversing e from r_k or r_l (depending on the definition of e) to b_1, let $I = \{\alpha, \beta, \dots, \rho\}$ be the set of indices of the edges of the star of r crossed by edge e in precisely that order, and let $y_\alpha, \dots, y_\rho$ be the corresponding crossing points on e. Note that, when orienting e from r_k or r_l to b_1, the edges $rb_\xi, \xi \in I$, oriented from r to b_ξ, cross e in counterclockwise direction, since they can neither cross $r_k b_t$ (by definition of M) nor rb_1.

The three arcs ry_α (on rb_α), $y_\alpha b_1$ (on e), and $b_1 r$ divide the plane into two (closed) regions, Π_{left}, containing vertex r_k, and Π_{right}, containing vertex b_t. For each $\xi \in I$, let M_ξ be the set of red edges of M incident to some red vertex in Π_{right} and to b_ξ. Note that all the edges in M_ξ (if any) must cross the edge e. When traversing e from r_k or r_l to b_1, we denote by x_ξ, z_ξ the first and the last crossing points of e with the edges of $M_\xi \cup rb_\xi$, respectively. We remark that both x_ξ and z_ξ might coincide with y_ξ and, in particular, if $M_\xi = \emptyset$ then $x_\xi = y_\xi = z_\xi$.

We now define some regions in the drawing D. Suppose first that there is an edge in M (oriented from the red to the blue endpoint) that crosses rb_1 (oriented from r to b_1) in clockwise direction. Let $r_s b_\eta$ be such an edge in M whose crossing with rb_1 at a point x' is the closest to x (recall that the arc $b_1 x$ on rb_1 is not be crossed by edges in M). Then, if $\eta \notin I$, we denote by W_η the region bounded by the arcs rx' (on rb_1), $x'b_\eta$ (on $r_s b_\eta$), and rb_η and not containing b_1; see Figure 7.9 (top) for an illustration. If $\eta \in I$, we define W_η as the region bounded by the arcs rx' (on rb_1), $x'b_\eta$ (on $r_s b_\eta$), $b_\eta z_\eta$, $z_\eta y_\eta$ (on e), and $y_\eta r$ (on rb_η) and not containing b_1; see Figure 7.9 (bottom) for an illustration. Otherwise, if such crossing point x' does not exist, W_η denotes the edge rb_1. Moreover, for each $\xi \in I \setminus \{\eta\}$, we define W_ξ as the triangular region bounded by the arcs $x_\xi b_\xi$, $b_\xi z_\xi$, and $z_\xi x_\xi$ (and not containing b_1); see Figure 7.9 for an illustration.

We can finally define the region Γ for this case. It is the region obtained from Π_{left} by removing the interior of all the regions W_ξ, $\xi \in I$ plus wedge W_η if $\eta \notin I$ (otherwise it is already contained in $\bigcup_{\xi \in I} W_\xi$). Now, consider the set of red and blue vertices contained in region Γ. Let J denote the set of indices such that for all $j \in J$ the blue vertex b_j lies in Γ (note that $1 \in J$). Since b_t is not in Γ, by the inductive hypothesis, we can find a set of edges M_Γ connecting each red vertex in Γ with a blue vertex b_j, $j \in J$ satisfying that $M_\Gamma \cup \{\bigcup_{j \in J} rb_j\}$ does not contain any crossing. We now show that all the edges in M_Γ lie entirely in Γ. An edge in M_Γ cannot cross any of the edges rb_j, with $j \in J$. Thus, it cannot leave Π_{left}, as otherwise it would cross e twice. Moreover, if it enters one of the regions in $\bigcup_{\xi \in I} W_\xi \cup W_\eta$, it would have to leave it crossing e, and then it cannot reenter Γ.

Consider the set M' of red edges obtained from M by replacing, for each red vertex $r_i \in \Gamma$, the edge $r_i b_\xi$ in M by the edge $r_i b_j$, $j \in J$, in M_Γ, and keeping the other edges in M unchanged. In particular, the edge $r_k b_t$ has been replaced by some edge $r_k b_j$, $j \in J$. The edges in M_Γ neither cross each other nor cross any of the red edges rb_j, $j \in J$ nor any of the other ones, lying entirely outside Γ. Moreover, the non-replaced edges in M cannot enter Γ since the only part of its boundary that they can cross are arcs on e. Therefore, M' satisfies that $M' \cup \{\bigcup_{j=2}^{n} rb_j\}$ does not contain any crossing, and has less crossings with rb_1 than M (at least crossing x is not present and none of the edges in M_Γ crosses rb_1). This contradicts the definition of M as the one with the minimum amount of crossings with rb_1. $\square$

7.6. Chapter summary

We have shown that every simple drawing of the complete bipartite graph $K_{m,n}$ contains a plane spanning tree. Though this is a basic aspect of simple drawings of $K_{m,n}$, to our best knowledge this is the first proof. We show that, in particular, simple drawings of $K_{m,n}$ contain shooting stars rooted at every vertex. However, our proof does not provide a polynomial algorithm to find them. We have also presented alternative proofs showing the existence of shooting stars for particular cases of simple drawings of $K_{m,n}$, namely all simple drawings of $K_{2,n}$ and $K_{3,n}$, as well as all circular drawings of $K_{m,n}$. We remark that these proofs for particular cases prove the existence of shooting stars rooted at any vertex in one specific vertex partition out of the two, and provide polynomial algorithms to find them.

8. Semi-simple drawings of K_n

An extended abstract containing the results in this chapter has appeared in [19]. Moreover, the results in Sections 8.3, 8.4, 8.6, and part of the ones in Section 8.7 appeared in the Master's thesis [82], co-supervised by the author of this thesis.

8.1. Introduction

Two simple drawings D and D' are *weakly isomorphic* if there exists an incidence-preserving bijection between their vertices, such that two edges of D cross if and only if the corresponding two edges of D' cross. We will consider the classes of weakly isomorphic simple drawings of the complete graph K_n. We can efficiently handle them using rotation systems. The *rotation* of a vertex in a drawing is the (clockwise) cyclic order of all edges incident to it. The rotations of all the vertices of a drawing form its *rotation system*. A rotation system is said to be *realizable* if it is a rotation system of a simple drawing. Two simple drawings of K_n are weakly isomorphic if and only if they have the same rotation system (up to reflection) [135].

A *semi-simple drawing* of a graph is a drawing in which edges that share a vertex do not cross, but edges not sharing a vertex are allowed to (properly) cross an arbitrary number of times; see [163] and also [42]. Semi-simple drawings can have regions that are bounded only by two continuous pieces of (two) edges, which we denote *lenses*. If a lens contains no vertex, we call it *empty*. We call a rotation system *semi-realizable* if it can be drawn as a semi-simple drawing. It is known that every semi-realizable rotation system can be drawn without empty lenses [157]. We call semi-simple drawings without empty lenses *minimal*.

Motivated by the amount of structure of simple drawings determined by rotation systems, we investigate the properties of (minimal) semi-simple drawings with respect to their rotation systems. As for simple drawings, a semi-simple drawing with three vertices can only be a simple 3-cycle. A simple drawing of K_4 has either no crossing or one crossing; a semi-simple drawing with the latter rotation system has exactly one pair of edges that crosses an odd number of times. We observe that for K_4, all minimal semi-simple drawings are simple. For K_n, the rotation system of a semi-simple drawing determines whether two edges cross an even or an odd number of times. This is related to the Hanani-Tutte theorem, which states that a graph is planar if and only if in a drawing any two non-adjacent edges cross an even number of times, and which motivates the notion of the *independent odd crossing number* [179]. Pach and Tóth [163] showed that semi-simple drawings can be redrawn such that no two edges that cross every edge an even number of times intersect, a result improved in [168], who showed that, additionally, the drawings have the same rotation system. Note that this is not equivalent to obtaining a drawing with no two edges crossing an even number of times. While the Hanani-Tutte theorem (like probably most results considering the odd crossing number) considers all drawings of a given graph, we are interested in drawings of K_n with a given rotation system.

Given a rotation system, we call the rotation systems restricted to 4 or 5 elements *4-tuples* and *5-tuples*, respectively. Kynčl showed that a rotation system of K_n is realizable if and only if all its 4- and 5-tuples are realizable [138]. Recently, Cardinal and Felsner [61] proved a similar result for outer drawings of $K_{m,n}$.

Outline. In Section 8.2 we show that checking the realizability of all 4-tuples in a rotation system of K_n can be done in $O(n^3)$. As a 4-tuple is semi-realizable if and only if it is realizable, it was conjectured that realizability of all 4-tuples is sufficient for semi-realizability. We refute this conjecture in Section 8.3. In Section 8.4 we present our backtracking approach for deciding whether a given rotation system with few elements is semi-realizable. It is based on an algorithm for generating simple drawings [1]. As the computations require an upper bound on the number of crossings between two edges in a minimal semi-simple drawing, in Section 8.5 we tackle this problem. The results from the computations are presented in Section 8.6. In Section 8.7 we explore semi-simple drawings in relation with Conway's thrackle conjecture and the existance of plane Hamiltonian cycles. Finally, in Section 8.8 we present a summary of the chapter and conclusions.

8.2. Realizability of rotation systems

To decide the realizability of a rotation system, it suffices to check that all 5-tuples (and thus 4-tuples) are realizable [138]. While a brute-force approach to check all the 4-tuples yields an $O(n^4)$ time algorithm, we show that only $O(n^3)$ checks are needed.

Observation 8.1. *In a realizable 4-tuple, the rotations of three of its vertices determine the rotation system.*

Lemma 8.2. *In any rotation system of five vertices, the number of non-realizable 4-tuples is even.*

Proof. A *flip* in a rotation system is the exchange of the positions of two neighboring vertices in the rotation of a vertex. Let $\mathcal{R}$ be a rotation system of the vertices $\{a, b, c, d, e\}$ and consider an arbitrary flip of two neighboring vertices b and c in the rotation of a. This flip only affects the two 4-tuples on $\{a, b, c, d\}$ and on $\{a, b, c, e\}$. Since in each of these two 4-tuples only one rotation changes, the realizability of the two 4-tuples switches due to Observation 8.1. Hence, by one single flip, the number of non-realizable 4-tuples in $\mathcal{R}$ either changes by two or stays the same. Using the argument of various sorting algorithms, we know that by multiple single flips we can obtain every possible rotation system. Since we know that a 5-tuple without a non-realizable 4-tuple exists the number of non-realizable 4-tuples in a 5-tuple is always even. $\square$

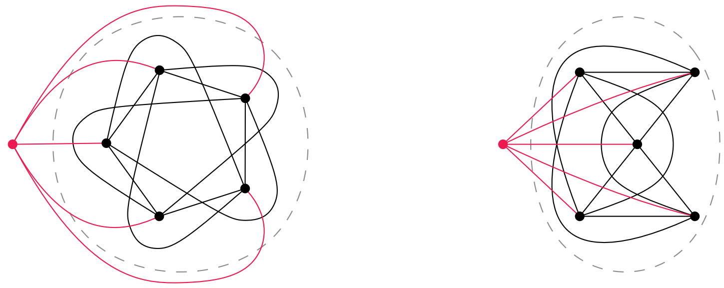

Figure 8.1.: Semi-simple drawings extending a non-realizable 5-tuple (depicted in black).

Lemma 8.3. *Let $\mathcal{R}$ be a rotation system of n vertices. If $\mathcal{R}$ contains a non-realizable 4-tuple, then every vertex of $\mathcal{R}$ is contained in a non-realizable 4-tuple.*

Proof. Assume that there are four vertices v, x, y, z whose sub-rotation system is non-realizable. Then a 5-tuple $\{u, v, x, y, z\}$ with any fifth vertex u is also non-realizable. From Lemma 8.2 we know that there exists a second non-realizable 4-tuple in this 5-tuple. This 4-tuple must include u. □

Thus, to decide realizability of all 4-tuples it suffices to check all 4-tuples containing an arbitrary vertex u.

Corollary 8.4. *Checking the realizability of all 4-tuples in a rotation system can be done in cubic time.*

There are five realizable 5-tuples and two other semi-realizable ones; see the black sub-drawings in Figure 8.1. While for 4-tuples it was sufficient to check all those containing one fixed element, this approach no longer works for realizability of 5-tuples, as there are arbitrarily large semi-simple drawings containing only one non-realizable 5-tuple. Such drawings can for example be constructed modifying the drawings in Figure 8.1 by adding an arbitrary number of copies of the red vertex close to it.

8.3. Semi-realizability of rotation systems

From the computations we know that all rotation systems of five vertices that have only realizable 4-tuples are either realizable or semi-realizable; see Section 8.6. Moreover, each of them has a unique minimal semi-simple drawing (up to homeomorphism of the sphere). However, in general it is not the case that realizability of 4-tuples implies semi-realizability, which disproves the related conjecture mentioned in the introduction.

Theorem 8.5. *The semi-realizability of a rotation system does not follow from the realizability of all its 4-tuples.*

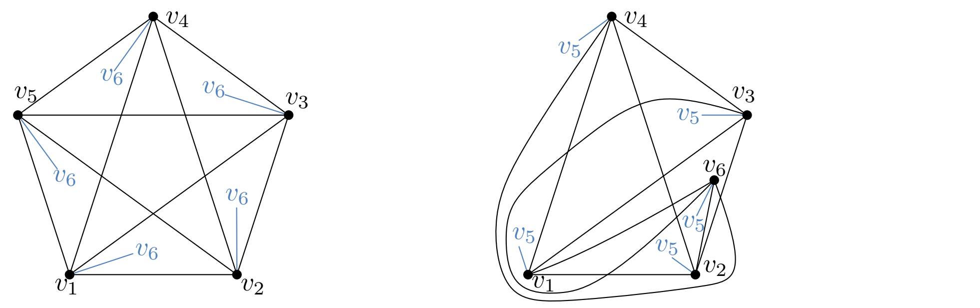

Figure 8.2.: Two partial drawings of the unique non-semi-realizable rotation system on six vertices.

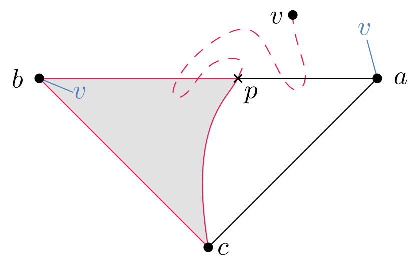

Figure 8.3.: The rotation system determines in which of the two regions bounded by a triangle lies a vertex.

Figure 8.2 shows two partial drawings of a rotation system with six vertices that is not semi-realizable, but whose 5-tuples are all semi-realizable. The left one is equivalent to the rotation system of the geometric drawing of K_5 with the five vertices in convex position; all other 5-tuples are equivalent to the rotation system of the semi-simple drawing to the right.

For proving that the given example in Figure 8.2 is not semi-realizable, we argue about area-containment of points in triangles.

We denote the directed cycle defined by three vertices a, b, c the *triangle* $\triangle(a, b, c)$. Every triangle $\triangle(a, b, c)$ decomposes the plane into two regions, one to the left of the directed edge ab and one to the right. We call the former the *interior* of $\triangle(a, b, c)$ and the latter the *exterior* of $\triangle(a, b, c)$. The following lemma generalizes a well-known property of simple drawings to semi-simple drawings.

Lemma 8.6. *For any rotation system of a semi-simple drawing of K_n and a triangle $\triangle(a, b, c)$ of this drawing, consider the two regions bounded by $\triangle(a, b, c)$. Any other vertex v of the drawing lies in the region in which at least two of the three edges to v emanate from a, b, and c.*

Proof. Figure 8.3 illustrates the proof. We assume without loss of generality that the edges from b and c to v start in the interior of $\triangle(a, b, c)$. We further assume for the sake of contradiction that vertex v lies in the exterior of $\triangle(a, b, c)$.

Since the edge cv has to leave the triangle $\triangle(a, b, c)$ and must not cross any edge incident to c, it has to cross the edge ab. We call the first such crossing point p. Then the curves cp,

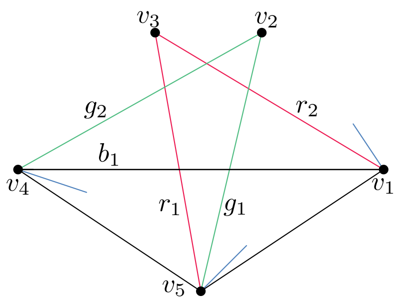

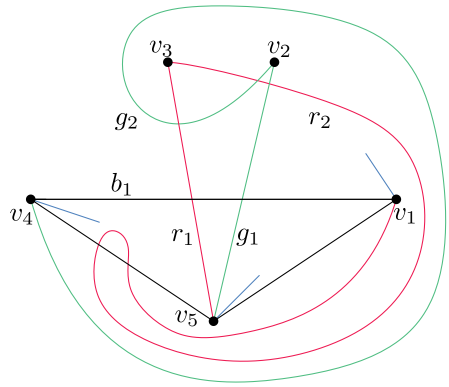

Figure 8.4.: Triangles $\triangle(v_1, v_4, v_5)$, $\triangle(v_1, v_3, v_5)$, and $\triangle(v_2, v_4, v_5)$ of the non-semi-realizable rotation system $\mathcal{R}$.

pb, and bc bound a region in the interior of $\triangle(a, b, c)$. The edge bv starts inside this region and cannot leave it. Since v lies in the exterior of $\triangle(a, b, c)$, we cannot draw the edge bv, a contradiction. □

Using Lemma 8.6 we can proof Theorem 8.5.

Proof of Theorem 8.5. We show that the rotation system $\mathcal{R}$ depicted in Figure 8.2 is not semi-realizable. We remark that the statement of Theorem 8.5 also follows from the computations; see Section 8.6.

Consider the triangles $\triangle(v_1, v_4, v_5)$, $\triangle(v_1, v_3, v_5)$, and $\triangle(v_2, v_4, v_5)$. From Lemma 8.6 it follows that v_6 has to lie in the interior of the black triangle $\triangle_b = \triangle(v_1, v_4, v_5)$, and to the exterior of the red triangle $\triangle_r = \triangle(v_1, v_3, v_5)$ and the green triangle $\triangle_g = \triangle(v_2, v_4, v_5)$.

If we can prove that in any semi-simple drawing of the 5-tuple $(v_1, v_2, v_3, v_4, v_5)$, the black triangle $\triangle(v_1, v_4, v_5)$ is always covered by the red $\triangle(v_1, v_3, v_5)$ and the green $\triangle(v_2, v_4, v_5)$, then it follows that there exists no semisimple drawing of $\mathcal{R}$, since we cannot draw vertex v_6. We sketch the configuration of the three triangles in Figure 8.4.

We label the edges like in Figure 8.4: We denote by r_1 the edge from v_5 to v_3, r_2 the edge from v_3 to v_1, and analogously g_1 the edge from v_5 to v_2, g_2 the edge from v_2 to v_4 and b_1 the edge from v_1 to v_4. Note that with the definition of the interior and exterior of a triangle it does not make a difference which region is the unbounded one in the drawing; see Figure 8.4 for an example. There, by our definition, the interior of the green triangle $\triangle(v_2, v_4, v_5)$ is the unbounded region.

The crossings shown in Figure 8.4 follow from the rotation system $\mathcal{R}$ (that is, edges crossing in this drawing cross an odd number of times in any semi-simple drawing of $\mathcal{R}$). The edge r_2 can intersect v_4v_5 an even number of times, forming inside $\triangle_b$ lenses bounded by r_2 and v_4v_5. We denote by $\triangle'_r$ the region inside $\triangle_r$ and outside these lenses. Analogously, we define $\triangle'_g$ as the region contained in $\triangle_g$ outside the lenses in $\triangle_b$ bounded by g_2 and v_5v_1. Notice that by definition $\triangle_r \subseteq \triangle'_r$ and $\triangle_g \subseteq \triangle'_g$. We further define the following regions inside the

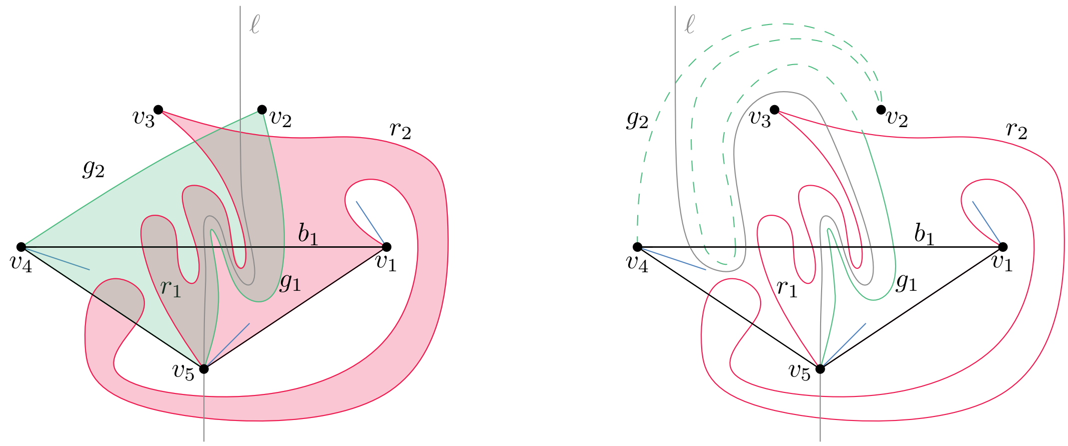

Figure 8.5.: Left: Illustration of the proof of Theorem 8.5. Right: Not a possible subdrawing of $\mathcal{R}$ since in the drawing v_3 lies inside the green triangle.

black tirangle $R_r := \triangle'_r \cap \triangle_b$ and $L_g := \triangle'_g \cap \triangle_b$. From the above definitions it follows that $R_r \subseteq \triangle'_r \subseteq \triangle_r$ and $L_g \subseteq \triangle'_g \subseteq \triangle_g$; see Figure 8.5 for an illustration.

Given the drawings of r_1 and g_1, since the two edges do not intersect, we can draw a pseudoline ℓ through v_5 separating the two edges and dividing the plane into the left halfplane L (where r_1 lies) and the right halfplane R (where g_1 lies). ℓ divides the black triangle into two not necessarily connected subsets: $L_\triangle := L \cap \triangle_b$ and $R_\triangle := R \cap \triangle_b$. Notice that from the rotation system we know that v_2 and v_3 are outside the black triangle and, respectively, the red and the green ones. In particular this means that we can draw ℓ such that in the drawing of the black triangle, r_1, g_1, and ℓ the vertices v_1, v_2, v_3, v_4 and v_5 lie in the unbounded face of the subdrawing. This implies that $L_\triangle \subseteq L_g$ and $R_\triangle \subseteq R_r$.

Summing up, we have that

$$\triangle = L_\triangle \cup R_\triangle \subseteq L_g \cup R_r \subseteq \triangle'_g \cup \triangle'_r \subseteq \triangle_g \cup \triangle_r$$

So, as we wanted to show, the black triangle is contained in the union of the green and the red ones. □

8.4. Computer-assisted results

To decide algorithmically whether a rotation system is semi-realizable, we used a backtracking approach based on the algorithm for realizing simple drawings used in [1]. We modified it allowing multiple proper crossings per edge pair. It thus requires an upper bound on the maximum number of proper crossings per edge pair in a minimal semi-simple drawing of K_n. From computations we get that five and ten crossings are such upper bounds for $n = 6$ and $n = 7$, respectively; see Table 8.1.

Using these parameters, we verified that the example in Figure 8.2 is the only non-semi-realizable one with six vertices where all 4-tuples are realizable. For seven vertices we exhaustively analyzed

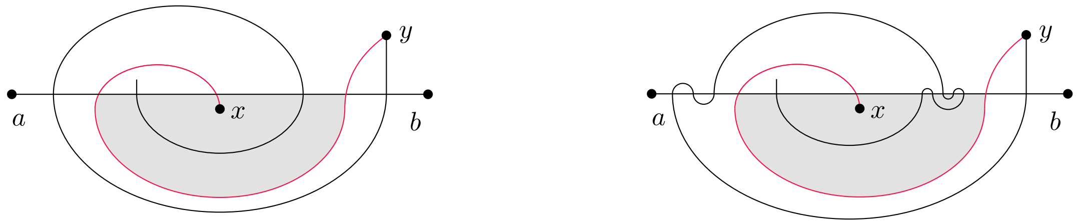

Figure 8.6.: Two edges forming a basic spiral (left) and a general spiral (right).

n	5	6	7	8	9	10	11	12	13	14	15
cr	2	5	10	27	35	59	83	143	197	323	589
K_n	2	4	10								

Table 8.1.: Maximum number of crossings per edge pair for n points (cr), and maximum number of pairwise crossings needed to draw K_n semi-simple.

all possible rotation systems. We found 480 non-semi-realizable rotation systems such that the sub-rotation system of every proper subset of vertices is semi-realizable.

To determine an upper bound on the maximum number of crossings per edge pair, we use another backtracking algorithm. It enumerates all different ways how two edges can cross multiple times without creating some forbidden patterns. We proceed with an overview of how the algorithm works; a detailed description can be found in [82].

The drawings with two edges in which we are interested are sub-drawings of minimal semi-simple drawings of K_n, that is, we want them to be completable to a minimal semi-simple drawing of K_n. In particular, the configuration shown in Figure 8.6, which we call a *spiral*, is forbidden, as it is not completable: Vertex x lies inside a lens bounded by the edges ab and xy (gray region in Figure 8.6), and thus, edges xa and xb cannot be added to the drawing.

The algorithm starts with a single edge. We also fix the starting vertex of the second edge. In every step the drawing is extended by adding a crossing (and maybe a vertex) in all possible ways avoiding spirals and empty lenses. It ends when we are forced to create an empty lens for which we do not have a point remaining to place in.

With this algorithm we can compute an upper bound on the maximal number of crossings of an edge pair in a semi-simple drawing of K_n. In Table 8.1 we show the results for up to 15 vertices.

8.5. Bounds on the number of crossings of edge pairs

In Section 8.4 we presented upper bounds on the number of times that two edges can cross in a minimal semi-simple drawing obtained by computer. The numbers in Table 8.1 seem to grow exponentially fast on the number of vertices. The next result shows that these numbers will grow at least exponentially fast.

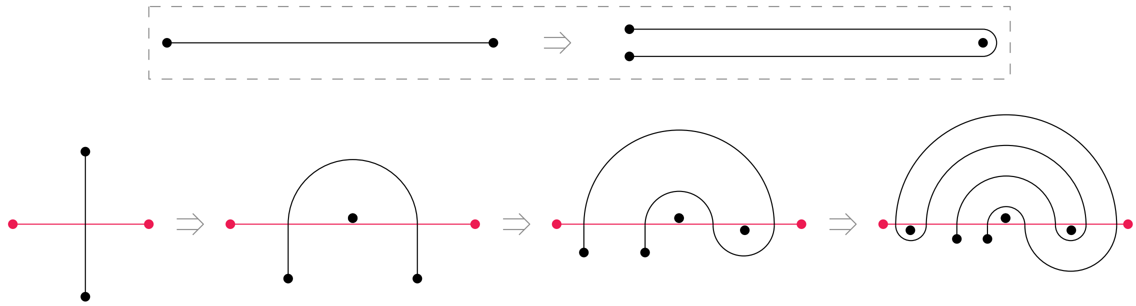

Figure 8.7.: Construction of one edge pair crossing 2^{n-4} times.

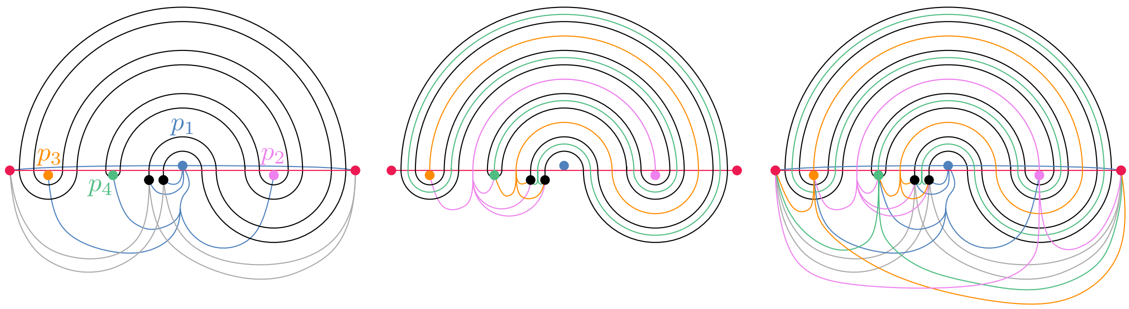

Figure 8.8.: Extension of the construction to a drawing of K_n.

Theorem 8.7. *There exist minimal semi-simple drawings of K_n in which two edges cross 2^{n-4} times.*

Proof. We construct a drawing of a pair of edges crossing an exponential number of times in the amount empty lenses created. To have a minimal drawing we place a point per empty lens. The idea of the iterative construction is, in every step, to replace one edge by a folded edge twice as long, and to add a new point where the new edge is folded; see Figure 8.8. We start with two edges crossing once. In every step we thus duplicate the number of crossings and add one more point, so after $n-4$ steps we have drawn n points and the two edges cross 2^{n-4} times.

Now we show how to add the remaining edges to complete the drawing to a drawing of K_n. In our construction we refer to the horizontal edge as the red edge, with endpoints r_1 and r_2, and to the other edge, that we replace in every iteration, as the black edge, with endpoints b_1 and b_2. Any other point is denoted by p_i with i being the iteration in which it was added. Connections between the black and the red endpoints are done in a direct manner; see gray edges in Figure 8.8 (left). Connecting a point p_i with the red endpoints is also done in a direct manner; see Figure 8.8 (right). The edge from p_i to a black endpoint goes first in parallel to the black edge (until reaching the unbounded face in the subdrawing of the red and the black edges) and then it is directly connected; see Figure 8.8 (right). Finally, the edge from p_i to p_j with $i < j$ also goes first in parallel to the black edge, until reaching the unbounded face in the subdrawing of the red and the black edges, and then it is directly connected; see Figure 8.8 (middle). $\square$

Number of vertices	5	6	7	8
Rotation systems	7	173	39 349	42 336 167
Realizable	5	102	11 556	5 370 725
Semi-realizable, needs 2 crossings	2	62	20 634	21 657 419
Semi-realizable, needs 3 crossings	-	5	3 379	$\geq$ 5 500 000
Semi-realizable, needs 4 crossings	-	3	2 152	$\geq$ 1 000 000
Semi-realizable, needs ≥ 5 crossings	-	-	808	?
Contains a non-semi-realizable subset	-	-	340	2 634 993
Minimal non-semi-realizable	0	1	480	?

Table 8.2.: Number of rotation systems (with all 4-tuples realizable) and their types of realizations.

We remark that the construction of the two edges crossing an exponential number of times has been presented before in the context of picture hanging [139, arXiv version, Section 5]. Concerning the completion to a semi-simple drawing, note that the intuitive straight-forward approach of connecting all pairs of blocking points with simple arcs "from below" (similar as the edges p_1p_i in Figure 8.8) does not work, as this would cause crossings between edges incident to the same vertex.

8.6. Results from computations

Table 8.2 summarizes the results obtained from the computations described in Section 8.4. We started with all rotation systems entirely consisting of realizable 4-tuples, and subtracted the realizable ones; see [1].

For the remaining sets we checked whether they admit a semi-simple drawing with a predefined maximal number of crossings per edge pair. Increasing this maximal number stepwise from 2 to the maximum given by Table 8.1, we obtained semi-simple drawings with the minimum maximal number of crossings per edge pair. The number of rotation systems requiring 2, 3, 4 or more crossings is also given in Table 8.2. For $n = 6$ there is precisely one rotation system which cannot be drawn semi-simple. For $n = 7$ there are 340 rotation systems which contain this rotation system as a subset, and are thus not semi-realizable. For $n = 8$ there are 2 634 993 rotation systems which contain a non-semi-realizable set of cardinality 6 or 7. The computations for $n = 8$ are still ongoing. Thus, we currently do not know if there exists a non-semi-realizable rotation system containing only semi-realizable proper subsets.

Let us finally remark that for $n = 6$, two edges can cross up to 5 times and there exists a minimal semi-simple drawing obtaining this number; see Figure 8.9 (left). However, for this rotation system there is a semi-simple drawing with only 3 crossings per edge pair; see Figure 8.9 (right). This shows that the bounds obtained in Table 8.1 are in general not tight.

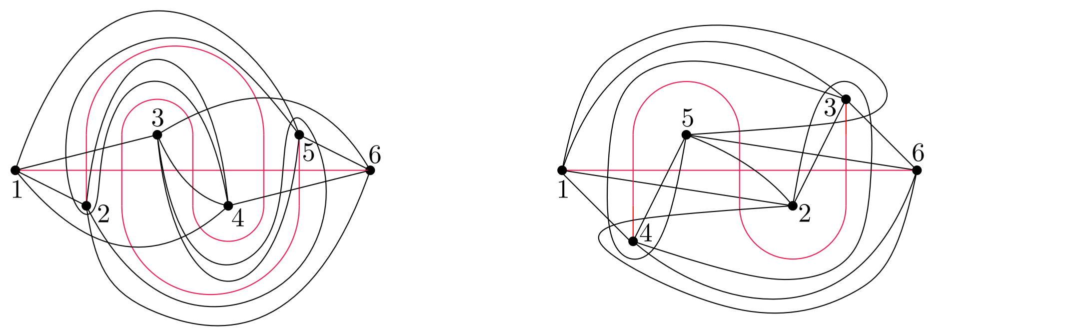

Figure 8.9.: Two semi-simple drawings of the same rotation system with up to five (left) and only three (right) crossings per edge pair.

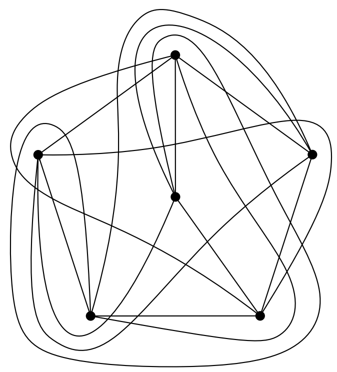

Figure 8.10.: A semi-simple drawing of K_6 without a plane Hamiltonian cycle.

We also investigated generalizations of the conjecture on plane Hamiltonian cycles and thrackles in simple drawings for the class of semi-simple drawings; see the following two sections.

8.7. Plane Hamiltonian cycles and thrackles

A *plane Hamiltonian cycle* of a drawing of G is a Hamiltonian cycle that does not cross itself. It is conjectured that every simple drawing of K_n contains a plane Hamiltonian cycle [171], but so far no proof was found. We show that a similar conjecture is not true for semi-simple drawings of K_n by providing an example with 6 vertices that does not contain a plane Hamiltonian cycle (where two edges of the cycle are allowed to cross an even number of times, but not an odd number of times); see Figure 8.10.

A *thrackle* is a simple drawing of a graph where each pair of edges meets exactly once, either at a common vertex or at a proper crossing [143]. A *generalized thrackle* is a drawing of a graph in which each pair of edges meets an odd number of times. For semi-simple drawings we can consider a notion of thrackle in between the above two. A *quasi-thrackle* [96] is a semi-simple drawing of a graph where each pair of edges either is incident to a common vertex or crosses an odd number of times.

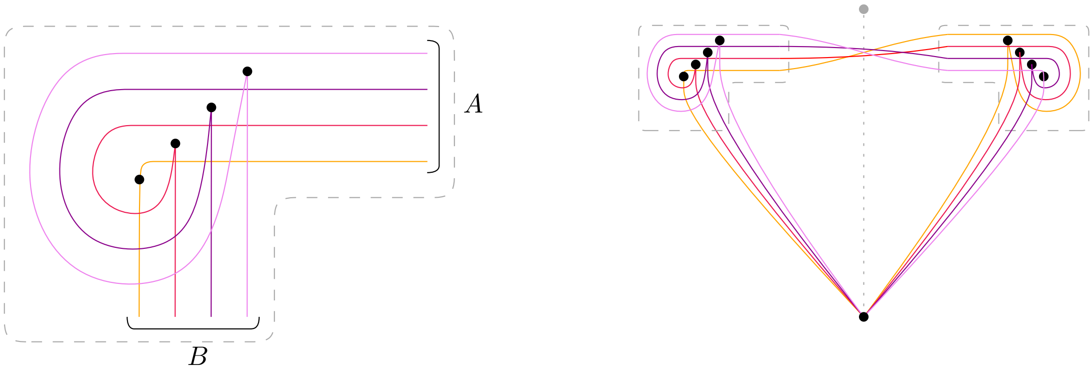

Figure 8.11.: Family of semi-simple thrackles matching the upper bound.

Conway conjectured that, for simple drawings, the number of edges of a thrackle cannot exceed the number of its vertices. However, despite of the efforts to try to prove the conjecture, the best known upper bound is $m \leq \frac{3}{2}(n-1)$ [60]. For generalized thrackles the tight upper bound is $m \leq 2n-2$ [60]. However, none of the examples matching the upper bound [60] is a semi-simple thrackle. Moreover, the proof of the upper bound of $m \leq \frac{3}{2}(n-1)$ [60] for thrackles translates direclty to quasi-thrackles. A construction matching this upper bound, and thus, showing that it is tight, is represented in Figure 8.11. This construction, that we obtained independently, is very similar to the one in [96], that was used for showing the stated tight bound of $m \leq \frac{3}{2}(n-1)$. Notice that, if Conway's conjecture holds, the upper bound on the number of edges in quasi-thrackles would be exactly half way (up to a constant) between the one for thrackles and the one for generalized thrackles.

We can also consider thrackles in drawings completable to a semi-simple drawing of K_n. For them, a direct translation of Conway's conjecture would imply that no semi-simple drawing of K_n can contain a quasi-thrackle with $n+1$ edges. In Figure 8.12 we show a semi-simple drawing of K_7 that contains a subgraph with eight edges that pairwise either share a common vertex or cross an odd number of times. This shows that a generalization of Conway's thrackle conjecture does also not directly translate for completable semi-simple thrackles.

8.8. Chapter summary

In this chapter we studied semi-simple drawings of K_n. For deciding if a rotation system is realizable as a simple drawing it is known that only (4-tuples and) 5-tuples need to be checked [138]. For the problem of characterizing semi-realizability, we showed that there is a rotation system with 6 vertices that is not semi-realizable, but where every 5-tuple is realizable. By an exhaustive computer search we found 480 rotation systems with 7 vertices such that the rotation system of any proper subset of vertices is semi-realizable. This indicates that checking semi-realizability is harder than checking realizability. Moreover, it seems plausible that there are arbitrarily large minimal non-semi-realizable rotation systems. The computational complexity of deciding semi-realizability remains an open problem. Concerning the maximum number of times that two edges can cross in a minimal semi-simple drawing of K_n, we computed

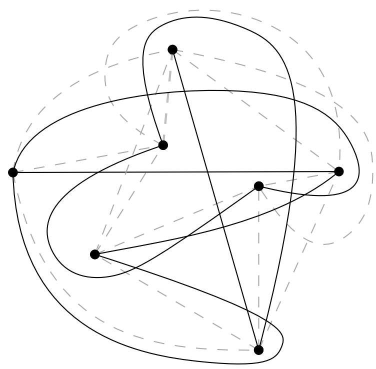

Figure 8.12.: A quasi-thrackle with eight edges in a semi-simple drawing of K_7.

upper bounds for up to 15 vertices. A general (finite) upper bound is still unknown. Finally, we showed that in semi-simple drawings of K_n there is not always a plane Hamiltonian cycle. And we showed that in semi-simple drawings of K_n there might be thrackles with more edges than vertices.

9. Open problems

Several open problems have emerged from the work presented in this thesis. We propose here a selection of them. They are sorted by the order of the chapters they are related to.

Problem 1. *Given a straight-line drawing of K_n, what is the computational complexity of finding a 2-edge-coloring that minimizes the number of monochromatic crossings?*

This problem is related to the max-cut problem of segment intersection graphs, which is NP-complete for general graphs [28]. However, for the intersection graph of a straight-line drawing of K_n the complexity is unknown.

Problem 2. *Is the straight-line convex drawing D_n of K_n, among all straight-line drawings of K_n, the one giving the largest ratio between the minimum number of monochromatic crossings given by a 2-edge-coloring and the number of crossings?*

In Section 2.4.2 we proved that for these drawings the ratio approaches $3/8$ from below when $n \to \infty$.

Problem 3. *Does the combinatorial description of an exit drawing encode the abstract order type?*

In Section 3.6 we showed that information given by the set of exit edges together with their witnesses is not enough to determine the order type.

Problem 4. *What is the computational complexity of deciding whether a given simple drawing is saturated, that is, no further edge can be inserted?*

In Section 5.5 we presented a reduction showing that deciding if one particular edge can be inserted into a simple drawing is NP-complete. Thus, the straight-forward approach of just trying to insert every possible edge is hopeless.

Problem 5. *What is the computational complexity of deciding whether (any) k edges can be inserted into a given 1-plane drawing?*

An FPT algorithm with respect to the parameter k was presented in Section 6.3.

Problem 6. *Does every simple drawing of $K_{n,n}$ contain a plane perfect matching as a sub-drawing?*

In Section 7.5 we proved that every simple drawing of $K_{m,n}$ contains a plane spanning tree, but the plane spanning tree that we construct might have few independent edges.

If this question is answered affirmatively, it implies that every simple drawing of K_n contains $\Theta(n)$ pairwise non-crossing edges. This would improve the currently best lower bound of $\Omega(n^{1/2-\varepsilon})$ pairwise non-crossing edges for any $\varepsilon > 0$ [178]. We remark that it is actually conjectured that every simple drawing of K_n contains a plane Hamiltonian cycle [171].

Problem 7. *What is the computational complexity of deciding semi-realizability of a given rotation system of K_n?*

The computational results in Section 8.6 seem to indicate that there is no finite set of obstructions, unlike for simple drawings.

List of Figures

Bibliography

[1] Bernardo M. Ábrego, Oswin Aichholzer, Silvia Fernández-Merchant, Thomas Hackl, Jürgen Pammer, Alexander Pilz, Pedro Ramos, Gelasio Salazar, and Birgit Vogtenhuber. "All good drawings of small complete graphs." In: *Proceedings of the 31st European Workshop on Computational Geometry (EuroCG'15)*. 2015, pages 57–60 (cited on pages 2, 27, 146, 150, 153).

[2] Bernardo M. Ábrego, Oswin Aichholzer, Silvia Fernández-Merchant, Dan McQuillan, Bojan Mohar, Petra Mutzel, Pedro Ramos, R. Bruce Richter, and Birgit Vogtenhuber. "Bishellable drawings of K_n." In: *SIAM Journal on Discrete Mathematics* 32.4 (2018), pages 2482–2492 (cited on page 1).

[3] Bernardo M. Ábrego, Oswin Aichholzer, Silvia Fernández-Merchant, Pedro Ramos, and Gelasio Salazar. "More on the crossing number of K_n: Monotone drawings." In: *Proceedings of the VII Latin-American Algorithms, Graphs, and Optimization Symposium (LAGOS'13)*. Volume 44. Electronic Notes in Discrete Mathematics. 2013, pages 411–414 (cited on page 1).

[4] Bernardo M. Ábrego, Oswin Aichholzer, Silvia Fernández-Merchant, Pedro Ramos, and Gelasio Salazar. "Shellable drawings and the cylindrical crossing number of K_n." In: *Discrete & Computational Geometry* 52.4 (2014), pages 743–753 (cited on page 1).

[5] Bernardo M. Ábrego, Oswin Aichholzer, Silvia Fernández-Merchant, Pedro Ramos, and Gelasio Salazar. "The 2-page crossing number of K_n." In: *Discrete & Computational Geometry* 49.4 (2013), pages 747–777 (cited on pages 1, 29).

[6] Bernardo M. Ábrego and Silvia Fernández-Merchant. "Geometric drawings of K_n with few crossings." In: *Journal of Combinatorial Theory, Series A* 114.2 (2007), pages 373–379 (cited on pages 10, 11, 15).

[7] Bernardo M. Ábrego, Silvia Fernández-Merchant, Jesús Leaños, and Gelasio Salazar. "A central approach to bound the number of crossings in a generalized configuration." In: *Proceedings of the IV Latin-American Algorithms, Graphs, and Optimization Symposium (LAGOS'07)*. Volume 30. Electronic Notes in Discrete Mathematics. 2008, pages 273–278 (cited on pages 9–11, 24).

[8] Bernardo M. Ábrego, Silvia Fernández-Merchant, and Gelasio Salazar. "The rectilinear crossing number of K_n: closing in (or are we?)" In: *Thirty Essays on Geometric Graph Theory*. Edited by János Pach. Springer, 2013, pages 5–18 (cited on page 3).

[9] Eyal Ackerman. "On topological graphs with at most four crossings per edge." In: *ArXiv e-Prints* (2013). arXiv: `1509.01932` (cited on page 25).

[10] Oswin Aichholzer. "[Empty] [colored] k-gons. Recent results on some Erdős–Szekeres type problems." In: *Proceedings of the XIII Encuentros de Geometría Computacional (EGC'09)*. 2009, pages 43–52 (cited on pages 4, 54).

[11] Oswin Aichholzer. *The order type database*. Last accessed: Aug. 12, 2019. URL: `http://www.ist.tugraz.at/aichholzer/research/rp/triangulations/ordertypes/` (cited on pages 3, 26).

[12] Oswin Aichholzer, Franz Aurenhammer, and Hannes Krasser. "Enumerating order types for small point sets with applications." In: *Order* 19.3 (2002), pages 265–281 (cited on pages 3, 26, 31, 44, 71).

[13] Oswin Aichholzer, Martin Balko, Thomas Hackl, Jan Kynčl, Irene Parada, Manfred Scheucher, Pavel Valtr, and Birgit Vogtenhuber. "A superlinear lower bound on the number of 5-holes." In: *Proceedings of the 33rd European Workshop on Computational Geometry (EuroCG'17)*. 2017, pages 69–72 (cited on page 4).

[14] Oswin Aichholzer, Martin Balko, Thomas Hackl, Jan Kynčl, Irene Parada, Manfred Scheucher, Pavel Valtr, and Birgit Vogtenhuber. "A superlinear lower bound on the number of 5-holes." In: *Proceedings of the 33rd International Symposium on Computational Geometry (SoCG'17)*. Volume 77. Leibniz International Proceedings in Informatics. Dagstuhl Publishing, 2017, 8:1–8:16 (cited on pages 4, 53).

[15] Oswin Aichholzer, Martin Balko, Thomas Hackl, Jan Kynčl, Irene Parada, Manfred Scheucher, Pavel Valtr, and Birgit Vogtenhuber. "A superlinear lower bound on the number of 5-holes." In: *Journal of Combinatorial Theory, Series A* (2019). Accepted (cited on pages 4, 53).

[16] Oswin Aichholzer, Martin Balko, Michael Hoffmann, Jan Kynčl, Wolfgang Mulzer, Irene Parada, Alexander Pilz, Manfred Scheucher, Pavel Valtr, Birgit Vogtenhuber, and Emo Welzl. "Minimal geometric graph representations of order types." In: *Proceedings of the 34th European Workshop on Computational Geometry (EuroCG'18)*. 2018, 21:1–21:6 (cited on page 4).

[17] Oswin Aichholzer, Martin Balko, Michael Hoffmann, Jan Kynčl, Wolfgang Mulzer, Irene Parada, Alexander Pilz, Manfred Scheucher, Pavel Valtr, Birgit Vogtenhuber, and Emo Welzl. "Minimal representations of order types by geometric graphs." In: *27th International Symposium on Graph Drawing and Network Visualization (GD'19)*. Lecture Notes in Computer Science. Springer, 2019. To appear (cited on pages 4, 33).

[18] Oswin Aichholzer, Frank Duque, Óscar Eduardo García-Quintero, Ruy Fabila-Monroy, and Carlos Hidalgo-Toscano. "An ongoing project to improve the rectilinear and pseudolinear crossing constants." In: *ArXiv e-Prints* (2019). arXiv: `1907.07796` (cited on pages 3, 9, 11, 22, 23).

[19] Oswin Aichholzer, Florian Ebenführer, Irene Parada, Alexander Pilz, and Birgit Vogtenhuber. "On semi-simple drawings of the complete graph." In: *Abstracts of the XVII Spanish Meeting on Computational Geometry (EGC'17)*. 2017, pages 25–28 (cited on pages 4, 145).

[20] Oswin Aichholzer, Ruy Fabila, Adrian Fuchs, Carlos Hidalgo, Irene Parada, Birgit Vogtenhuber, and Francisco Zaragoza. "On the 2-colored crossing number." In: *27th International Symposium on Graph Drawing and Network Visualization (GD'19)*. Lecture Notes in Computer Science. Springer, 2019. To appear (cited on pages 4, 9).

[21] Oswin Aichholzer, Ruy Fabila, Adrian Fuchs, Carlos Hidalgo, Irene Parada, Birgit Vogtenhuber, and Francisco Zaragoza. "On the 2-colored crossing number." In: *Proceedings of the 35th European Workshop on Computational Geometry (EuroCG'19)*. 2019, 56:1–56:7 (cited on page 4).

[22] Oswin Aichholzer, Ruy Fabila-Monroy, Thomas Hackl, Clemens Huemer, Alexander Pilz, and Birgit Vogtenhuber. "Lower bounds for the number of small convex k-holes." In: *Computational Geometry: Theory and Applications* 47.5 (2014), pages 605–613 (cited on pages 53, 54).

[23] Oswin Aichholzer, Alfredo García, Irene Parada, Javier Tejel, Birgit Vogtenhuber, and Alexandra Weinberger. "Plane spanning trees in simple drawings of $K_{m,n}$." Manuscript. 2019 (cited on pages 4, 131).

[24] Oswin Aichholzer, Thomas Hackl, Alexander Pilz, Pedro A. Ramos, Vera Sacristán, and Birgit Vogtenhuber. "Empty triangles in good drawings of the complete graph." In: *Graphs and Combinatorics* 31.2 (2015), pages 335–345 (cited on pages 2, 131).

[25] Oswin Aichholzer, Thomas Hackl, and Birgit Vogtenhuber. "On 5-gons and 5-holes." In: *Computational Geometry: Revised Selected Papers of the XIV Spanish Meeting on Computational Geometry (EGC'11)*. Volume 7579. Lecture Notes in Computer Science. Springer, 2012, pages 249–257 (cited on page 53).

[26] Oswin Aichholzer and Hannes Krasser. "Abstract order type extension and new results on the rectilinear crossing number." In: *Computational Geometry: Theory and Applications* 36.1 (2007), pages 2–15 (cited on pages 3, 26, 31, 44, 71).

[27] Oswin Aichholzer, Vincent Kusters, Wolfgang Mulzer, Alexander Pilz, and Manuel Wettstein. "An optimal algorithm for reconstructing point set order types from radial orderings." In: *International Journal of Computational Geometry & Applications* 27 (2017), pages 57–83 (cited on page 3).

[28] Oswin Aichholzer, Wolfgang Mulzer, Partick Schnider, and Birgit Vogtenhuber. "NP-completeness of max-cut for segment intersection graphs." In: *Proceedings of the 34th European Workshop on Computational Geometry (EuroCG'18)*. 2018, pages 1–6 (cited on pages 10, 157).

[29] Oswin Aichholzer, Irene Parada, Manfred Scheucher, Birgit Vogtenhuber, and Alexandra Weinberger. "Shooting stars in simple drawings of $K_{m,n}$." In: *Proceedings of the 35th European Workshop on Computational Geometry (EuroCG'19)*. 2019, 59:1–59:6 (cited on pages 4, 131).

[30] Martin Aigner and Günter M. Ziegler. *Proofs from THE BOOK*. 6th edition. Springer, 2018 (cited on page 96).

[31] Soroush Alamdari, Patrizio Angelini, Fidel Barrera-Cruz, Timothy M. Chan, Giordano Da Lozzo, Giuseppe Di Battista, Fabrizio Frati, Penny Haxell, Anna Lubiw, Maurizio Patrignani, Vincenzo Roselli, Sahil Singla, and Bryan T. Wilkinson. "How to morph planar graph drawings." In: *SIAM Journal on Computing* 46.2 (2017), pages 824–852 (cited on pages 34, 47).

[32] Paola Alimonti and Viggo Kann. "Some APX-completeness results for cubic graphs." In: *Theoretical Computer Science* 237.1 (2000), pages 123–134 (cited on pages 91, 95).

[33] Patrizio Angelini, Giuseppe Di Battista, Fabrizio Frati, Vít Jelínek, Jan Kratochvíl, Maurizio Patrignani, and Ignaz Rutter. "Testing planarity of partially embedded graphs." In: *ACM Transactions on Algorithms* 11.4 (2015), 32:1–32:42 (cited on pages 2, 86).

[34] Sanjeev Arora and Boaz Barak. *Computational Complexity - A Modern Approach*. Cambridge University Press, 2009 (cited on page 5).

[35] Alan Arroyo, Martin Derka, and Irene Parada. "Extending simple drawings." In: *27th International Symposium on Graph Drawing and Network Visualization (GD'19)*. Lecture Notes in Computer Science. Springer, 2019. To appear (cited on pages 4, 85).

[36] Alan Arroyo, Martin Derka, and Irene Parada. "Extending simple drawings." In: *Proceedings of the 34th European Workshop on Computational Geometry (EuroCG'19)*. 2019, 60:1–60:7 (cited on page 4).

[37] Alan Arroyo, Fabian Klute, Irene Parada, Raimund Seidel, Birgit Vogtenhuber, and Tilo Wiedera. "Extending simple drawings with one edge is hard." In: *ArXiv e-Prints* (2019). arXiv: `1909.07347` (cited on pages 4, 85).

[38] Alan Arroyo, Dan McQuillan, R. Bruce Richter, and Gelasio Salazar. "Drawings of K_n with the same rotation scheme are the same up to Reidemeister moves (Gioan's Theorem)." In: *Australasian Journal of Combinatorics* 67 (2017), pages 131–144 (cited on page 2).

[39] Christopher Auer, Franz J. Brandenburg, Andreas Gleißner, and Josef Reislhuber. "1-planarity of graphs with a rotation system." In: *Journal of Graph Algorithms and Applications* 19.1 (2015), pages 67–86 (cited on page 117).

[40] Alireza Bagheri and Mohammadreza Razzazi. "Planar straight-line point-set embedding of trees with partial embeddings." In: *Information Processing Letters* 110.12-13 (2010), pages 521–523 (cited on pages 2, 86).

[41] Martin Balko. `http://kam.mff.cuni.cz/~balko/superlinear5Holes`. Last accessed: Aug. 12, 2019 (cited on page 71).

[42] Martin Balko, Radoslav Fulek, and Jan Kynčl. "Crossing numbers and combinatorial characterization of monotone drawings of K_n." In: *Discrete & Computational Geometry* 53.1 (2015), pages 107–143 (cited on pages 1, 71, 145).

[43] József Balogh, Bernard Lidický, and Gelasio Salazar. "Closing in on Hill's conjecture." In: *SIAM Journal on Discrete Mathematics* 33.3 (2019), pages 1261–1276 (cited on page 1).

[44] Michael J. Bannister, Sergio Cabello, and David Eppstein. "Parameterized complexity of 1-planarity." In: *Journal of Graph Algorithms and Applications* 22.1 (2018), pages 23–49 (cited on page 117).

[45] Michael J. Bannister and David Eppstein. "Crossing minimization for 1-page and 2-page drawings of graphs with bounded treewidth." In: *Journal of Graph Algorithms and Applications* 22.4 (2018), pages 577–606 (cited on page 29).

[46] Imre Bárány and Zoltán Füredi. "Empty simplices in Euclidean space." In: *Canadian Mathematical Bulletin* 30.4 (1987), pages 436–445 (cited on pages 45, 53).

[47] Imre Bárány and Gyula Károlyi. "Problems and results around the Erdős–Szekeres convex polygon theorem." In: *Proceedings of the 5th Japanese Conference on Discrete and Computational Geometry (JCDCG'00)*. Volume 2098. Lecture Notes in Computer Science. Springer, 2001, pages 91–105 (cited on page 53).

[48] Imre Bárány and Pavel Valtr. "A positive fraction Erdős–Szekeres theorem." In: *Discrete & Computational Geometry* 19.3 (1998), pages 335–342 (cited on page 28).

[49] Imre Bárány and Pavel Valtr. "Planar point sets with a small number of empty convex polygons." In: *Studia Scientiarum Mathematicarum Hungarica* 41.2 (2004), pages 243–266 (cited on page 53).

[50] Joseph Battle, Frank Harary, and Yukihiro Kodama. "Every planar graph with nine points has a nonplanar complement." In: *Bulletin of The American Mathematical Society* 68.6 (1962), pages 569–571 (cited on page 27).

[51] Michael A. Bekos, Martin Gronemann, and Chrysanthi N. Raftopoulou. "Two-page book embeddings of 4-planar graphs." In: *Algorithmica* 75.1 (2016), pages 158–185 (cited on page 118).

[52] Mark de Berg and Amirali Khosravi. "Optimal binary space partitions in the plane." In: *International Journal of Computational Geometry & Applications* 22.03 (2010), pages 187–205 (cited on page 87).

[53] Jérémy Blanc. "The best polynomial bounds for the number of triangles in a simple arrangement of n pseudo-lines." In: *Geombinatorics*. Volume 21. 2011, pages 5–17 (cited on page 44).

[54] Jaroslav Blažek and Milan Koman. "A minimal problem concerning complete plane graphs." In: *Proceedings of the Symposium on the Theory of Graphs and Its Applications*. Czechoslovak Academy of Sciences, 1964, pages 113–117 (cited on page 1).

[55] Jürgen Bokowski and Bernd Sturmfels. "On the coordinatization of oriented matroids." In: *Discrete & Computational Geometry* 1 (1986), pages 293–306 (cited on page 34).

[56] Peter Brass, William O. J. Moser, and János Pach. *Research problems in discrete geometry*. Springer, 2005 (cited on page 54).

[57] Guido Brückner and Ignaz Rutter. "Partial and constrained level planarity." In: *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'17)*. 2017, pages 2000–2011 (cited on pages 2, 86).

[58] Sergio Cabello. "Planar embeddability of the vertices of a graph using a fixed point set is NP-hard." In: *Journal of Graph Algorithms and Applications* 10.2 (2006), pages 353–363 (cited on page 34).

[59] Sergio Cabello and Bojan Mohar. "Adding one edge to planar graphs makes crossing number and 1-planarity hard." In: *SIAM Journal on Computing* 42.5 (2013), pages 1803–1829 (cited on pages 86, 117).

[60] Grant Cairns and Yuri Nikolayevsky. "Bounds for generalized thrackles." In: *Discrete & Computational Geometry* 23.2 (2000), pages 191–206 (cited on page 155).

[61] Jean Cardinal and Stefan Felsner. "Topological drawings of complete bipartite graphs." In: *Journal of Computational Geometry* 9.1 (2018), pages 213–246 (cited on pages 85, 137, 146).

[62] Yair Caro. *New results on the independence number.* Technical report. Tel Aviv University, 1979 (cited on page 96).

[63] Steven Chaplick, Paul Dorbec, Jan Kratochvíl, Mickaël Montassier, and Juraj Stacho. "Contact representations of planar graphs: extending a partial representation is hard." In: *Proceedings of the 40th International Workshop on Graph-Theoretic Concepts in Computer Science (WG'14).* Volume 8747. Lecture Notes in Computer Science. Springer, 2014, pages 139–151 (cited on page 2).

[64] Steven Chaplick, Radoslav Fulek, and Pavel Klavík. "Extending partial representations of circle graphs." In: *Journal of Graph Theory* 91.4 (2019), pages 365–394 (cited on page 2).

[65] Steven Chaplick, Grzegorz Guśpiel, Grzegorz Gutowski, Tomasz Krawczyk, and Giuseppe Liotta. "The partial visibility representation extension problem." In: *Algorithmica* 80.8 (2018), pages 2286–2323 (cited on page 2).

[66] Fan R. K. Chung and Ron L. Graham. "Forced convex n-gons in the plane." In: *Discrete & Computational Geometry* 19.3 (1998), pages 367–371 (cited on page 4).

[67] Fan R. K. Chung, F. Thomson Leighton, and Arnold L. Rosenberg. "Embedding graphs in books: a layout problem with applications to VLSI design." In: *SIAM Journal on Algebraic Discrete Methods* 8.1 (1987), pages 33–58 (cited on page 29).

[68] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms.* 3rd edition. The MIT Press, 2009 (cited on page 5).

[69] Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms.* Springer, 2015 (cited on page 5).

[70] Éva Czabarka, Ondrej Sýkora, László A. Székely, and Imrich Vrt'o. "Biplanar crossing numbers I: a survey of results and problems." In: *More sets, graphs and numbers.* Edited by Ervin Győri, Gyula O. H. Katona, László Lovász, and Tamás Fleiner. Volume 15. Bolyai Society Mathematical Studies. Springer, 2006, pages 57–77 (cited on page 9).

[71] Éva Czabarka, Ondrej Sýkora, László A. Székely, and Imrich Vrt'o. "Biplanar crossing numbers II. Comparing crossing numbers and biplanar crossing numbers using the probabilistic method." In: *Random Structures & Algorithms* 33.4 (2008), pages 480–496 (cited on pages 9, 25).

[72] Giordano Da Lozzo, Giuseppe Di Battista, and Fabrizio Frati. "Extending upward planar graph drawings." In: *Proceedings of the 16th International Symposium Algorithms and Data Structures (WADS'19)*. Volume 11646. Lecture Notes in Computer Science. Springer, 2019, pages 339–352 (cited on pages 2, 86).

[73] Etienne de Klerk, John Maharry, Dmitrii V. Pasechnik, R. Bruce Richter, and Gelasio Salazar. "Improved bounds for the crossing numbers of $K_{m,n}$ and K_n." In: *SIAM Journal on Discrete Mathematics* 20.1 (2006), pages 189–202 (cited on page 1).

[74] Etienne de Klerk, Dmitrii Pasechnik, and Alexander Schrijver. "Reduction of symmetric semidefinite programs using the regular *-representation." In: *Mathematical Programming, Series B* 109.2–3 (2007), pages 613–624 (cited on page 1).

[75] Knut Dehnhardt. "Leere konvexe Vielecke in ebenen Punktmengen." In German. PhD thesis. TU Braunschweig, Germany, 1987 (cited on page 53).

[76] Erik Demaine. *6.892 Algorithmic lower bounds: Fun with hardness proofs*. `http://courses.csail.mit.edu/6.892/spring19/`. Last accessed: Aug. 12, 2019. 2019 (cited on page 5).

[77] Walter Didimo, Giuseppe Liotta, and Fabrizio Montecchiani. "A survey on graph drawing beyond planarity." In: *ACM Computing Surveys* 52.1 (2019), 4:1–4:37 (cited on pages 9, 117).

[78] Reinhard Diestel. *Graph theory*. 5th edition. Graduate texts in mathematics. Springer, 2016 (cited on page 5).

[79] Rodney G. Downey and Michael R. Fellows. *Fundamentals of parameterized complexity*. Texts in Computer Science. Springer, 2013 (cited on page 5).

[80] Rodney G. Downey and Michael R. Fellows. *Parameterized complexity*. Monographs in Computer Science. Springer, 1999 (cited on page 5).

[81] Vida Dujmović. "The utility of untangling." In: *Journal of Graph Algorithms and Applications* 21.1 (2017), pages 121–134 (cited on page 46).

[82] Florian Ebenführer. "Realizability of rotation systems." Master's thesis. Graz University of Technology, Austria, 2017 (cited on pages 145, 151).

[83] Paul Erdős. "Some more problems on elementary geometry." In: *Australian Mathematical Society Gazette* 5.2 (1978), pages 52–54 (cited on page 4).

[84] Paul Erdős. "Some old and new problems in combinatorial geometry." In: *Annals of Discrete Mathematics*. Edited by Moshe Rosenfeld and Joseph Zaks. Volume 20. North-Holland, 1984, pages 129–136 (cited on page 4).

[85] Paul Erdős and George Szekeres. "A combinatorial problem in geometry." In: *Compositio Mathematica* 2 (1935), pages 463–470 (cited on page 3).

[86] Kapali P. Eswaran and Robert E. Tarjan. “Augmentation problems.” In: *SIAM Journal on Computing* 5.4 (1976), pages 653–665 (cited on page 86).

[87] Ruy Fabila-Monroy and Clemens Huemer. “Order types of random point sets can be realized with small integer coordinates.” In: *Abstracts of the XVII Spanish Meeting on Computational Geometry (EGC’17)*. 2017, pages 73–77 (cited on page 34).

[88] Ruy Fabila-Monroy and Jorge López. “Computational search of small point sets with small rectilinear crossing number.” In: *Journal of Graph Algorithms and Applications* 18.3 (2014), pages 393–399 (cited on pages 3, 9).

[89] Stefan Felsner and Jacob E. Goodman. “Pseudoline arrangements.” In: *Handbook of Discrete and Computational Geometry*. Edited by Csaba D. Tóth, Joseph O’Rourke, and Jacob E. Goodman. 3rd edition. CRC Press, 2017, pages 125–157 (cited on pages 3, 34, 41).

[90] Stefan Felsner and Helmut Weil. “A theorem on higher Bruhat orders.” In: *Discrete & Computational Geometry* 23.1 (2000), pages 121–127 (cited on pages 51, 52).

[91] Stefan Felsner and Helmut Weil. “Sweeps, arrangements and signotopes.” In: *Discrete Applied Mathematics* 109.1–2 (2001), pages 67–94 (cited on page 71).

[92] Jörg Flum and Martin Grohe. *Parameterized complexity theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006 (cited on page 5).

[93] Jon Folkman and Jim Lawrence. “Oriented matroids.” In: *Journal of Combinatorial Theory, Series B* 25.2 (1978), pages 199–236 (cited on page 41).

[94] Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. *Kernelization: theory of parameterized preprocessing*. Cambridge University Press, 2019 (cited on page 5).

[95] Adrian Fuchs. “On the number of monochromatic crossings in rectilinear embeddings of complete graphs.” Master’s thesis. University of Technology Graz, Austria, 2019 (cited on page 26).

[96] Radoslav Fulek and János Pach. “Thrackles: an improved upper bound.” In: *Discrete Applied Mathematics* 259 (2019), pages 226–231 (cited on pages 154, 155).

[97] Radoslav Fulek and Andres J. Ruiz-Vargas. “Topological graphs: empty triangles and disjoint matchings.” In: *Proceedings of the 29th Annual Symposium on Computational Geometry (SoCG’13)*. ACM, 2013, pages 259–266 (cited on page 131).

[98] Zoltán Füredi and Ilona Palásti. “Arrangements of lines with a large number of triangles.” In: *Proceedings of the American Mathematical Society* 92.4 (1984), pages 561–566 (cited on page 43).

[99] Alfredo García. “A note on the number of empty triangles.” In: *Computational Geometry: Revised Selected Papers of the XIV Spanish Meeting on Computational Geometry (EGC’11)*. Volume 7579. Lecture Notes in Computer Science. Springer, 2012, pages 249–257 (cited on pages 53, 54, 56, 82).

[100] Michael R. Garey and David S. Johnson. *Computers and intractability: A guide to the theory of NP-completeness*. W. H. Freeman & Co., 1979 (cited on page 5).

[101] Michael R. Garey and David S. Johnson. "The rectilinear Steiner tree problem is NP-complete." In: *SIAM Journal on Applied Mathematics* 32.4 (1977), pages 826–834 (cited on page 118).

[102] Michael Garey and David S. Johnson. "Crossing number is NP-complete." In: *SIAM Journal on Algebraic and Discrete Methods* 4.3 (1983), pages 312–316 (cited on page 2).

[103] Tobias Gerken. "Empty convex hexagons in planar point sets." In: *Discrete & Computational Geometry* 39.1–3 (2008), pages 239–272 (cited on page 4).

[104] Emeric Gioan. "Complete graph drawings up to triangle mutations." In: *Proceedings of the 31st International Workshop on Graph-Theoretic Concepts in Computer Science (WG'05).* Volume 3787. Lecture Notes in Computer Science. Springer, 2005, pages 139–150 (cited on page 2).

[105] Jacob E. Goodman and Richard Pollack. "Multidimensional sorting." In: *SIAM Journal on Computing* 12.3 (1983), pages 484–507 (cited on pages 3, 33).

[106] Jacob E. Goodman, Richard Pollack, and Bernd Sturmfels. "Coordinate representation of order types requires exponential storage." In: *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC'89).* ACM, 1989, pages 405–410 (cited on page 34).

[107] Alexander Grigoriev and Hans L. Bodlaender. "Algorithms for graphs embeddable with few crossings per edge." In: *Algorithmica* 49.1 (2007), pages 1–11 (cited on page 117).

[108] Branko Grünbaum. *Arrangements and spreads.* American Mathematical Society, 1972 (cited on pages 34, 42, 43).

[109] Carsten Gutwenger, Petra Mutzel, and René Weiskircher. "Inserting an edge into a planar graph." In: *Algorithmica* 41.4 (2005), pages 289–308 (cited on page 86).

[110] Richard K. Guy. "Crossing numbers of graphs." In: *Graph Theory and Applications: Proceedings of the Conference at Western Michigan University.* Springer, 1972, pages 111–124 (cited on page 1).

[111] Richard K. Guy. "The decline and fall of Zarankiewicz's theorem." In: 1969, pages 87–91 (cited on page 1).

[112] Péter Hajnal, Alexander Igamberdiev, Günter Rote, and André Schulz. "Saturated simple and 2-simple topological graphs with few edges." In: *Journal of Graph Algorithms and Applications* 22.1 (2018), pages 117–138 (cited on page 86).

[113] Frank Harary and Anthony Hill. "On the number of crossings in a complete graph." In: *Proceedings of the Edinburgh Mathematical Society* 13.4 (1963), pages 333–338 (cited on page 1).

[114] Heiko Harborth. "Empty triangles in drawings of the complete graph." In: *Discrete Mathematics* 191.1–3 (1998), pages 109–111 (cited on page 131).

[115] Heiko Harborth. "Konvexe Fünfecke in ebenen Punktmengen." In: *Elemente der Mathematik* 33 (1978). In German, pages 116–118 (cited on pages 4, 53, 55, 60, 77).

[116] Heiko Harborth. "Some simple arrangements of pseudolines with a maximum number of triangles." In: *Annals of the New York Academy of Sciences* 440.1 (1985), pages 31–33 (cited on page 44).

[117] Heiko Harborth and Ingrid Mengersen. "Edges without crossings in drawings of complete graphs." In: *Journal of Combinatorial Theory, Series B* 17.3 (1974), pages 299–311 (cited on page 132).

[118] Lenwood Scott Heath. "Algorithms for embedding graphs in books." PhD thesis. University of North Carolina (USA), 1985 (cited on page 118).

[119] César Hernández-Vélez, Jesús Leaños, and Gelasio Salazar. "On the pseudolinear crossing number." In: *Journal of Graph Theory* 84.3 (2016), pages 297–310 (cited on page 3).

[120] Sok-Hee Hong, Michael Kaufmann, Stephen G. Kobourov, and János Pach. "Beyond-planar graphs: algorithmics and combinatorics (Dagstuhl seminar 16452)." In: *Dagstuhl Reports* 6.11 (2017), pages 35–62 (cited on page 9).

[121] John E. Hopcroft and Richard M. Karp. "An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs." In: *SIAM Journal on computing* 2.4 (1973), pages 225–231 (cited on page 127).

[122] Joseph D. Horton. "Sets with no empty convex 7-gons." In: *Canadian Mathematical Bulletin* 26.4 (1983), pages 482–484 (cited on pages 4, 53).

[123] Vít Jelínek, Jan Kratochvíl, and Ignaz Rutter. "A Kuratowski-type theorem for planarity of partially embedded graphs." In: *Computational Geometry: Theory and Applications* 46.4 (2013), pages 466–492 (cited on pages 2, 86).

[124] Gyula Károlyi, János Pach, and Géza Tóth. "Ramsey-type results for geometric graphs, I." In: *Discrete & Computational Geometry* 18.3 (1997), pages 247–255 (cited on page 131).

[125] Pavel Klavík, Jan Kratochvíl, Tomasz Krawczyk, and Bartosz Walczak. "Extending partial representations of function graphs and permutation graphs." In: *Proceedings of the 20th European Symposium on Algorithms (ESA'12)*. Volume 7501. Lecture Notes in Computer Science. Springer, 2012, pages 671–682 (cited on page 2).

[126] Pavel Klavík, Jan Kratochvíl, Yota Otachi, Ignaz Rutter, Toshiki Saitoh, Maria Saumell, and Tomás Vyskocil. "Extending partial representations of proper and unit interval graphs." In: *Algorithmica* 77.4 (2017), pages 1071–1104 (cited on page 2).

[127] Pavel Klavík, Jan Kratochvíl, Yota Otachi, and Toshiki Saitoh. "Extending partial representations of subclasses of chordal graphs." In: *Theoretical Computer Science* 576 (2015), pages 85–101 (cited on page 2).

[128] Pavel Klavík, Jan Kratochvíl, Yota Otachi, Toshiki Saitoh, and Tomás Vyskocil. "Extending partial representations of interval graphs." In: *Algorithmica* 78.3 (2017), pages 945–967 (cited on page 2).

[129] Daniel J. Kleitman. "The crossing number of $K_{5,n}$." In: *Journal of Combinatorial Theory* 9.4 (1971), pages 315–323 (cited on page 1).

[130] Daniel J. Kleitman and Lior Pachter. "Finding convex sets among points in the plane." In: *Discrete & Computational Geometry* 19.3 (1998), pages 405–410 (cited on page 4).

[131] Fabian Klute, Irene Parada, and Thekla Hamm. "Extending to 1-plane drawings." In: *Abstracts of the XVIII Spanish Meeting on Computational Geometry (EGC'19)*. 2019, pages 30–33 (cited on pages 4, 117).

[132] Stephen G. Kobourov, Giuseppe Liotta, and Fabrizio Montecchiani. "An annotated bibliography on 1-planarity." In: *Computer Science Review* 25 (2017), pages 49–67 (cited on pages 9, 117).

[133] Vladimir P. Korzhik. "Minimal non-1-planar graphs." In: *Discrete Mathematics* 308.7 (2008), pages 1319–1327 (cited on page 123).

[134] Vladimir P. Korzhik and Bojan Mohar. "Minimal obstructions for 1-immersions and hardness of 1-planarity testing." In: *Journal of Graph Theory* 72.1 (2013), pages 30–71 (cited on page 117).

[135] Jan Kynčl. "Enumeration of simple complete topological graphs." In: *European Journal of Combinatorics* 30.7 (2009), pages 1676–1685 (cited on page 145).

[136] Jan Kynčl. "Improved enumeration of simple topological graphs." In: *Discrete & Computational Geometry* 50.3 (2013), pages 727–770 (cited on pages 85, 86).

[137] Jan Kynčl. "Simple realizability of complete abstract topological graphs in P." In: *Discrete & Computational Geometry* 45.3 (2011), pages 383–399 (cited on page 2).

[138] Jan Kynčl. "Simple realizability of complete abstract topological graphs simplified." In: *Proceeding of the 23rd International Symposium on Graph Drawing and Network Visualization (GD'15).* Volume 9411. Lecture Notes in Computer Science. Springer, 2015, pages 309–320 (cited on pages 2, 85, 146, 155).

[139] Jan Kynčl. "Simple realizability of complete abstract topological graphs simplified." In: *ArXiv e-Prints* (2016). arXiv: `1608.05867` (cited on page 153).

[140] Jan Kynčl, János Pach, Radoš Radoičić, and Géza Tóth. "Saturated simple and k-simple topological graphs." In: *Computational Geometry: Theory and Applications* 48.4 (2015), pages 295–310 (cited on pages 85, 86, 100, 101).

[141] Friedrich Levi. "Die Teilung der projektiven Ebene durch Gerade oder Pseudogerade." In: *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig, Mathematisch-Physische Klasse* 78 (1926). In German, pages 256–267 (cited on page 42).

[142] Dragoslav Ljubić, Jean-Pierre Roudneff, and Bernd Sturmfels. "Arrangements of lines and pseudolines without adjacent triangles." In: *Journal of Combinatorial Theory, Series A* 50.1 (1989), pages 24–32 (cited on page 44).

[143] László Lovász, János Pach, and Mario Szegedy. "On Conway's thrackle conjecture." In: *Discrete & Computational Geometry* 18.4 (1997), pages 369–376 (cited on page 154).

[144] László Lovász, Katalin Vesztergombi, Uli Wagner, and Emo Welzl. "Convex quadrilaterals and k-sets." In: *Towards a theory of geometric graphs.* Edited by János Pach. Volume 342. Contemporary Mathematics. American Mathematical Society, 2004 (cited on page 3).

[145] Jiří Matoušek. *Lectures on discrete geometry.* Volume 212. Graduate Texts in Mathematics. Springer, 2002 (cited on pages 25, 28).

[146] Tamara Mchedlidze, Martin Nöllenburg, and Ignaz Rutter. "Extending convex partial drawings of graphs." In: *Algorithmica* 76.1 (2015), pages 47–67 (cited on pages 2, 86).

[147] Silvio Micali and Vijay V. Vazirani. "An $O(\sqrt{|V|}|E|)$ algorithm for finding maximum matching in general graphs." In: *Proceedings of the 21st Annual Symposium on Foundations of Computer Science (FOCS'80)*. Volume 21. IEEE, 1980, pages 17–27 (cited on page 127).

[148] Nicolai E. Mnëv. "The universality theorems on the classification problem of configuration varieties and convex polytope varieties." In: *Topology and Geometry–Rohlin Seminar*. Volume 1346. Lecture Notes in Mathematics. Springer, 1988, pages 527–544 (cited on pages 34, 38, 41).

[149] Petra Mutzel and Lutz Oettershagen. "The crossing number of seq-shellable drawings of complete graphs." In: *Proceedings of the International Workshop on Combinatorial Algorithms (IWOCA'18)*. Volume 10979. Lecture Notes in Computer Science. Springer, 2018, pages 273–284 (cited on page 1).

[150] Hossein Nassajian Mojarrad and Georgios Vlachos. "An improved upper bound for the Erdős–Szekeres conjecture." In: *ArXiv e-Prints* (2015). arXiv: `1510.06255` (cited on page 4).

[151] Carlos M. Nicolás. "The empty hexagon theorem." In: *Discrete & Computational Geometry* 38.2 (2007), pages 389–397 (cited on page 4).

[152] Rolf Niedermeier. *Invitation to fixed-parameter algorithms*. Oxford University Press, 2006 (cited on page 5).

[153] Sergey Norin. *Turán brickyard problem and flag algebras*. Presentation at the BIRS workshop on geometric and topological graph theory (13w5091). Last accessed: Aug. 12, 2019. 2013 (cited on page 1).

[154] Sergey Norin and Yelena Yuditsky. "Erdős–Szekeres without induction." In: *Discrete & Computational Geometry* 55.4 (2016), pages 963–971 (cited on page 4).

[155] Mark Overmars. "Finding sets of points without empty convex 6-gons." In: *Discrete & Computational Geometry* 29.1 (2002), pages 153–158 (cited on page 78).

[156] Andrew Owens. "On the biplanar crossing number." In: *IEEE Transactions on Circuit Theory* 18.2 (1971), pages 277–280 (cited on page 9).

[157] János Pach, Radoš Radoičić, and Géza Tóth. "Relaxing planarity for topological graphs." In: *Proceedings of the 7th Japanese Conference on Discrete and Computational Geometry (JCDCG'02)*. Volume 2866. Lecture Notes in Computer Science. Springer, 2003, pages 221–232 (cited on page 145).

[158] János Pach, József Solymosi, and Géza Tóth. "Unavoidable configurations in complete topological graphs." In: *Discrete & Computational Geometry* 30.2 (2003), pages 311–320 (cited on pages 85, 131).

[159] János Pach, László A. Székely, Csaba D. Tóth, and Géza Tóth. "Note on k-planar crossing numbers." In: *Computational Geometry: Theory and Applications* 68 (2018), pages 2–6 (cited on pages 9, 24, 26, 27).

[160] János Pach and Géza Tóth. "Disjoint edges in topological graphs." In: *Journal of Combinatorics* 1.3 (2010), pages 335–344 (cited on page 131).

[161] János Pach and Géza Tóth. "Graphs drawn with few crossings per edge." In: *Combinatorica* 17.3 (1997), pages 427–439 (cited on page 128).

[162] János Pach and Géza Tóth. "How many ways can one draw a graph?" In: *Combinatorica* 26.5 (2006), pages 559–576 (cited on page 2).

[163] János Pach and Géza Tóth. "Which crossing number is it anyway?" In: *Journal of Combinatorial Theory, Series B* 80.2 (2000), pages 225–246 (cited on page 145).

[164] Shengjun Pan and R. Bruce Richter. "The crossing number of K_{11} is 100." In: *Journal of Graph Theory* 56.2 (2007), pages 128–134 (cited on pages 1, 2).

[165] Christos H. Papadimitriou. *Computational complexity.* Addison-Wesley, 1994 (cited on page 5).

[166] Christos H. Papadimitriou and Mihalis Yannakakis. "Optimization, approximation, and complexity classes." In: *Journal of Computer and System Sciences* 43.3 (1991), pages 425–440 (cited on page 95).

[167] Maurizio Patrignani. "On extending a partial straight-line drawing." In: *International Journal of Foundations of Computer Science* 17.5 (2006), pages 1061–1070 (cited on pages 2, 86).

[168] Michael J. Pelsmajer, Marcus Schaefer, and Daniel Stefankovic. "Removing even crossings." In: *Journal of Combinatorial Theory, Series B* 97.4 (2007), pages 489–500 (cited on page 145).

[169] Michael J. Pelsmajer, Marcus Schaefer, and Daniel Štefankovič. "Crossing numbers of graphs with rotation systems." In: *Algorithmica* 60.3 (2011), pages 679–702 (cited on page 2).

[170] Rom Pinchasi, Radoš Radoičić, and Micha Sharir. "On empty convex polygons in a planar point set." In: *Journal of Combinatorial Theory, Series A* 113.3 (2006), pages 385–419 (cited on page 54).

[171] Nabil H. Rafla. "The good drawings D_n of the complete graph K_n." PhD thesis. McGill University, Canada, 1988 (cited on pages 131, 154, 158).

[172] R. Bruce Richter and Carsten Thomassen. "Relations between crossing numbers of complete and complete bipartite graphs." In: *The American Mathematical Monthly* 104.2 (1997), pages 131–137 (cited on page 24).

[173] Jürgen Richter-Gebert and Günter M. Ziegler. "Oriented matroids." In: *Handbook of Discrete and Computational Geometry.* Edited by Csaba D. Tóth, Joseph O'Rourke, and Jacob E. Goodman. 3rd edition. CRC Press, 2017, pages 159–184 (cited on page 41).

[174] Gerhard Ringel. "Ein Sechsfarbenproblem auf der Kugel." In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg.* Volume 29. 1. Springer, 1965, pages 107–117 (cited on page 2).

[175] Gerhard Ringel. "Teilungen der Ebene durch Geraden oder topologische Geraden." In: *Mathematische Zeitschrift* 64 (1956), pages 79–102 (cited on pages 34, 38).

[176] Adrian Riskin. "The crossing number of a cubic plane polyhedral map plus an edge." In: *Studia Scientiarum Mathematicarum Hungarica* 31.4 (1996), pages 405–414 (cited on page 86).

[177] Jean-Pierre Roudneff. “On the number of triangles in simple arrangements of pseudolines in the real projective plane.” In: *Discrete Mathematics* 60 (1986), pages 243–251 (cited on page 44).

[178] Andres J. Ruiz-Vargas. “Many disjoint edges in topological graphs.” In: *Computational Geometry: Theory and Applications* 62 (2017), pages 1–13 (cited on pages 2, 131, 158).

[179] Marcus Schaefer. “The graph crossing number and its variants: a survey.” In: *Electronic Journal of Combinatorics* Dynamic Survey 21 (2013/2017) (cited on pages 9, 145).

[180] Manfred Scheucher. `http://www.ist.tugraz.at/scheucher/5holes`. Last accessed: Aug. 12, 2019 (cited on pages 54, 71).

[181] Manfred Scheucher. “Counting convex 5-holes.” In German. Bachelor’s thesis. Graz University of Technology, Austria, 2013 (cited on page 54).

[182] Felix Schröder. *Personal communication.* 2019 (cited on page 132).

[183] Farhad Shahrokhi, Ondrej Sykora, László A. Székely, and Imrich Vrt’o. “On k-planar crossing numbers.” In: *Discrete Applied Mathematics* 155.9 (2007), pages 1106–1115 (cited on page 9).

[184] Michael Sipser. *Introduction to the theory of computation.* 3rd edition. Cengage learning, 2012 (cited on page 5).

[185] William Steiger and Jihui Zhao. “Generalized ham-sandwich cuts.” In: *Discrete & Computational Geometry* 44.3 (2010), pages 535–545 (cited on page 55).

[186] Andrew Suk. “Disjoint edges in complete topological graphs.” In: *Discrete & Computational Geometry* 49.2 (2013), pages 280–286 (cited on page 131).

[187] Andrew Suk. “On the Erdős–Szekeres convex polygon problem.” In: *Journal of the American Mathematical Society* 30 (2017), pages 1047–1053 (cited on page 4).

[188] Csaba D. Tóth. *Personal communication.* 2018 (cited on page 26).

[189] Géza Tóth and Pavel Valtr. “Note on the Erdős–Szekeres theorem.” In: *Discrete & Computational Geometry* 19.3 (1998), pages 457–459 (cited on page 4).

[190] Paul Turán. “A note of welcome.” In: *Journal of Graph Theory* 1 (1977), pages 7–9 (cited on page 1).

[191] Paul Turán. “On an extremal problem in graph theory.” In: *Matematikai és Fizikai Lapok* 48 (1941), pages 436–452 (cited on pages 96, 123).

[192] William T. Tutte. “The non-biplanar character of the complete 9-graph.” In: *Canadian Mathematical Bulletin* 6.3 (1963), pages 319–330 (cited on page 27).

[193] Pavel Valtr. “Convex independent sets and 7-holes in restricted planar point sets.” In: *Discrete & Computational Geometry* 7.2 (1992), pages 135–152 (cited on pages 4, 53).

[194] Pavel Valtr. “On empty pentagons and hexagons in planar point sets.” In: *Proceedings of Computing: The Eighteenth Australasian Theory Symposium (CATS’12).* Volume 128. ACS, 2012, pages 47–48 (cited on page 53).

[195] Pavel Valtr. “Sets in $\mathbb{R}^d$ with no large empty convex subsets.” In: *Discrete Mathematics* 108.1 (1992), pages 115–124 (cited on pages 4, 53).

[196] Georgios Vlachos. "On a conjecture of Erdős and Szekeres." In: *ArXiv e-Prints* (2015). arXiv: `1505.07549` (cited on page 4).

[197] Ingo Wegener and R. Pruim. *Complexity theory: Exploring the limits of efficient algorithms*. Springer, 2005 (cited on page 5).

[198] Victor K. Wei. *A lower bound on the stability number of a simple graph*. Technical report 81–11217–9. Bell Laboratories, 1981 (cited on page 96).

[199] Alexandra Weinberger. "Simple drawings and rotations systems beyond the complete graph." Master's thesis. Graz University of Technology, Austria, 2019 (cited on page 131).

[200] Douglas R. Woodall. "Cyclic-order graphs and Zarankiewicz's crossing-number conjecture." In: *Journal of Graph Theory* 17.6 (1993), pages 657–671 (cited on page 1).

Curriculum vitae of Irene Parada

Contact and personal data

Address:	Inffeldgasse 16b, 8010 Graz (Austria).	Date and place of birth:	November 10, 1992. Spain.
		Sex:	Female.
Email:	iparada@ist.tugraz.at	Website:	www.ist.tugraz.at/parada

Education

Aug. 2015–present	PhD student in *Discrete and Computational Geometry*, Graz University of Technology, Austria. Advisor: Prof. O. Aichholzer.
Sep. 2014–Jul. 2015	Master's degree in *Advanced Mathematics and Mathematical Engineering* (60 ECTS), Universitat Politècnica de Catalunya (UPC), Spain.
Fall 2013	Exchange semester at the *Royal Institute of Technology (KTH)*, Sweden.
Sep. 2010–Jul. 2014	Bachelor's degree in *Mathematical Engineering* (240 ECTS), Universidad Complutense de Madrid (UCM), Spain.
Sep. 2011–present	Bachelor's degree in *Physics* (155/240 ECTS), Universidad Nacional de Educación a Distancia (UNED), the public open university of Spain.

Awards and scholarships

Spring 2015	Scholarship for academic support activities (AAD3) at the Universitat Politècnica de Catalunya (UPC).
2013–2014	Best academic record of the bachelor's degree.
Fall 2013	Erasmus grant.
2010–2011	Excellence scholarship of Madrid.
2010	Bronze medal in the Spanish Physics Olympiad.

Projects

Jul. 2019–present	*FWF DACH project "Arrangements and Drawings".* Project assistant.
2017–present	*H2020-MSCA-RISE project "CONNECT".* Project member and management support.
Aug. 2015–Jul. 2019	*FWF DK Doctoral Program "Discrete Mathematics".* Project assistant and PhD candidate.
2018	*FWF-JSPS bilateral Japanese-Austrian Joint Seminar.* Project member.
2015–2016	*OeAD-CZ "Erdős-Szekeres type questions for point sets".* Project member.

Academic activities

Talks

Invited talks	Congreso Bienal de la RSME 2019, Spain. Taller de Teoría de Gráficas y Geometría Discreta 2017, Mexico.
Conference talks	2 × GD'19, EGC'19, EuroCG'19, EGC'17, ICRA'16, EGC'15.
Seminar talks	2 × TU Graz'19, UAB'18, 2 × TU Graz'18, UASLP'17, Cinestav'17, TU Graz'17, TU Graz'15.

Invitation-only workshops

2015–2019	2nd, 3rd, 4th, 5th, and 6th *Austrian-Japanese-Mexican-Spanish Workshop on Discrete Geometry, Spain-Mexico-Spain-Spain-Austria.*
2016–2019	8th, 9th, and 11th *Crossing Numbers Workshop, Austria-Germany-Czech Republic.*
2017–2019	32nd, 33rd, and 34th *Bellairs Winter Workshop on Computational Geometry, Barbados.*
2018–2019	6th and 7th *Annual Workshop on Geometry and Graphs, Barbados.*
2017–2018	16th and 17th *Workshop: Routing in Ensenada, Mexico.*
2017–2018	2nd and 3rd *Renunión de Optimización, Matemáticas y Algoritmos, Mexico.*
2016	13th *European Research Week on Geometric Graphs, Spain.*
2015	*Workshop on Sidedness Queries, Austria.*

Extended research stays

2017–2019	2 months visiting Prof. *R. Fabila* at Cinvestav, Mexico.
2017–2019	2 months visiting Prof. *G. Salazar* at UASLP, Mexico.
2017–2018	2 months visiting Prof. *J. Urrutia* at UNAM, Mexico.
2015–2016	4 weeks visiting Prof. *P. Valtr* at Charles University, Czech Republic.

Teaching and student supervision

2017–2019	Co-supervision of two master's theses.
Fall 2019	Lecturer in a master's course.
2018–2019	Guest lecturer in a bachelor's and a master's course.
Spring 2015	Teaching assistant in three bachelor's courses.

Reviews

2019 Ars Combinatoria, EGC, EuroCG, GD, SoCG.

2017 CCCG, EGC, EuroCG, IROS.

2016 EuroCG, FUN.

Publications

In this list, publications which have appeared in multiple versions, for example in proceedings of a conference and in a journal, are grouped together. Following the standard practice in the discrete and computational geometry community, the co-authors of a paper are sorted alphabetically. The publications are sorted chronologically (from the newest to the oldest) and alphabetically by the authors' last names, in that order. Accepted papers that have not yet appeared are considered to be the most recent ones.

Publications marked in red are part of this thesis.

Fourteen different peer reviewed publications

- **A superlinear lower bound on the number of 5-holes.**
  Oswin Aichholzer, Martin Balko, Thomas Hackl, Jan Kynčl, Irene Parada, Manfred Scheucher, Pavel Valtr, and Birgit Vogtenhuber. In: *Journal of Combinatorial Theory, Series A*. 2019. Accepted.

  Also appeared as conference paper:
  A superlinear lower bound on the number of 5-holes.
  Oswin Aichholzer, Martin Balko, Thomas Hackl, Jan Kynčl, Irene Parada, Manfred Scheucher, Pavel Valtr, and Birgit Vogtenhuber. In: *Proceedings of the 33rd International Symposium on Computational Geometry (SoCG'17)*. Vol. 77. Leibniz International Proceedings in Informatics. Dagstuhl Publishing, 2017, 8:1–8:16.

  Also appeared as extended abstract:
  A superlinear lower bound on the number of 5-holes.
  Oswin Aichholzer, Martin Balko, Thomas Hackl, Jan Kynčl, Irene Parada, Manfred Scheucher, Pavel Valtr, and Birgit Vogtenhuber. In: *Proceedings of the 33rd European Workshop on Computational Geometry (EuroCG'17)*. 2017, pp. 69–72.

- **Minimal representations of order types by geometric graphs.**
  Oswin Aichholzer, Martin Balko, Michael Hoffmann, Jan Kynčl, Wolfgang Mulzer, Irene Parada, Alexander Pilz, Manfred Scheucher, Pavel Valtr, Birgit Vogtenhuber, and Emo Welzl. In: *Proceedings of the 27th International Symposium on Graph Drawing and Network Visualization (GD'19)*. Lecture Notes in Computer Science. Springer, 2019. To appear.

Also appeared as extended abstract:
Minimal geometric graph representations of order types.
Oswin Aichholzer, Martin Balko, Michael Hoffmann, Jan Kynčl, Wolfgang Mulzer, Irene Parada, Alexander Pilz, Manfred Scheucher, Pavel Valtr, Birgit Vogtenhuber, and Emo Welzl. In: *Proceedings of the 34th European Workshop on Computational Geometry (EuroCG'18)*. 2018, 21:1–21:6.

- **On the 2-colored crossing number.**
  Oswin Aichholzer, Ruy Fabila, Adrian Fuchs, Carlos Hidalgo, Irene Parada, Birgit Vogtenhuber, and Francisco Zaragoza. In: *Proceedings of the 27th International Symposium on Graph Drawing and Network Visualization (GD'19)*. Lecture Notes in Computer Science. Springer, 2019. To appear.

  Also appeared as extended abstract:
  On the 2-colored crossing number.
  Oswin Aichholzer, Ruy Fabila, Adrian Fuchs, Carlos Hidalgo, Irene Parada, Birgit Vogtenhuber, and Francisco Zaragoza. In: *Proceedings of the 35th European Workshop on Computational Geometry (EuroCG'19)*. 2019, 56:1–56:7.

- **Graphs with large total angular resolution.**
  Oswin Aichholzer, Matias Korman, Yoshio Okamoto, Irene Parada, Daniel Perz, André van Renssen, and Birgit Vogtenhuber. In: *Proceedings of the 27th International Symposium on Graph Drawing and Network Visualization (GD'19)*. Lecture Notes in Computer Science. Springer, 2019. To appear.

- **Extending simple drawings.**
  Alan Arroyo, Martin Derka, and Irene Parada. In: *Proceedings of the 27th International Symposium on Graph Drawing and Network Visualization (GD'19)*. Lecture Notes in Computer Science. Springer, 2019. To appear.

  Also appeared as extended abstract:
  Extending simple drawings.
  Alan Arroyo, Martin Derka, and Irene Parada. In: *Proceedings of the 35th European Workshop on Computational Geometry (EuroCG'19)*. 2019, 60:1–60:7.

- **Hamiltonian meander paths and cycles on bichromatic point sets.**
  Oswin Aichholzer, Carlos Alegría, Irene Parada, Alexander Pilz, Javier Tejel, Csaba D. Tóth, Jorge Urrutia, and Birgit Vogtenhuber. In: *Abstracts of the XVIII Spanish Meeting on Computational Geometry (EGC'19)*. 2019, pp. 35–38.

- **Shooting stars in simple drawings of $K_{m,n}$.**
  Oswin Aichholzer, Irene Parada, Manfred Scheucher, Birgit Vogtenhuber, and Alexandra Weinberger. In: *Proceedings of the 35th European Workshop on Computational Geometry (EuroCG'19)*. 2019, 59:1–59:6.

- **Universal reconfiguration of facet-connected modular robots by pivots: the O(1) musketeers**.
  Hugo Akitaya, Esther Arkin, Mirela Damian, Erik D. Demaine, Vida Dujmović, Robin Flatland, Matias Korman, Belen Palop, Irene Parada, André van Renssen, and Vera Sacristán. In: *Proceedings of the 27th Annual European Symposium on Algorithms (ESA'19)*. Vol. 144. Leibniz International Proceedings in Informatics. 2019, 3:1–3:14.

  Also appeared as extended abstract:
  Reconfiguring edge-connected pivoting modular robots.
  Hugo Akitaya, Esther Arkin, Mirela Damian, Erik D. Demaine, Vida Dujmović, Robin Flatland, Matias Korman, Belen Palop, Irene Parada, André van Renssen, and Vera Sacristán. In: *Abstracts of the XVIII Spanish Meeting on Computational Geometry (EGC'19)*. 2019, p. 64.

- **Efficient segment folding is hard**.
  Takashi Horiyama, Fabian Klute, Matias Korman, Irene Parada, Ryuhei Uehara, and Katsuhisa Yamanaka. In: *Proceedings of the 31st Canadian Conference in Computational Geometry (CCCG'19)*. 2019, pp. 182–188.

- **Extending to 1-plane drawings**.
  Fabian Klute, Irene Parada, and Thekla Hamm. In: *Abstracts of the XVIII Spanish Meeting on Computational Geometry (EGC'19)*. 2019, pp. 30–33.

- **Bounding the number of crossings for a particular class of drawings of $K_{n,n}$**.
  Carolina Medina, Irene Parada, Gelasio Salazar, and Birgit Vogtenhuber. In: *Abstracts of the XVIII Spanish Meeting on Computational Geometry (EGC'19)*. 2019, p. 34.

- **How to fit a tree in a box**.
  Hugo Akitaya, Maarten Löffler, and Irene Parada. In: *Proceedings of the 26th International Symposium on Graph Drawing and Network Visualization (GD'18)*. Vol. 11282. Lecture Notes in Computer Science. Springer, 2018, pp. 361–367.

- **On semi-simple drawings of the complete graph**.
  Oswin Aichholzer, Florian Ebenführer, Irene Parada, Alexander Pilz, and Birgit Vogtenhuber. In: *Abstracts of the XVII Spanish Meeting on Computational Geometry (EGC'17)*. 2017, pp. 25–28.

- **A new meta-module for efficient reconfiguration of hinged-units modular robots**.
  Irene Parada, Vera Sacristán, and Rodrigo I. Silveira. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'16), IEEE*. 2016, pp. 5197–5202.

  Also appeared as extended abstract:
  A new meta-module for efficient reconfiguration of hinged-units modular robots.
  Irene Parada, Vera Sacristán, and Rodrigo I. Silveira. In: *Abstracts of the XVI Spanish Meeting on Computational Geometry (EGC'15)*. 2015, pp. 49–52.

Other manuscripts and preprints related to this thesis

- **Plane spanning trees in simple drawings of $K_{m,n}$.**
  Oswin Aichholzer, Alfredo García, Irene Parada, Javier Tejel, Birgit Vogtenhuber, and Alexandra Weinberger. 2019. Manuscript.

- **Extending simple drawings with one edge is hard.**
  Alan Arroyo, Fabian Klute, Irene Parada, Raimund Seidel, Birgit Vogtenhuber, and Tilo Wiedera. In: *ArXiv e-Prints*. 2019. arXiv: `1909.07347`.

Theses

- **Systematic strategies for 3-dimensional modular robots**. Irene Parada. Master's thesis. Universitat Politècnica de Catalunya, 2015. 48 pp.

- **Optimización de estrategias de vuelo de vehículos pilotados por control remoto**. Irene Parada. Bachelor's thesis. Universidad Complutense de Madrid, 2014. 25 pp.